

Research Article

Six Birds Foundations II: Admissibility, Meta-Theory, and the Exact-Six Program

Ioannis Tsiokos¹

1. Automorph Inc., United States

Six Birds Foundations I ^[1] introduced a stratified closure calculus with six primitive closure roles. The present paper extends that calculus into a theorem-level framework in three parts: a primitive role architecture fixing the meanings, levels, local boundaries, and selected lifts of the six typed closure roles $P_1, \dots, P_6$; an admissibility and meta-theory regime in which honest bookkeeping, typed non-collapse, level profiles, forgetting with selected lifts, activation thresholds, instrument-relative visibility, and empirical-bridge admissibility are stated as theorem schemas; and a scoped exact-six theorem chain over the domain FATCD of finite audited typed closure descriptions. Within this domain we prove six role-lower-bound witnesses, an upper-bound classification with corrected role-projection alignments (closing the probability, causality, and agency threats in the covered scope), a decomposition/exhaustion theorem, and an integrated stress certification. The terminal theorem asserts that every $D \in \text{FATCD}$ admits a well-formed decomposition/exhaustion record with exactly six role channels, each in one of five recorded statuses, with residual features classified under an explicit scheme and, under an explicit recognition hypothesis (the three principal threat classes — probability, causality, and agency — closed unconditionally), no seventh irreducible role required in the covered scope. The claim concerns role channels rather than active projections, and the decomposition is existential, not unique. We do not claim unrestricted exact-six, all-emergence coverage, a full six-primitives algebra, global primitive independence, empirical realization, internal instrument-family completeness, or recognition-completeness of the six roles.

Correspondence: papers@team.qeios.com — Qeios will forward to the authors

1. Introduction

Six Birds Foundations I ^[1] introduced a stratified closure-calculus perspective ^[2] on finite theory packages, $\mathcal{T} = (Z, f, \Sigma_f, E, \mathcal{A})$, in which theories arise as fixed points of idempotent operators ^[3] and six primitive closure roles $P_1, \dots, P_6$ emerge from composability combined with limited-interface observation. The present paper extends that calculus into a theorem-level framework. It fixes the primitive role architecture, specifies the admissibility and meta-theory regime under which role talk is licensed, and proves a scoped exact-six terminal result over a precisely delimited domain. The exact-six theorem is the consequence of that stack, not its whole content.

1.1. The three parts

Each of the three parts below supplies structure the other two do not. The *primitive role architecture* (Section 2) fixes the six role meanings of $P1-P6$ at their corrected alignments, records the common behavioral, verification, and structural-categorical level pattern, enumerates the local boundary distinctions under scoped bookkeeping, and states the selected-lift atlas vocabulary. The *admissibility and meta-theory regime* (Section 3) elevates honest bookkeeping from a per-description condition to a meta-theoretic discipline, replaces flat primitive independence by typed non-collapse, and records level, forgetting, threshold, visibility, and empirical-bridge theorem schemas. The *exact-six theorem chain* (Sections 4–8) quantifies the exact-six question precisely over the domain FATCD of finite audited typed closure descriptions and proves the terminal result via lower-bound witnesses, an upper-bound classification with corrected role-projection alignments, a decomposition/exhaustion theorem, and an integrated stress certification.

1.2. The terminal theorem, informally

A finite audited typed closure description D is a finitely presented collection of typed raw records and annotations subject to finiteness, closure, and honest-bookkeeping conditions. The role vocabulary $P1-P6$ enters D not as a constraint on admissibility but as a family of derived projections $\pi_i(D)$, each available when threshold, witness, level, and audit data are attached to raw features. The terminal theorem (Theorem 9.1) states that for every $D \in \text{FATCD}$ there exists a well-formed decomposition/exhaustion record $\text{Dec}(D)$ carrying exactly the six typed closure-role channels $P1, \dots, P6$, each in exactly one of five statuses, with all residual features covered by an explicit classification and, under an explicit recognition hypothesis, no seventh irreducible role required in the covered scope. The claim is about role channels, not active projections, and the decomposition is existential, not unique.

1.3. What the paper does and does not claim

The terminal theorem is scoped exact-six, not unrestricted exact-six. It makes none of the following stronger claims:

- unrestricted exact-six, or any claim about all emergence phenomena;
- that every $D \in \text{FATCD}$ activates all six roles or carries six active nontrivial witnesses;
- that the decomposition/exhaustion record $\text{Dec}(D)$ is unique or canonical;
- a full six-primitives algebra or global primitive independence;
- empirical realization;
- internal instrument-family completeness;
- any statement about broader domains without further work.

These are scope boundaries, collected in Section 11 and enforced locally by the theorem-chain results that would otherwise overread them.

1.4. Organization

After this introduction, Section 2 fixes the primitive role meanings, level trichotomy, boundary distinctions, and selected-lift atlas. Section 3 states the admissibility and meta-theory regime. Section 4 defines FATCD and formulates the exact-six

question over it. Section 5 proves the six-role lower bounds. Section 6 proves the repaired upper-bound classification theorem, closing the probability, causality, and agency threats within the covered scope. Section 7 proves decomposition existence, six-channel exactness, and residual exhaustion, and records the channel-versus-activation discipline and non-uniqueness. Section 8 records the anti-overclaim stress families and the integrated stress certification. Section 9 states and proves the terminal theorem. Sections 11 and 12 record scope boundaries and future directions. Appendix A records development provenance and evidence hygiene, and Appendix B documents the Lean mechanization.

2. Primitive role architecture

This section fixes the meanings of the six typed closure roles $P1$, ..., $P6$, the behavioral, verification, and structural-categorical level pattern they share, the local boundary distinctions between them, and the selected-lift atlas vocabulary. These meanings are fixed for the remainder of the paper and govern every subsequent definition, lower bound, and upper-bound classification.

2.1. Role meanings

Definition 2.1 (Primitive closure roles). The six primitive closure roles are

- $P1$ — update-vs-quotient compatibility / descent obstruction [\[4\]](#), [\[5\]](#);
- $P2$ — macro-specification representability [\[6\]](#);
- $P3$ — enriched route/coherence mismatch [\[7\]](#), [\[8\]](#); $P3$ is not a directionality certificate;
- $P4$ — stage/order/transition/refinement structure [\[9\]](#), [\[10\]](#);
- $P5$ — quotient packaging / quotienting [\[11\]](#);
- $P6$ — audit/bookkeeping/provenance/replay/ checkability [\[12\]](#), [\[13\]](#).

Remark 2.2 (Corrected alignments). Three alignments that Definition 2.1 rules out are used throughout later classification: generic transition data is not $P3$; generic path or derivation data is not $P5$; and generic semantic or interpretation data is not $P1$. These constraints determine which feature clusters can later be interpreted as active projections of each role.

Remark 2.3 (Drive face of $P6$). Arguments about affinities, detailed balance, or honest directionality use the narrower $P6_{drive}$ specialization of $P6$: its affinity-carrying, detailed-balance-breaking face [\[14\]](#). $P6_{drive}$ is a face of $P6$, not a replacement for it. A $P3$ route/coherence-mismatch witness does not establish directionality without a $P6$ -side audit.

2.2. Behavioral, verification, and structural-categorical levels

Each role admits a common three-level pattern.

Definition 2.4 (Level trichotomy). For each role $P1$, ..., $P6$, we distinguish

- the *behavioral level*: status, occurrence, or observed behavior of the role in a description;
- the *verification level*: proof, witness, certificate, or checkable record supporting a behavioral claim [\[15\]](#);
- the *structural-categorical level*: reusable structure supporting the primitive role.

Forgetting proceeds from structural-categorical to verification to behavioral and is lossy. Lifting proceeds from behavioral to verification to structural-categorical only after adding explicitly recorded data, and selected lifts are not unique.

Remark 2.5 (Level content per primitive). The level trichotomy instantiates as follows.

- $P1$: descent success or descent obstruction status; descent proof or obstruction witness; quotient data equipped with a descended map, or an explicit no-descended-map obstruction record.
- $P2$: representability or non-representability status; representability or non-representability certificate; macro-specification layer together with a realization relation.
- $P3$: witness behavior of route disagreement; generated route terms and non-reducibility certificates; categorical route grammar or a reduced categorical universe.
- $P4$: stage status, transition occurrence, or refinement behavior; order/preorder, transition, or refinement certificates; stage/refinement architecture.
- $P5$: packaging or identification behavior; equivalence and quotient-validity certificate; quotient object with a packaging map and universal-property or tower/coarsening data.
- $P6$: audit-relevant event occurrence or record–event association; checkable audit record; audit schema with provenance, replay, and check relations.

The meta-theoretic properties of this trichotomy — lossy forgetting, selected non-unique lifts, and the absence of lower-to-higher determinacy without added data — are formalized in Section 3.

2.3. Local boundary distinctions

With role meanings and level pattern fixed, we turn to how the six roles remain distinct. Table 2 records the pairwise local rules. These are local distinctions under scoped bookkeeping, not a global primitive-independence theorem: two roles may interact legitimately in a single description, but neither collapses into the other. The table is reproduced in landscape orientation on the page immediately preceding the references.

Proposition 2.6 (Boundary distinctions under scoped bookkeeping). *For every pair of roles listed in Table 2, the corresponding “forbidden collapse” is ruled out as a locally admissible identification of those two roles under the repaired alignments of Remark 2.2 together with the honest bookkeeping required by a FATCD.¹*

Proof. We argue row by row. Each row of Table 2 names an ordered pair (P_i, P_j) , a permitted interaction, and a forbidden collapse. The forbidden collapse always has the same logical form: it would license treating witness or structural data carried under the meaning of one role as if it were, by itself, witness or structural data for the other role — that is, it would identify the two roles’ admissible data types. We establish the distinction by a general schema, then prove that the schema provably applies to every row by exhibiting the common form of its forbidden-collapse column, and finally discharge all fifteen rows of Table 2 by a compact enumeration that maps each row’s forbidden collapse to the schema, with the role-specific data types read off Definition 2.1 and the per-primitive level content of Remark 2.5. Three rows are then worked in full as illustrations; they add nothing the enumeration leaves open.

Schema. Fix a row (P_i, P_j) . By Definition 2.1 the two roles have fixed, distinct meanings, and by Remark 2.5 each meaning prescribes a distinct kind of behavioral status, verification witness, and structural object. A forbidden collapse would assert

that data admissible as a P_j -witness (or P_j -structure) is, with no further data, admissible as a P_i -witness (or P_i -structure), or conversely. Suppose such an identification were locally admissible inside a FATCD. Honest bookkeeping requires every active projection to carry a role-aligned feature cluster together with the witness, threshold, level, and audit data appropriate to its role, and requires the channel status of each role to be recorded on its own role-aligned content rather than borrowed from another role. Under the collapse, the bookkeeping for one channel would be discharged by content whose recorded role meaning is the other channel's. The two role meanings being distinct and non-interchangeable, that recorded content does not meet the activation conditions of the channel it would be credited to: the required witness type is simply not present. The honest record must then carry that channel as `below_threshold` or `collapsed_boundary_case`, not as `active_projection`, which contradicts the supposed identification. Hence no such identification is locally admissible. The permitted interaction of the same row is untouched, because it composes P_i -content with P_j -content while keeping each role's bookkeeping on its own content; composition is not identification.

The schema applies to every row. It remains to verify that the hypothesis of the schema — that the forbidden collapse credits one channel on content whose recorded role meaning is the other channel's — holds for each of the fifteen rows, and not merely for chosen representatives. We read this off the structure of the forbidden-collapse column. Every entry in that column of Table 2 has one of two shapes, both instances of the schema. (*Carried form.*) The entry has the form “ X alone is not P_i ” or “ X is not P_i ”, where X is content named by the distinction-rule column as the meaning of the partner role P_j (packaging, stage index, quotient object, audit record, and so on). Such an entry is exactly the assertion the schema forbids: it would credit the P_i channel on the strength of P_j -content. (*Non-implication form.*) The entry has the form “ P_j -status S is not automatically P_i -status S' ” (a representability failure read as a descent obstruction, a route undefinedness read as a route mismatch, a non-representability read as a route mismatch). Such an entry is the symmetric case named in the schema: it would credit the P_i channel's status on the strength of a status belonging to P_j 's meaning. In both shapes the credited channel P_i and the content/status X (or S) carried under P_j 's meaning are fixed by the two columns of the row, and by Definition 2.1 together with Remark 2.5 the witness type required to activate P_i is distinct from, and not supplied by, the P_j -content the collapse offers. The schema therefore forces the honest P_i -status to `below_threshold` or `collapsed_boundary_case`, contradicting the identification. This holds for every row, because every forbidden-collapse entry falls under one of the two shapes just described; the only remaining task is to record, per row, which channel is credited and which witness type is thereby missing.

Enumeration of all fifteen rows. We list each row as P_i/P_j , name the credited channel, the borrowed content, and the distinct witness type the credited channel requires but does not receive, so that the schema applies with no residue. Throughout, “the honest status is `below_threshold`” marks the case in which the credited channel's own witness type is simply not present in the borrowed content, while “the honest status is `collapsed_boundary_case`” marks the case in which a candidate cluster of the partner role's data is offered as the credited channel's witness and fails to instantiate that witness schema.

1. P_1/P_5 , “packaging alone is not P_1 ”: credits P_1 on P_5 packaging content; the required descended-map witness or no-descended-map obstruction record is absent, so the honest P_1 status is `below_threshold`.
2. P_1/P_4 , “stage index alone is not P_1 ”: credits P_1 on P_4 stage-index content; no descent or obstruction witness is thereby present, so the honest P_1 status is `below_threshold`.

3. $P1/P2$, “representability failure is not automatically descent obstruction”: credits $P1$ on a $P2$ non-representability status; a non-representability certificate is not a descended-map obstruction record and does not instantiate the $P1$ obstruction witness, so the honest $P1$ status is `collapsed_boundary_case`.
4. $P1/P3$, “route undefinedness is not route mismatch”: credits $P3$ on a $P1$ route-blocking status; a non-reducibility certificate between *already-defined* routes is required, and an undefined route supplies no such certificate, so the candidate cluster does not instantiate the $P3$ witness schema and the honest $P3$ status is `collapsed_boundary_case`.
5. $P1/P6$, “audit record is not descent obstruction”: credits $P1$ on a $P6$ audit record; an audit record is not a descended-map or obstruction witness, so the honest $P1$ status is `below_threshold`.
6. $P2/P5$, “nontrivial $P2$ is not mere packaging”: credits $P2$ on $P5$ packaging content in the nontrivial case; the realization relation beyond quotient equality that makes the $P2$ content nontrivial is absent from the packaging map, so the honest $P2$ status is `below_threshold`. (In the trivial case the row permits the coincidence and forbids nothing.)
7. $P2/P4$, “stage variation alone is not $P2$ ”: credits $P2$ on $P4$ stage content; a representability certificate with realization relation is required, and stage variation supplies none, so the honest $P2$ status is `below_threshold`.
8. $P2/P3$, “non-representability is not route mismatch”: credits $P3$ on a $P2$ non-representability status; a non-representability record is not a non-reducibility certificate between defined routes and does not instantiate the $P3$ witness schema, so the honest $P3$ status is `collapsed_boundary_case`.
9. $P2/P6$, “audit trail is not representability”: credits $P2$ on a $P6$ audit trail; an audit record is not a representability certificate, so the honest $P2$ status is `below_threshold`.
10. $P3/P4$, “stage-indexed route mismatch is not $P4$ itself”: credits $P4$ on a $P3$ indexed route-disagreement witness; a non-reducibility certificate carries no order, transition, or refinement structure and does not instantiate the $P4$ witness schema, so the honest $P4$ status is `collapsed_boundary_case`.
11. $P3/P5$, “quotient object is not route mismatch”: credits $P3$ on a $P5$ quotient object; a quotient object is not a non-reducibility certificate between defined routes and does not instantiate the $P3$ witness schema, so the honest $P3$ status is `collapsed_boundary_case`.
12. $P3/P6$, “audit record is not route mismatch”: credits $P3$ on a $P6$ audit record; an audit record is not a non-reducibility certificate, so the honest $P3$ status is `below_threshold`.
13. $P4/P5$, “packaging is not stage structure”: credits $P4$ on $P5$ packaging content; a packaging map is not an order/transition/refinement certificate and does not instantiate the $P4$ witness schema, so the honest $P4$ status is `collapsed_boundary_case`.
14. $P4/P6$, “audit record is not staging”: credits $P4$ on a $P6$ audit record of staged histories; an audit record is not an order/transition/refinement certificate, so the honest $P4$ status is `below_threshold`.
15. $P5/P6$, “audit record is not packaging map”: credits $P5$ on a $P6$ audit record of packaging events; an audit record is not a quotient object with a packaging map, so the honest $P5$ status is `below_threshold`.

Clauses (1)–(15) exhaust the rows of Table 2: each names the credited channel, the borrowed partner-role content, and the distinct witness type the credited channel requires and does not receive, and in each the schema forces the honest status to `below_threshold` or `collapsed_boundary_case`. Nothing is left to follow “*mutatis mutandis*”.

We now work three of these rows in full; they illustrate the carried and non-implication shapes and add no obligation beyond clauses (1)–(15).

Row $P1/P5$ (*packaging alone is not $P1$*), clause (1). By Definition 2.1, $P5$ is quotient packaging and $P1$ is update-vs-quotient compatibility, i.e. descent. By Remark 2.5 a $P5$ -witness is a quotient object with a packaging map (and universal-property or tower data), whereas a $P1$ -witness is quotient data *equipped with a descended map*, or else an explicit no-descended-map obstruction record. Possession of the packaging map is precisely the data that says *what* the quotient is; it says nothing about *whether* a selected update descends through it. The collapse “packaging alone is $P1$ ” would credit the $P1$ channel on the strength of the packaging map alone, but no descended-map witness and no obstruction record is thereby present, so by the schema the honest $P1$ status is `below_threshold`, not `active_projection`: the descent question has not been answered. The permitted interaction — $P1$ using $P5$ when checking descent or obstruction — is untouched, since there $P1$ consumes the quotient packaging as the object *over which* a descended map or obstruction is then exhibited, and the descent witness remains the $P1$ channel’s own content.

Row $P3/P4$ (*stage-indexed route mismatch is not $P4$ itself*), clause (10). By Definition 2.1, $P4$ is stage/order/transition/refinement structure and $P3$ is enriched route/coherence mismatch; by Remark 2.2 generic transition data is *not* $P3$ and, symmetrically, a route-mismatch witness is *not* $P4$. A $P4$ -witness is an order/transition/refinement certificate or a stage architecture; a $P3$ -witness is a non-reducibility certificate exhibiting disagreement between already-defined routes. When route mismatch is indexed by stages, the stage index is $P4$ -content and the route disagreement at each index is $P3$ -content; these are recorded as data of different roles. The collapse “stage-indexed route mismatch is $P4$ ” would credit the $P4$ channel on the strength of the indexed route-disagreement witness, but a non-reducibility certificate is not an order/transition/refinement certificate — it carries no stage, order, or refinement structure of its own — so by the schema the $P4$ status it would license is `collapsed_boundary_case`: the candidate cluster does not instantiate the required $P4$ witness schema. The permitted interaction — route systems varying over valid stages — is untouched, since there the $P4$ stage structure is supplied independently and the $P3$ mismatch is merely indexed by it, each channel keeping its own witness.

Row $P2/P5$ (*nontrivial $P2$ is not mere packaging*), clause (6). By Definition 2.1, $P2$ is macro-specification representability and $P5$ is quotient packaging. As the row itself records, when a specification is nothing more than a quotient class the $P2$ content *may* legitimately coincide with $P5$ packaging; the forbidden collapse is scoped to the *nontrivial* case. By Remark 2.5 a $P2$ -witness is a representability certificate for a macro-specification layer together with a realization relation, while a $P5$ -witness is a quotient object with a packaging map. Nontriviality of the $P2$ content means exactly that its representability certificate carries a realization relation not reducible to quotient equality — data beyond the packaging map. The collapse “nontrivial $P2$ is mere packaging” would credit the $P2$ channel on the strength of the packaging map alone; but the realization relation that makes the content nontrivial is by hypothesis absent from the packaging data, so by the schema the honest $P2$ status is `below_threshold`: the representability witness required for activation is not present. This argument is confined to the nontrivial case; in the trivial case the row permits the coincidence and there is nothing to forbid. The permitted interaction — nontrivial $P2$ carried over $P5$ quotient structure — is untouched, since there the realization relation is supplied as the $P2$ channel’s own content layered over the $P5$ packaging.

The enumeration (1)–(15) covers every row of Table 2, and the three worked rows confirm both shapes in detail. The argument is local to each listed pair under the covered scope and the honest bookkeeping a FATCD requires; it establishes pairwise non-collapse, not global independence of the six roles (cf. Remark 2.7). $\square$

Remark 2.7 (Not a global independence theorem). The boundary matrix asserts local distinctions, not global independence. Primitives interact legitimately in a single description: $P1$ descent is audited by $P6$; $P3$ route terms may be indexed by $P4$ stages; nontrivial $P2$ representability is typically carried over $P5$ quotient structure. No claim is made that the six roles are independent generators of a global algebra. The typed-non-collapse counterpart of this observation appears in Section 3.

2.4. Selected-lift atlas and global nonclaims

The remaining architectural data are the selected lifts between levels and the nonclaims that accompany them. Upgrading from behavioral or verification content to the structural-categorical level requires explicitly recorded added data. The resulting lifts are selected, not unique, and the atlas carries structural content that is lost under forgetting.

Definition 2.8 (Selected-lift atlas). The *selected-lift atlas* for the six roles records, for each $P1, \dots, P6$:

- $P1$: proof or witness together with quotient structure, carrying either a descended map or an explicit no-map obstruction;
- $P2$: representability or non-representability certificates together with the macro-specification layer and realization relation;
- $P3$: the selected lift data recorded in the promoted atlas, carrying generated route terms, non-reducibility certificates, and categorical route grammar or a reduced categorical universe;
- $P4$: validity certificates together with structural stage/refinement architecture;
- $P5$: quotient-validity certificates together with the structural quotient data, a packaging map, and universal-property or tower/coarsening data;
- $P6$: record-event or check evidence together with the audit schema and its provenance and replay/check relations.

Each lift is selected and not unique, and forgetting from a higher level loses structural content that is not recoverable without re-adding data.

Remark 2.9 (Global nonclaims for the role architecture). The role architecture preserves the following nonclaims, all carried forward into the admissibility regime of Section 3 and the theorem chain of Sections 5–7:

- no full six-primitives algebra is claimed;
- no global primitive-independence theorem is claimed;
- no empirical realization is claimed;
- no uniqueness of primitive presentations is claimed;
- no uniqueness of selected lifts is claimed;
- no lower-level-to-higher-level determinacy is claimed;
- no instrument-free formulation is claimed;
- direct cross-level comparison remains inadmissible without translation.

3. Admissibility and meta-theory

The admissibility regime is honest bookkeeping, elevated from a per-description condition on $D \in \text{FATCD}$ to a meta-theoretic discipline governing which claims about admissible descriptions enter the theorem chain. Under honest bookkeeping, the flat primitive-independence question is replaced by typed non-collapse, level transitions are controlled by forgetting and selected lifts, thresholds govern activation, instruments govern visibility, and empirical-bridge claims are admitted only as threshold-dependent, level-indexed, instrument-visible assertions. The admissibility regime, recorded in the definitions and propositions of this section, fixes these requirements and preserves the global nonclaims of Remark 2.9.

3.1. Honest bookkeeping as an admissibility regime

The audit condition of Definition 4.6 records enough data for each claim, promotion, translation, lift, forgetting, visibility, and empirical-bridge assertion to be interrogated on its own terms. The following definition stipulates that these record requirements jointly constitute the admissibility regime, in a spirit analogous to the axiomatic specification discipline of Hoare logic [\[16\]](#).

Definition 3.1 (Admissible claim under honest bookkeeping). Honest bookkeeping is the admissibility regime for claims about a description $D \in \text{FATCD}$. A claim that enters the theorem chain is declared *admissible* only if every assertion below carries an explicit record in $\text{Audit}(D)$:

- the claim itself, with its statement, hypotheses, and scope;
- every promotion, with its source level, target level, added data, losses, and preserved nonclaims;
- every translation, with its source context, target context, preserved data, suppressed data, and claim strength;
- every selected lift, marked selected and non-unique;
- every forgetting map, with its record of losses;
- every visibility claim, with both its visible and its suppressed data;
- every empirical-bridge assertion, with its substrate, observable, instrument, level map, threshold-witness relation, and suppressed-structure nonclaims.

Accordingly, unrecorded claims, implicit promotions, silent translations, asserted-but-unselected lifts, lossless forgetting without proof, visibility overreads, and bare empirical-realization assertions are all inadmissible.

Remark 3.2 (Why the regime is necessary). The bookkeeping requirement is a structural check, not a stylistic one. Without explicit records, passive observation can be misread as $P6$, mere packaging as $P1$ descent, and route failure as $P3$ mismatch. The regime prevents those misreadings from entering the theorem chain.

3.2. Typed non-collapse

The regime does not include flat primitive independence — the claim that $P1, \dots, P6$ are independent generators of a global algebra. Its meta-theoretic counterpart here is typed non-collapse.

Proposition 3.3 (Typed non-collapse). Under honest bookkeeping (Definition 3.1) and the repaired alignments of Remark 2.2, for every pair of primitive roles (P_i, P_j) with $i \neq j$, the “forbidden collapse” of the corresponding row in Table 2 is not admissible as an

identification of P_i with P_j ; the “permitted interaction” of the same row is not thereby excluded and may remain admissible when its own bookkeeping conditions are met.²

Proof. Fix a pair (P_i, P_j) with $i \neq j$. The corresponding row of Table 2 records a “forbidden collapse”: an identification asserting that the P_i -content and the P_j -content of a description are the same closure role, so that one projection may be read off the other without further data. To be admissible under Definition 3.1, such an identification would have to enter the theorem chain as a claim with a recorded statement and scope. We show that no such record can be honest.

The two roles have distinct meanings fixed in Definition 2.1: each of $P_1, \dots, P_6$ carries a different closure function (descent obstruction, macro-specification representability, enriched route mismatch, stage/refinement structure, quotient packaging, and audit/provenance/replay, respectively). An identification of P_i with P_j would therefore equate two roles whose meanings differ. Proposition 2.6 of Section 2 establishes that, under the repaired alignments of Remark 2.2 together with the honest bookkeeping required by a FATCD, the forbidden collapse of every pair listed in Table 2 is ruled out as a locally admissible identification of those two roles. In particular, the repaired alignments forbid exactly the substitutions on which a collapse would rely — reading generic transition data as P_3 , generic path or derivation data as P_5 , or generic semantic data as P_1 — so the cross-role record demanded by the collapse cannot be supplied without misalignment. Hence the collapse fails the audit condition of Definition 3.1 and is inadmissible.

It remains to check that ruling out the collapse does not also rule out the row’s “permitted interaction”. The permitted interaction does not identify P_i with P_j ; it records a legitimate joint occurrence of the two distinct roles in a single description — for instance, P_1 descent audited by P_6 , or P_3 route terms indexed by P_4 stages, as in Remark 2.7. Such an interaction carries its own honest-bookkeeping record for each role separately and asserts no identification. Nothing in the argument above touches it: the inadmissibility derived for the collapse rests solely on the equation $P_i = P_j$, which the permitted interaction never asserts. Therefore the permitted interaction remains admissible whenever its own bookkeeping conditions are met. $\square$

Remark 3.4 (Typed non-collapse versus global independence). Typed non-collapse is strictly weaker than global primitive independence. It does not assert that the six roles are logically or categorically independent in a global algebra; it asserts only that the pairwise identifications listed as forbidden collapses in Table 2 are inadmissible under honest bookkeeping, leaving the corresponding typed interactions available.

3.3. Level-profile architecture

Typed non-collapse fixes which identifications are inadmissible; the regime must also record the level at which each admissible claim is made. The level trichotomy of Definition 2.4 is therefore lifted from a role-by-role template to a regime-level structure.

Definition 3.5 (Level profile of an admissible claim). Every admissible claim about a description $D \in \text{FATCD}$ carries a *level profile*: each claim is located on the behavioral, verification, or structural-categorical level, and each cross-level assertion is accompanied by a translation record in the sense of Definition 3.1. Direct comparison of claims at different levels without a translation record is inadmissible.

3.4. Forgetting and selected lifts

Level transitions are controlled by forgetting and selected lifts.

Definition 3.6 (Admissible forgetting and selected lifts). Forgetting from the structural-categorical level to verification, and from verification to behavioral, is admissible and is always lossy in principle: an admissible forgetting carries a record of the losses in $Audit(D)$. Lifting from a lower level to a higher level is admissible only when accompanied by

- the added data required to support the higher-level statement;
- the mark that the lift is selected and non-unique;
- explicit preservation of lower-level invariants and explicit losses where the lift is partial.

Lossless forgetting and unique lifts are inadmissible unless the losslessness or uniqueness is itself proved and recorded.

Remark 3.7 (Lower-to-higher non-determinacy). Definition 3.6 formalizes the non-determinacy of lower-level data: given behavioral or verification content alone, no choice of structural-categorical content is forced by the admissibility regime, and any asserted choice must be marked as a selected lift.

3.5. Activation thresholds

The regime must also state when a projected role is active. The attachment protocol of Definition 4.9 is lifted to a regime-level principle.

Remark 3.8 (Activation thresholds). At the regime level, the threshold-attachment discipline of Definition 4.9 reads as follows. An admissible active derived role projection $\pi_i(D)$ exists only when D carries the threshold, witness, level, collapse, and audit data required by Definition 4.9. Below-threshold clusters carry an explicit `below_threshold` status, and boundary-collapsed clusters carry `collapsed_boundary_case`. Arbitrary enrichment, adding structure solely to force a role projection, is inadmissible unless the additions are themselves honestly bookkept as raw features or annotations. This is a restatement of the attachment protocol at regime level and adds no content beyond Definition 4.9.

3.6. Instrument-relative visibility and suppression

Threshold-governed activation still leaves open what an instrument makes visible or suppresses. Admissibility is therefore instrument-relative.

Definition 3.9 (Admissible visibility record). Every admissible claim carries an explicit *visibility record* comprising its visible content, its suppressed content, the instruments under which the visibility and suppression are asserted, and the resulting instrument-relative claim strength. Overreading a suppressed claim as visible under a different instrument is inadmissible. Partial visibility, in which a claim is visible only up to some coarsening or translation, is admitted when the visibility record itself makes the partiality explicit.

Remark 3.10 (No instrument-free truth). Definition 3.9 does not claim that instruments are free of distortion or that any single instrument is canonical. It asserts only that admissible claims carry an instrument-relative visibility record. Instrument completeness is deferred to Section 12.

3.7. Empirical-bridge admissibility

Any claim crossing from the closure-algebraic side of the framework to a substrate-level observable must pass through the empirical-bridge discipline, related to but narrower than the empirical-adequacy view of theories ^[17].

Definition 3.11 (Admissible empirical-bridge claim). An empirical-bridge claim about $D \in \text{FATCD}$ is admissible only when it carries, in $\text{Audit}(D)$,

- the substrate and observable to which the bridge relates;
- the instrument under which the observable is measured;
- the level map between the closure-algebraic content and the substrate-level content;
- a threshold-witness relation recording when the bridge is crossed;
- the nonclaims suppressed by the bridge record, in particular any structural content of D that is not tracked by the substrate or instrument.

Empirical realization — the claim that D is itself a substrate observation — is stronger than empirical-bridge admissibility and is not asserted by any empirical-bridge claim.

3.8. Nonclaim closure across the admissibility regime

The admissibility regime, recorded in the definitions and propositions of this section — honest bookkeeping (Definition 3.1), typed non-collapse (Proposition 3.3), the level-profile architecture (Definition 3.5), admissible forgetting and selected lifts (Definition 3.6), activation thresholds (Remark 3.8), admissible visibility (Definition 3.9), and admissible empirical-bridge claims (Definition 3.11) — together preserve the global nonclaims of Remark 2.9 and block a corresponding class of overclaims.

Proposition 3.12 (Admissibility-regime nonclaim closure). *Under the admissibility regime, the following remain non-claims:*

- no full six-primitives algebra is claimed;
- no global primitive-independence theorem is claimed;
- no empirical realization is claimed;
- no unique lift theorem is claimed;
- no lower-level-to-higher-level determinacy is claimed;
- no instrument-free truth claim is made;
- direct cross-level comparison remains inadmissible without translation;
- forgetting remains lossy unless losslessness is explicitly proved;
- empirical-bridge claims require threshold data, level maps, instrument visibility records, and suppressed-structure nonclaims.

3

Proof. Each requirement of the regime blocks a specific class of overclaims, and the listed non-claims are exactly the residue.

Honest bookkeeping (Definition 3.1) declares inadmissible every unrecorded claim, implicit promotion, silent translation, asserted-but-unselected lift, lossless forgetting without proof, visibility overread, and bare empirical-realization assertion;

this is the audit condition of Definition 4.6 read as an admissibility filter.

Typed non-collapse (Proposition 3.3) blocks the forbidden primitive collapses of Table 2 and, with them, the flat-independence overclaim that would treat the six roles as independent generators of a global algebra; by the same proposition it leaves the restricted typed boundary interactions available, so the standing non-claim is the global primitive-independence theorem and the full six-primitives algebra, not the absence of interaction.

Level-profile architecture (Definition 3.5) blocks direct cross-level comparison without a translation record, which is exactly the non-claim that cross-level comparison remains inadmissible without translation.

Admissible forgetting and selected lifts (Definition 3.6) block lossless forgetting, unique lifts, and lower-to-higher determinacy: forgetting is lossy unless losslessness is proved and recorded, no lift is admitted as unique, and, by Remark 3.7, no structural-categorical content is forced by behavioral or verification content alone. These are the non-claims that forgetting remains lossy, that no unique-lift theorem is claimed, and that no lower-to-higher determinacy is claimed.

Activation thresholds (Remark 3.8) block arbitrary enrichment and threshold-free nontriviality claims, since an active projection is admissible only with the threshold, witness, level, collapse, and audit data of Definition 4.9.

Instrument-relative visibility (Definition 3.9) blocks instrument-free truth claims and visible-to-suppressed overreads, which yields the non-claim that no instrument-free truth claim is made.

Empirical-bridge admissibility (Definition 3.11) blocks empirical-realization overclaims and any bridge claim lacking threshold data, a level map, an instrument visibility record, or suppressed-structure nonclaims; this yields the non-claim that no empirical realization is claimed and that empirical-bridge claims carry the required data.

Collecting these residues gives exactly the listed non-claims. Those that restate a global nonclaim of Remark 2.9 are thereby carried forward into the regime; the two further items — that forgetting remains lossy absent a recorded proof of losslessness, and that empirical-bridge claims carry their threshold, level-map, visibility, and suppressed-structure data — are additional disciplines the regime imposes at this level. $\square$

4. The FATCD domain and the exact-six question

The scoped exact-six theorem quantifies over a specific class of admissible descriptions; this section defines that class. A *finite audited typed closure description*, abbreviated FATCD, is a finitely presented collection of typed raw records and annotations subject to finiteness, closure, and honest-bookkeeping conditions, sitting in the broader finite-structure tradition of finite model theory [18]. Threshold annotations attach to raw features before any role projection is attempted; the active projections $\pi_1(D), \dots, \pi_6(D)$ for the primitives $P1-P6$ are derived interpretations of raw material, not conditions of domain membership. Candidate raw features not absorbed into such projections are carried into an explicit residual classification, so seventh-role threats are not excluded by definition. The section closes with non-circularity and non-overbreadth schemas and formulates the exact-six question precisely over FATCD: whether every $D \in \text{FATCD}$ admits a decomposition/exhaustion record with exactly six role channels. The theorem chain answering that question appears in Sections 5–7.

The domain is framed in the finite theory-package setting of Foundations I ^[1], where one works with tuples $\mathcal{T} = (Z, f, \Sigma_f, E, \mathcal{A})$ carrying a lens, a completion operator, and an audit functional. FATCD does not identify a description with such a tuple; it extracts from that setting the structural requirements relevant to the scoped exact-six question: finite typed raw data, closure under references, sources, targets, endpoints, and annotations, and honest bookkeeping in the sense of Section 3. Optimization, objective, and richer semantic material are not excluded at the outset; they belong in the domain when finitely represented and honestly recorded as raw or annotated data.

4.1. Raw grammar

We begin with the neutral typed raw grammar from which every element of FATCD is built. The alphabet is fixed in advance of the role vocabulary and follows standard typed-record conventions ^[19].

Definition 4.1 (Raw record). A *raw record* is a finite typed syntactic or relational feature of a description. The admissible record kinds are

- *object* and *state* records;
- *relation* records, *map* records, and *partial-map* records;
- *predicate* records and *comparison* records;
- *path* or *derivation* records;
- *index*, *order*, *transition*, and *refinement* records;
- *trace*, *event*, and *provenance* records;
- *replay* and *check* records.

A raw record is not, by itself, a $P1-P6$ role.

Definition 4.2 (Annotations). In addition to raw records, a description admits *annotations* of six declared kinds:

- level annotations Ann_L ;
- threshold annotations Ann_Θ ;
- lift annotations Ann_{Lift} ;
- forgetting annotations Ann_{Forget} ;
- instrument-visibility annotations Ann_{Vis} ;
- empirical-bridge annotations Ann_{Bridge} .

These annotation kinds belong to the raw grammar; the finiteness and closure conditions on how they attach to declared records are imposed in the next subsection. For a description D , we write $Raw(D)$ for the finite collection of raw records and annotations of D .

Remark 4.3 (Independence of the feature alphabet). The record and annotation kinds in Definitions 4.1 and 4.2 are syntactic or relational entities. No kind is named $P1, \dots, P6$. Consequently, membership in FATCD does not require any derived role projection, does not require an exact-six decomposition/exhaustion record, and does not exclude candidate seventh-role features by definition. The role vocabulary will enter only in §4.5 as a derived interpretation of raw features under threshold and witness data.

4.2. Finiteness, closure, and audit conditions

Three further restrictions turn the raw grammar into a domain of admissible descriptions: finiteness, closure, and audit.

Definition 4.4 (Finite description). A description D is *finite* when

- $Raw(D)$ contains finitely many records and annotations;
- every record or annotation carries finite typed fields;
- every relation, map, path, trace, or annotation is finitely represented;
- every reference in a record points to a declared finite record of D ;
- no infinite semantic, probabilistic, topological, geometric, causal, optimization, or resource structure is admitted unless it is finitely represented and bookkept as a raw feature of D .

Definition 4.5 (Closed description). A finite description D satisfies the *closure condition* $Closed(D)$ when

- every source and target of every relation or map record is a declared record of D ;
- every endpoint of every path or derivation record is a declared record of D ;
- every index, transition, order, or refinement reference is a declared record of D ;
- every trace, event, provenance, replay, or check record references declared records of D ;
- every level, threshold, lift, forgetting, visibility, or empirical-bridge annotation references declared records of D ;
- every claim carried by D has a bookkeeping record.

Definition 4.6 (Audited description). A description D is *audited* when every claim carried by D has an honest bookkeeping record. The collection of such records is denoted $Audit(D) \subseteq Raw(D)$, subject to the following discipline:

- every promotion records its source level, target level, added data, losses, and nonclaims;
- every translation records its source context, target context, preserved data, suppressed data, and claim strength;
- every lift is marked as selected, not unique;
- every forgetting map records its losses;
- every visibility claim records both visible and suppressed data;
- every empirical-bridge claim records substrate, observable, instrument, level map, threshold witness relation, and suppressed-structure nonclaims.

4.3. Finite audited typed closure descriptions

Combining the grammar with the three admissibility conditions yields the domain over which the exact-six question is posed.

Definition 4.7 (Finite audited typed closure description). A *finite audited typed closure description* is a finitely presented, closed, audited collection of typed raw records and annotations, equipped with honest bookkeeping for claims, levels, thresholds, visibility, forgetting, lifts, and empirical-bridge assertions where present. Equivalently, it is a description D whose raw content $Raw(D)$ is drawn from the kinds declared in Definitions 4.1 and 4.2, such that D is finite in the sense of Definition 4.4, satisfies the closure condition $Closed(D)$ of Definition 4.5, and is audited in the sense of Definition 4.6. We write FATCD for the class of all such descriptions, and $D \in \text{FATCD}$ for membership.

Remark 4.8 (No minimum richness). FATCD imposes no lower bound on description size. A description consisting of a single object record together with the bookkeeping required by its claims may be admissible, as may a richer description carrying transitions, paths or derivations, annotations of every kind, and empirical-bridge annotations. The domain is defined by structural discipline, not by the presence of any particular feature kind.

4.4. Threshold attachment

Membership in FATCD is fixed; the next step is to specify how role-theoretic claims may attach to the raw data of an admissible description. Role projections in FATCD enter only through a disciplined attachment of threshold and witness data to raw features. The condition below is the structural precondition for the role vocabulary introduced in the next subsection.

Definition 4.9 (Threshold attachment). A threshold annotation $Ann_\Theta \in Raw(D)$ attaches first to a finite cluster $C \subseteq Raw(D)$ of raw features of D , where C may consist of a single raw feature. No role projection is attempted on C until all of the following are present in $Raw(D)$:

- the raw features of C themselves;
- a threshold annotation Ann_Θ attached to C ;
- a witness schema specifying what in C would support the attempted projection;
- a record of the collapse cases that would invalidate the witness;
- a level annotation Ann_L attached to C ;
- an audit record in $Audit(D)$ honestly bookkeeping the claim.

Remark 4.10 (Pre-projection discipline). The items of Definition 4.9 form a structural precondition, not a process narrative. What matters is that each item is present explicitly in $Raw(D)$ before a role projection is asserted, so that no threshold, witness, collapse case, or level is supplied tacitly by the act of projection itself.

4.5. Derived role projections

The role vocabulary $P1-P6$ now enters FATCD, not as a constraint on membership but as a derived interpretation of raw features conditioned on the attachment data of §4.4.

Definition 4.11 (Derived role projection). A *derived role projection* for a description $D \in \text{FATCD}$ at role index $i \in \{1, \dots, 6\}$ is an active derived interpretation $\pi_i(D)$ of raw features of D as the role P_i , supported by threshold data, witness schema, level assignment, and honest bookkeeping. Concretely, $\pi_i(D)$ is specified by

- the proposed role name, one of $P1, P2, P3, P4, P5, P6$;
- a raw feature cluster $C \subseteq Raw(D)$ on which the role is interpreted;
- the threshold data Θ_i attached to C in the sense of Definition 4.9;
- a witness schema W_i specifying what constitutes a witness for the candidate role under its fixed meaning;
- a level assignment $Level_i$;
- an evidence or certificate type;
- a record of collapse cases avoided or accepted;
- loss and lift annotations for any level crossings involved;

- the explicit nonclaims preserved by the projection.

The supporting data for $\pi_i(D)$ are carried by $Raw(D)$ and certified by $Audit(D)$ in the sense of Definition 4.6.

Remark 4.12 (Fixed role meanings). The role names $P1–P6$ are defined by their meanings in Section 2: $P1$ concerns update-
vs-quotient compatibility and descent obstruction; $P2$ concerns macro-specification representability; $P3$ concerns enriched
route and coherence mismatch, not a directionality certificate; $P4$ concerns stage, order, transition, and refinement
structure; $P5$ concerns quotient packaging; and $P6$ concerns audit, bookkeeping, provenance, replay, and checkability. The
witness schema W_i in Definition 4.11 must be aligned with these meanings; a cluster whose raw features do not support the
meaning of the proposed role contributes no $\pi_i(D)$.

Remark 4.13 (Projection and membership are independent). Whether a description D supports a given derived role
projection $\pi_i(D)$ is not part of the admissibility conditions of Definitions 4.4–4.6. A description $D \in \text{FATCD}$ may carry no
derived role projection, projections for some roles, or projections for all six; its membership in FATCD is unaffected. The
decomposition/exhaustion record of Section 7 accordingly assigns a status to every role channel of D , only some of which
need be active.

4.6. Residual classification scheme

Some raw features remain role-theoretically unsettled at the raw-grammar level. FATCD provides an explicit scheme for
classifying them.

Definition 4.14 (Residual feature; residual classification scheme). Let $D \in \text{FATCD}$. A *residual feature* of D is a raw record or
annotation in $Raw(D)$ whose role-theoretic disposition is not fixed solely by its raw-grammar kind. The *residual classification
scheme* assigns to each residual feature r a single label $Classify(r)$ drawn from the list

- i. already absorbed into a derived $P1–P6$ projection $\pi_i(D)$;
- ii. composite of such projections;
- iii. presentation variant of an already classified feature;
- iv. level shift of an already classified feature;
- v. activation-threshold refinement;
- vi. instrument-visibility artifact;
- vii. empirical-bridge annotation;
- viii. outside-domain structure;
- ix. already-covered bookkeeping or annotation data;
- x. unresolved candidate seventh-role threat;
- xi. genuine seventh irreducible typed role.

The scheme is declarative. At the domain-definition level, FATCD records the available labels but neither proves that every
residual feature lands in categories (i)–(ix) nor proves that categories (x) and (xi) are empty. Those claims are deferred to
Section 6.

Remark 4.15 (Optimization and objective data are not excluded). Raw features expressing optimization, objective, or selection criteria pass through the scheme of Definition 4.14; they are not dismissed by definition. Under the fixed role meanings (Remark 4.12), typical classifications include: a predicate or specification record possibly supporting a $P2$ projection; a map or transition-selection record possibly supporting a $P1$ or $P4$ projection; a route-comparison or route-choice record possibly supporting a $P3$ projection; an audit or check record possibly supporting a $P6$ projection; an empirical-bridge annotation; or an outside-domain structure when the feature requires non-finite or unbookkept semantics. A feature that fits none of these is recorded under category (x).

4.7. Non-circularity and non-overbreadth

Two proposition schemas record the definition-theoretic facts used by the later theorem chain: the domain does not presuppose the answer to the exact-six question, and it is structurally restrictive enough for that question to remain substantive.

Proposition 4.16 (Non-circularity schema for FATCD). *Membership in FATCD is not defined by requiring a decomposition/exhaustion record or by requiring any derived role projection to exist.*⁴

Proof. We inspect the clauses that govern membership, namely those of Definitions 4.1–4.7, and the residual classification scheme of Definition 4.14.

- i. No clause of Definitions 4.1, 4.2, 4.4, 4.5, or 4.6 quantifies over the role index i or requires a role-channel assignment. The record and annotation kinds (Definitions 4.1 and 4.2) are syntactic or relational entities; the finiteness, closure, and audit conditions (Definitions 4.4–4.6) constrain cardinality, field types, reference resolution, and bookkeeping discipline. None of these conditions mentions π_i , a role channel, or a decomposition/exhaustion record. Consequently membership cannot presuppose that any such projection or record exists.
- I. No membership clause excludes a candidate seventh-role feature by definition. A candidate seventh-role feature is, at the raw-grammar level, a record or annotation of one of the declared kinds; nothing in Definitions 4.4–4.6 filters a feature out on the ground that it might resist absorption into a $P1$ – $P6$ projection. Such a feature is therefore admissible whenever it is finite, closed, and audited, and is carried forward to the residual classification scheme rather than barred at the door.
- 1. The residual classification scheme (Definition 4.14) is declarative and carries the categories (x) and (xi)—an unresolved candidate seventh-role threat, and a genuine seventh irreducible typed role—as explicit failure options. At the domain-definition level the scheme records the available labels but does not assert that those two categories are empty, so it cannot smuggle the exact-six answer into the definition of the domain.

Since none of (i)–(iii) ties admissibility to the existence of a role projection or a six-channel record, membership in FATCD is independent of the exact-six question, which is what was claimed. $\square$

Proposition 4.17 (Non-overbreadth schema for FATCD). *FATCD is structurally restricted and does not admit arbitrary syntactic material: arbitrary unaudited syntactic material cannot enter FATCD, and mentioning a role name without the supporting threshold, witness, level, collapse, and bookkeeping data does not by itself produce a derived role projection.*⁵

Proof. Let $D \in \text{FATCD}$. We inspect the admissibility conditions of Definitions 4.4–4.9 in turn.

- i. By Definition 4.4, $\text{Raw}(D)$ contains finitely many records and annotations, each record or annotation carries finite typed fields, and no infinite semantic, probabilistic, topological, geometric, causal, optimization, or resource structure is admitted unless finitely represented and bookkept as a raw feature. Thus unbounded or unrepresented material is excluded at the level of finiteness.
1. By Definition 4.5, every source, target, endpoint, index, transition, order, refinement, trace, event, provenance, replay, check, and annotation reference is bound to a declared record of D . Hence no dangling or externally defined reference can occur, and the syntactic material of D is self-contained.
2. By Definition 4.6, every claim carried by D is bookkept, and every promotion, translation, lift, forgetting, visibility, and empirical-bridge assertion carries the structured audit record required there. Therefore an assertion lacking its honest bookkeeping record cannot be part of an admissible D .
3. By Definition 4.9, a role projection cannot be asserted on a cluster C until the raw features of C , a threshold annotation, a witness schema, a record of collapse cases, a level annotation, and an audit record are all present in $\text{Raw}(D)$. Hence a role name supplied without this attachment data does not, by the act of naming alone, produce a projection.

Combining (i)–(iii), arbitrary unaudited syntactic material cannot enter FATCD , since every admissible D is finite, reference-closed, and honestly bookkept. By (iv), mentioning a role name without the supporting threshold, witness, level, collapse, and bookkeeping data does not by itself produce a derived role projection. This establishes both halves of the claim. $\square$

4.8. The scoped exact-six question over FATCD

We can now state the scoped exact-six question precisely.

Definition 4.18 (Scoped exact-six question). The *scoped exact-six question* over FATCD asks whether, for every $D \in \text{FATCD}$, there exists at least one decomposition/exhaustion record $\text{Dec}(D)$ whose role-channel structure consists of exactly the six typed closure-role channels $P1, P2, P3, P4, P5, P6$, each assigned exactly one of the five statuses later formalized in Section 7, with all residual features of D classified under the scheme of Definition 4.14 without recourse to categories (x) or (xi). No uniqueness is built into this question. The affirmative answer is denoted $\text{ScopedExactSix}(\text{FATCD})$, and is established in Section 9 conditionally on the recognition hypothesis introduced in Section 6 (Definition 6.7), with the three principal threat classes settled unconditionally.

Remark 4.19 (Channels, not active projections). An affirmative answer to the scoped exact-six question asserts the existence of six role channels per description, not six active derived projections per description. The distinction introduced in Remark 4.13 and formalized in Section 7 is preserved throughout.

Remark 4.20 (Relation to the Foundations theorem). Proving $\text{ScopedExactSix}(\text{FATCD})$ is not a re-derivation of the canonical unavoidability result of Foundations I ^[1], and $\text{ScopedExactSix}(\text{FATCD})$ does not follow from that result alone. The Foundations theorem derives the $P1$ – $P6$ vocabulary under minimal hypotheses of composability with limited interface access; the scoped exact-six claim asserts that, within the finite audited typed closure regime of FATCD and under the recognition hypothesis of Section 6, every admissible description admits a six-channel decomposition/exhaustion record

without a seventh irreducible typed role in the covered scope. The two results share vocabulary but address different questions.

5. Six-role lower bounds

The exact-six question of §4.8 has two directions. The lower-bound direction asserts that each of $P_1, \dots, P_6$ is instantiated in FATCD; the upper-bound classification of Section 6 will assert, under a recognition hypothesis, that no covered seventh role is needed. This section establishes the lower-bound direction by constructing one admissible witness for each role. The content is occurrence, not exhaustion: lower bounds alone do not answer the exact-six question, but they rule out the trivialization in which some role has no admissible witness in the domain.

5.1. Lower-bound theorem and common witness requirements

Theorem 5.1 (Six-role lower bounds). *For each role index $i \in \{1, \dots, 6\}$, there exists a description $D_i \in \text{FATCD}$ and a derived active projection $\pi_i(D_i)$ witnessing the role channel P_i in the sense of Definition 4.11. In particular, each of $P_1, P_2, P_3, P_4, P_5, P_6$ is instantiated by at least one admissible description in FATCD. This is an occurrence theorem, not an exhaustion theorem.*⁶

A witness for role P_i consists of the data specified below, following the role-projection discipline of §4.5 and the threshold-attachment discipline of §4.4.

Definition 5.2 (Common witness data for role P_i). *A witness of role P_i in FATCD is a tuple*

$$\left(D_i, C_i, \Theta_i, W_i, Level_i, v_i^{fin}, v_i^{cl}, v_i^{aud}, n_i, \pi_i(D_i) \right)$$

where

- $D_i \in \text{FATCD}$ is the ambient admissible description;
- $C_i \subseteq \text{Raw}(D_i)$ is the raw feature cluster on which the projection is interpreted;
- Θ_i is the threshold annotation attached to C_i ;
- W_i is the witness schema for the role meaning of P_i fixed in Definition 2.1;
- $Level_i$ is the level assignment in the sense of Definition 2.4;
- $v_i^{fin}, v_i^{cl}, v_i^{aud}$ are the finite-description, closure, and audit/bookkeeping verifications for D_i ;
- n_i is the explicit nonclaim record carried by the witness;
- $\pi_i(D_i)$ is the derived role projection of Definition 4.11.

A witness is admissible only if every item in this tuple is present and correctly aligned with the role meaning of P_i .

The proof of Theorem 5.1 proceeds role by role: for each $i \in \{1, \dots, 6\}$ we exhibit an admissible tuple of the form in Definition 5.2. The six constructions are collected in §5.2.

5.2. Per-role witnesses

Each construction below specifies a raw feature cluster, threshold annotation, witness schema, level assignment, derived active projection, and explicit nonclaim record. The finite-description, closure, and audit/bookkeeping verifications are

routine once these data are placed in a description built from the raw grammar of §4.1 and checked against the admissibility conditions of §4.2.

Proposition 5.3 (*P1 lower bound*). *There exists $D_1 \in \text{FATCD}$ with a derived active projection $\pi_1(D_1)$ witnessing P1.*⁷

Proof. Let D_1 carry object records x, y , a relation record $x \sim y$, a selected update map record F , and a comparison record exhibiting $F(x) \sim' F(y)$ in the target relation $\sim'$. Attach a threshold annotation $\Theta_1 = \text{Ann}_\Theta[\text{descent obstruction}]$ to the cluster C_1 comprising these records, and record the witness schema W_1 that detects update-vs-quotient incompatibility between $\sim$ and $\sim'$ under F . Take Level_1 on the verification level, and let $\pi_1(D_1)$ be the derived active projection interpreting C_1 via $(\Theta_1, W_1, \text{Level}_1)$. The finiteness, closure, and audit verifications $v_1^{fin}, v_1^{cl}, v_1^{aud}$ then attach routinely. The nonclaim record n_1 excludes any P5-packaging identification (Remark 2.2) and any claim of an unconditional descended map in the obstruction case. $\square$

Proposition 5.4 (*P2 lower bound*). *There exists $D_2 \in \text{FATCD}$ with a derived active projection $\pi_2(D_2)$ witnessing P2.*⁸

Proof. Let D_2 carry object records together with a predicate record s playing the role of a macro specification, a relation record Real connecting realizers to s , and a comparison record showing that the realizer relation is not reducible to quotient-class equality alone. Attach $\Theta_2 = \text{Ann}_\Theta[\text{nontrivial representability}]$ to the resulting cluster C_2 , and record W_2 as the schema checking that s has or lacks a realizer through Real in a way not reducible to quotient-class equality alone. Take Level_2 on the verification level, and let $\pi_2(D_2)$ be the derived active projection determined by $(C_2, \Theta_2, W_2, \text{Level}_2)$. The finiteness, closure, and audit checks then follow from the construction. The nonclaim n_2 excludes collapse to P5 quotient equality alone and excludes arbitrary relation promotion. $\square$

Proposition 5.5 (*P3 lower bound*). *There exists $D_3 \in \text{FATCD}$ with a derived active projection $\pi_3(D_3)$ witnessing P3.*⁹

Proof. Let D_3 carry object records u, v , two path records ρ_1, ρ_2 both from u to v generated by the enriched route grammar, and a comparison record recording that ρ_1 and ρ_2 are distinguishable by a non-reducibility certificate. Attach $\Theta_3 = \text{Ann}_\Theta[\text{route mismatch}]$ to the cluster $C_3 = \{u, v, \rho_1, \rho_2, \text{comparison}\}$, and record W_3 as the schema certifying that two generated same-source, same-target route records are distinguishable by a non-reducibility or coherence-mismatch certificate. Take Level_3 on the verification level, and let $\pi_3(D_3)$ be the derived active projection determined by $(C_3, \Theta_3, W_3, \text{Level}_3)$. The routine verifications $v_3^{fin}, v_3^{cl}, v_3^{aud}$ then follow. The nonclaim n_3 excludes conflation with P1 route undefinedness and with direct cross-level comparison. $\square$

Proposition 5.6 (*P4 lower bound*). *There exists $D_4 \in \text{FATCD}$ with a derived active projection $\pi_4(D_4)$ witnessing P4.*¹⁰

Proof. Let D_4 carry index records $i_0 < i_1 < i_2$ and transition records t_{01}, t_{12} linking them, together with a refinement record connecting two levels of indexing with explicit transition laws, and closure records for the endpoints and refinement references used by those transitions. Attach $\Theta_4 = \text{Ann}_\Theta[\text{nontrivial stage/refinement}]$ to the cluster C_4 comprising these records, and record W_4 as the schema certifying a non-identity order, transition, refinement, or stage tower. Take Level_4 on the structural-categorical level, and let $\pi_4(D_4)$ be the derived active projection determined by $(C_4, \Theta_4, W_4, \text{Level}_4)$. The routine verifications $v_4^{fin}, v_4^{cl}, v_4^{aud}$ then follow. The nonclaim n_4 excludes bare labels and arbitrary relations or transitions without laws. $\square$

Proposition 5.7 (*P5 lower bound*). *There exists $D_5 \in \text{FATCD}$ with a derived active projection $\pi_5(D_5)$ witnessing P5.*¹¹

Proof. Let D_5 carry distinct object records $x \neq y$, a relation record $x \sim y$, a packaging map record $pack$ with $pack(x) = pack(y)$, and a comparison record providing quotient-validity evidence. Attach $\Theta_5 = Ann_{\Theta}[\text{nontrivial quotient packaging}]$ to the cluster C_5 comprising these records, and record W_5 as the schema certifying that distinct raw objects are related and packaged together by a valid packaging map with quotient-validity evidence. Take $Level_5$ on the verification level, and let $\pi_5(D_5)$ be the derived active projection determined by $(C_5, \Theta_5, W_5, Level_5)$. The routine verifications $v_5^{fin}, v_5^{cl}, v_5^{aud}$ then follow. The nonclaim n_5 excludes identity packaging and weak-distinctness witnesses. $\square$

Proposition 5.8 (*P6 lower bound*). *There exists $D_6 \in \text{FATCD}$ with a derived active projection $\pi_6(D_6)$ witnessing P6.*¹²

Proof. Let D_6 carry event records, a trace record, a provenance record referencing the events, and a replay or check record. Record an audit entry that references the event, tracks provenance and dependency, and supports a replay or check relation. Attach $\Theta_6 = Ann_{\Theta}[\text{nontrivial audit schema}]$ to the cluster C_6 comprising these records, and record W_6 as the schema certifying that the audit record contains event-reference, provenance/dependency, and replay/check structure. Take $Level_6$ on the verification level, and let $\pi_6(D_6)$ be the derived active projection determined by $(C_6, \Theta_6, W_6, Level_6)$. The routine verifications $v_6^{fin}, v_6^{cl}, v_6^{aud}$ then follow. The nonclaim n_6 excludes passive logs, certificate-only content, and partial audit records. $\square$

Propositions 5.3–5.8 supply, for each $i \in \{1, \dots, 6\}$, an admissible witness tuple of the form in Definition 5.2. This proves Theorem 5.1.

Remark 5.9 (What the lower bounds do and do not establish). Theorem 5.1 asserts that each of the six roles is instantiated somewhere in FATCD. It does not prove exact-six, does not rule out candidate seventh-role threats, does not prove decomposition, and does not prove uniqueness of any later decomposition/exhaustion record. It also does not assert that any single description activates all six channels simultaneously. Those questions are handled in Sections 6 and 7.

6. Upper-bound classification: no seventh irreducible role

The lower bounds of Section 5 rule out the trivialization in which some primitive role lacks an admissible witness in FATCD. The upper-bound direction addresses the complementary question: whether some additional, seventh closure role is needed in the covered scope. This section shows that no such role is needed. The argument proceeds in three steps. First, we state the alignment rules that constrain how raw features may project to $P1, \dots, P6$. Second, we close the three high-priority candidate seventh-role threats: probability and uncertainty, residual causality, and residual agency. Third, we record the upper-bound classification theorem that collects these local classifications into a single covered-scope statement.

6.1. Role-projection alignment rules

The role meanings of Definition 2.1 carry alignment rules constraining which raw features may be interpreted as active projections of each role. These rules refine the corrected alignments of Remark 2.2 and form the discipline under which the classification theorem is proved.

Proposition 6.1 (Role-projection alignment). *Under honest bookkeeping and the witness discipline of Definition 5.2, each role projection is constrained as follows.*

- *P1: admissible only when the candidate cluster carries a selected update, a quotient or relation record, and a witness of descent compatibility, descent obstruction, or update-vs-quotient mismatch. Generic semantics, generic predicate meaning, generic selection criteria, and generic agency are not admissible as P1 projections in the absence of such a witness.*
- *P2: admissible only when the candidate carries a macro specification, a realization relation, specification-satisfaction data, and evidence of representability or non-representability. Macro-level predicates or objectives count only insofar as they are treated as content to be represented. Mere quotient equality, arbitrary relations without representability structure, and empirical observations without level and threshold data are not admissible as P2 projections.*
- *P3: admissible only when the candidate carries enriched route grammar, generated same-source, same-target route terms, and route-comparison data with a non-reducibility or coherence-mismatch certificate. Generic transitions, generic causal arrows, generic path records, and route undefinedness arising from P1 obstruction are not admissible as P3 projections unless the route/coherence threshold is explicitly present and the route comparison is specifically of the P3 type.*
- *P4: admissible only when the candidate carries a stage, order, transition, refinement, staged-dynamics, or regime-ordering record with laws beyond bare labels. Bare labels, arbitrary relations without order, transition, or refinement law, quotient packaging, and audit bookkeeping are not admissible as P4 projections.*
- *P5: admissible only when the candidate carries a quotient relation, a packaging map, the identification of distinct micro-elements, a same-package or quotient-object record, and quotient-validity evidence. Generic structural paths, derivations, topologies or geometries merely because they have paths, and causal chains are not admissible as P5 projections unless quotient or packaging data is present.*
- *P6: admissible only when the candidate carries an audit record with event reference, provenance or dependency trace, replay or check structure, and bookkeeping or checkability structure. Passive logs alone, arbitrary metadata, and certificate-only content without audit/provenance/replay/check structure are not admissible as P6 projections.*

13

Proof. The constraints are not fresh stipulations; each is read off the role meaning fixed in Definition 2.1 together with the witness discipline of Definition 5.2. By that definition an admissible projection $\pi_i(D)$ carries a witness schema W_i that is, by requirement, the witness schema for the role meaning of P_i ; a candidate cluster therefore qualifies as an active P_i projection only if it supplies data of the type that the meaning of P_i demands, and only then. We unpack this role by role.

For $P1$, the fixed meaning is update-vs-quotient compatibility and descent obstruction. The witness type for that meaning is exactly a selected update together with a quotient or relation record and a certificate of descent compatibility, descent obstruction, or update-vs-quotient mismatch; absent such a witness the candidate exhibits no descent datum to compare, so it cannot be an active $P1$ projection. In particular generic semantics, generic predicate meaning, generic selection criteria, and generic agency carry no descent witness and are inadmissible, which is the corrected alignment of Remark 2.2.

For $P2$, the fixed meaning is macro-specification representability. Its witness type is a macro specification, a realization relation, specification-satisfaction data, and a representability or non-representability certificate; a macro-level predicate or objective enters only when treated as content to be represented, since only then does it instantiate the representability relation. Mere quotient equality is $P5$ data, an arbitrary relation lacking representability structure supplies no realization

relation, and empirical observations lacking level and threshold data fail the witness discipline; none is an admissible $P2$ projection.

For $P3$, the fixed meaning is enriched route/coherence mismatch, and $P3$ is not a directionality certificate. Its witness type is enriched route grammar, generated same-source same-target route terms, and route-comparison data carrying a non-reducibility or coherence-mismatch certificate. A candidate without such a route comparison has nothing of the $P3$ type to compare. Generic transitions and generic causal arrows are $P4$ data, generic path records are not route-comparison data, and route undefinedness arising from a $P1$ obstruction is a $P1$ phenomenon; none projects to $P3$ unless the route/coherence threshold is explicitly present and the comparison is specifically of the $P3$ type. This is the corrected alignment that generic transition is not $P3$.

For $P4$, the fixed meaning is stage, order, transition, and refinement structure. Its witness type is a stage, order, transition, refinement, staged-dynamics, or regime-ordering record carrying laws beyond bare labels. Bare labels, arbitrary relations without an order, transition, or refinement law, quotient packaging (which is $P5$ data), and audit bookkeeping (which is $P6$ data) supply no such law and are inadmissible as $P4$ projections.

For $P5$, the fixed meaning is quotient packaging. Its witness type is a quotient relation, a packaging map, the identification of distinct micro-elements, a same-package or quotient-object record, and quotient-validity evidence. Generic structural paths, derivations, topologies or geometries that merely carry paths, and causal chains supply no quotient relation or packaging map, so they are inadmissible absent quotient or packaging data; this is the corrected alignment that generic path or derivation is not $P5$.

For $P6$, the fixed meaning is audit, provenance, replay, and checkability. Its witness type is an audit record with event reference, a provenance or dependency trace, a replay or check structure, and bookkeeping or checkability structure. Passive logs without check structure, arbitrary metadata, and certificate-only content lacking audit/provenance/replay/check structure supply no checkable audit datum and are inadmissible as $P6$ projections.

In each case the admissibility constraint is forced by the fixed role meaning through the witness discipline: a projection is admissible only with the witness type that the meaning of the role fixes. $\square$

Remark 6.2 (Three corrected alignments, held fixed). Proposition 6.1 fixes three alignments that are held throughout: generic transition data does not project to $P3$, generic path or derivation data does not project to $P5$, and generic semantic or interpretation data does not project to $P1$. These are carried into every step of the classification theorem below.

6.2. High-priority threat closures

Three classes of candidate seventh-role feature require particular scrutiny: probability and uncertainty, residual causality in the sense of the interventional and counterfactual hierarchy ^[20], and residual agency in the planning sense ^[21]. Each closure is recorded as a proposition. In every case the candidate is classified under the alignment rules of Proposition 6.1 together with the residual scheme of Definition 4.14; no candidate is dismissed by definition. These closures remove the remaining high-priority blockers before the general classification theorem.

Proposition 6.3 (Closure of probability and uncertainty). *Every admissible probability or uncertainty feature of a description $D \in \text{FATCD}$ whose content takes one of the recognized forms below classifies, within the covered scope, as one of*

- a *P2* projection, when the feature is a macro-specification or representability condition;
- a *P4* projection, when the feature is a stochastic transition or regime structure carrying valid stage or transition data;
- a *P6* projection, when the feature is an audited uncertainty-bookkeeping record;
- an activation-threshold refinement, when the uncertainty data modulates when an already admissible role projection activates without supplying a new witness schema;
- an empirical-bridge annotation, when the feature expresses measurement uncertainty bridged to a substrate observable;
- or a composite of the foregoing projections and annotations when the admissible bookkeeping carries several of these aspects at once.

Probability or uncertainty content of any of these recognized forms therefore requires no seventh role; content of none of these forms is governed by the recognition hypothesis (Definition 6.7).¹⁴

Proof. Let r be an admissible probability or uncertainty feature of $D \in \text{FATCD}$, and let it pass through the residual scheme of Definition 4.14 under the alignment rules of Proposition 6.1. By admissibility r carries threshold and bookkeeping data; we case on the content that data fixes.

If the probabilistic content is a macro specification or a representability or non-representability condition on a description, then it supplies the realization-relation and representability witness of *P2*, so by Proposition 6.1 it is a *P2* projection. If instead the content is a stochastic transition, regime, or staged structure carrying a valid stage, order, transition, or refinement law, then it supplies the *P4* witness and is a *P4* projection. If the content is an audited record of uncertainty bookkeeping — event reference, provenance, replay or check structure — it supplies the *P6* witness and is a *P6* projection. If the uncertainty data does not supply any new witness schema but only modulates when an already admissible projection activates, then by the witness discipline it adds no role content and is an activation-threshold refinement. If it expresses measurement uncertainty bridged to a substrate observable, it carries an empirical-bridge annotation rather than a role witness and is so classified. Finally, when the admissible bookkeeping simultaneously carries several of these aspects, r decomposes into the corresponding projections and annotations and is a composite of the foregoing. These are the recognized forms of probability and uncertainty content, and each reduces to the six roles as stated; a probability feature whose content is of none of these forms — for instance a bare normalization or conservation constraint — is not classified here and falls under the recognition hypothesis (Definition 6.7). A feature whose probabilistic content lies outside the covered scope is classified as an outside-domain structure by Definition 4.14, not absorbed by default. The closure is therefore scoped, as stated. $\square$

Proposition 6.4 (Closure of residual causality). *Every admissible residual-causality feature of a description $D \in \text{FATCD}$ whose content takes one of the recognized forms below classifies, within the covered scope, as one of*

- a *P4* projection, as the primary classification when the feature is a directed transition, refinement, or staged-dependence record;
- a *P3* projection, only when the feature is specifically a route/coherence comparison between generated same-source, same-target routes;
- a *P6* projection, when the feature is a provenance or dependency audit;
- a composite of *P4* and *P6*, when directed dependence and provenance audit occur together in the same admissible feature;
- an outside-domain structure, when the feature requires metaphysical causation not amenable to honest bookkeeping.

Classification as $P5$ is inadmissible in the absence of quotient packaging. Residual-causality content of these recognized forms therefore requires no seventh role; content of none of these forms is governed by the recognition hypothesis (Definition 6.7).¹⁵

Proof. Let r be an admissible residual-causality feature of $D \in \text{FATCD}$, causality being taken in the interventional and counterfactual sense of [20], and pass it through the residual scheme of Definition 4.14 under Proposition 6.1. A causal feature presents a directed dependence, so we case on how that directedness is carried.

If r records a directed transition, a refinement, or a staged-dependence law — the typical form of admissible directed dependence in FATCD — it supplies the order/transition/refinement witness of $P4$, so by Proposition 6.1 it is a $P4$ projection; this is the primary classification. The single exception in which directedness routes elsewhere is when r is specifically a route/coherence comparison between generated same-source, same-target routes carrying a non-reducibility or coherence-mismatch certificate; only then does it supply the $P3$ witness and classify as a $P3$ projection. We must not assign a generic causal arrow to $P3$, since by the alignment rules a generic transition is not a route comparison of the $P3$ type. If r is instead a provenance or dependency audit — event reference with a provenance or dependency trace — it supplies the $P6$ witness and is a $P6$ projection. When directed dependence and provenance audit occur together in the same admissible feature, r decomposes into both and is a composite of $P4$ and $P6$. Finally, if r asserts metaphysical causation not amenable to honest bookkeeping — carrying no finite directed-dependence, route, or audit witness — it lies outside the covered scope and is classified as an outside-domain structure.

Classification as a $P5$ projection is not available here: by Proposition 6.1 a $P5$ projection requires a quotient relation and packaging map, and a causal feature carrying no quotient packaging supplies none; the directedness of a causal chain is not packaging data. Hence the $P5$ -inadmissibility caveat holds. The cases above are the recognized forms of residual-causality content; a causality feature whose content is of none of these forms is not classified here and falls under the recognition hypothesis (Definition 6.7). The closure is scoped and makes no scope-free claim. $\square$

Proposition 6.5 (Closure of residual agency). *Every admissible residual-agency feature of a description $D \in \text{FATCD}$ whose content takes one of the recognized forms below classifies, within the covered scope, as one of*

- *a composite of $P2$, $P4$, and $P6$, when the feature decomposes into an objective or specification to be represented, a transition, regime, or staged selection event, and an audited decision-provenance record;*
- *an empirical-bridge annotation, when observed agency serves as substrate evidence for a bridged claim;*
- *an outside-domain structure, when irreducible intentional or teleological content is asserted beyond what admissible bookkeeping reduces to $P2$, $P4$, and $P6$.*

Classification as $P1$ is inadmissible in the absence of an explicit update-vs-quotient descent obstruction. Residual-agency content of these recognized forms therefore requires no seventh role; content of none of these forms is governed by the recognition hypothesis (Definition 6.7).¹⁶

Proof. Let r be an admissible residual-agency feature of $D \in \text{FATCD}$, agency being taken in the planning sense of [21], and pass it through the residual scheme of Definition 4.14 under Proposition 6.1. An agency feature in admissible form presents an objective pursued through selection and recorded for later inspection; we case on whether that content reduces to role witnesses or asserts more.

If r decomposes into an objective or specification to be represented, a transition, regime, or staged selection event, and an audited decision-provenance record, then each component supplies the corresponding witness: the objective-as-content supplies the $P2$ representability witness, the selection event supplies the $P4$ transition or staged-dynamics witness, and the decision-provenance record supplies the $P6$ audit witness. By Proposition 6.1 r is then a composite of $P2$, $P4$, and $P6$. If instead the observed agency serves only as substrate evidence for a bridged claim, it carries an empirical-bridge annotation and is so classified. If r asserts irreducible intentional or teleological content beyond what admissible bookkeeping reduces to $P2$, $P4$, and $P6$, then it carries no finite role witness and lies outside the covered scope, classified as an outside-domain structure.

Classification as a $P1$ projection is not available here merely because agency selects an update: by Proposition 6.1 a $P1$ projection requires a selected update together with a quotient or relation record and an explicit descent-compatibility, descent-obstruction, or update-vs-quotient mismatch witness. A selection event carrying no such descent obstruction supplies no $P1$ witness, so its update content classifies under $P4$ (and, where audited, $P6$) rather than $P1$. This is the stated $P1$ -inadmissibility caveat. The cases above are the recognized forms of residual-agency content; content of none of these forms falls under the recognition hypothesis (Definition 6.7). The closure is scoped and asserts nothing beyond it. $\square$

Remark 6.6 (Scope of the closures). Propositions 6.3–6.5 close the three threats within the covered scope; they do not make scope-free claims. A feature that falls outside the covered scope is classified as an outside-domain structure via Definition 4.14, not absorbed into a primitive role by default.

6.3. The recognition hypothesis and the classification theorem

Definition 6.7 (Recognition hypothesis). A description $D \in \text{FATCD}$ satisfies the *recognition hypothesis* when every residual feature $r \in \text{Residual}(D)$ is *recognized*: the closure content of r either supplies the witness type of one of the six role meanings of Definition 2.1 — so that r , or a finite composite of its constituents, projects to $P1, \dots, P6$ under Proposition 6.1 — or instantiates one of the six annotation kinds of Definition 4.2, or is substrate already accounted for by the audit discipline of FATCD. Equivalently, D carries no finite audited residual feature whose closure content is irreducible to the six role meanings, and in particular none whose role-theoretic content emerges only in an assembled cluster and is not the join of its constituents' role labels.

Remark 6.8 (The recognition hypothesis is assumed, not derived). The recognition hypothesis is an explicit assumption on D , not a consequence of the FATCD definitions. Two kinds of finite, audited, in-domain feature are not known to be recognized: a feature carrying a conservation or balance law over declared records — a finite resource invariant of the kind admitted by Definition 4.4 — whose content supplies none of the six witness types, and a feature whose closure content is an emergent obstruction or holonomy of an assembled cluster, not reducible to the role labels of its constituents. Whether every finite audited feature is recognized — equivalently, whether the six roles are *recognition-complete* for FATCD — is left open and is recorded with the endogenous-instrument meta-limit of Sections 11.3 and 12.1. The three threat closures of Propositions 6.3–6.5 establish the recognition hypothesis unconditionally for the three highest-priority threat classes; the classification theorem below assumes it in general.

The alignment rules of Proposition 6.1, the threat closures of Propositions 6.3–6.5, and the residual scheme of Definition 4.14 combine to yield the upper-bound classification theorem: alignment rules fix the admissible role meanings, threat closures

rule out the three misclassifications that otherwise block the argument, and the theorem, under the recognition hypothesis (Definition 6.7), records the classification codomain for every residual feature.

Theorem 6.9 (Upper-bound classification). *Let $D \in \text{FATCD}$ satisfy the recognition hypothesis (Definition 6.7), and let r be a residual feature of D in the sense of Definition 4.14. Then r classifies as one of*

1. (i) *already absorbed into an admissible $P1$, $P2$, $P3$, $P4$, $P5$, or $P6$ projection under the alignment rules of Proposition 6.1;*
2. (ii) *a composite of such projections;*
3. (iii) *a presentation variant of an already classified feature;*
4. (iv) *a level shift of an already classified feature;*
5. (v) *an activation-threshold refinement;*
6. (vi) *an instrument-visibility artifact;*
7. (vii) *an empirical-bridge annotation;*
8. (viii) *an outside-domain structure;*
9. (ix) *already-covered bookkeeping or annotation data.*

Equivalently, within the covered scope, categories (x) and (xi) of Definition 4.14, namely the unresolved-threat and genuine-seventh-role categories, are empty: no covered candidate seventh-role feature remains as an unresolved threat, and no covered candidate seventh-role feature requires a genuine seventh irreducible typed closure role.¹⁷

Proof. Let r be a residual feature of $D \in \text{FATCD}$ in the covered scope, in the sense of Definition 4.14: a raw record or annotation in $\text{Raw}(D)$ whose role-theoretic disposition is not fixed by its raw-grammar kind alone. By Definition 4.14 the scheme assigns r a single label $\text{Classify}(r)$ drawn from (i)–(xi); we show that, within the covered scope, applying the alignment discipline of Proposition 6.1 together with the three threat closures forces that label into (i)–(ix), leaving (x) and (xi) empty.

Disposition under the alignment discipline. Since r is admissible, it carries the threshold, witness, level, and bookkeeping data of Definition 5.2. We test that data against the six fixed role meanings using the alignment rules of Proposition 6.1, which (by that proposition) admit a projection to P_i exactly when r supplies the witness type that the meaning of P_i fixes. There are two possibilities. Either the witness, threshold, and level data of r align with one of the fixed role meanings, in which case r is absorbed into an admissible $P1$, $P2$, $P3$, $P4$, $P5$, or $P6$ projection and $\text{Classify}(r) = (i)$; or the data of r require several such aligned components at once — for instance an objective to be represented carried over a staged selection with an audit trail — in which case r decomposes into those projections and $\text{Classify}(r) = (ii)$. The corrected alignments of Remark 2.2, carried in through Remark 6.2, are used here to keep these absorptions honest: a generic transition does not slip into (i) as a $P3$ projection, a generic path or derivation does not slip in as a $P5$ projection, and generic semantics does not slip in as a $P1$ projection.

The bookkeeping and presentation categories. If r is not a fresh role witness but a re-presentation of an already classified feature, it is a presentation variant, $\text{Classify}(r) = (iii)$; if it restates an already classified feature at a different level of the trichotomy of Definition 2.4, it is a level shift, (iv). If its uncertainty or threshold data only modulates when an already admissible projection activates, without supplying a new witness schema, it is an activation-threshold refinement, (v). If it is visible or suppressed only relative to a declared instrument, it is an instrument-visibility artifact, (vi). If it bridges a substrate

observable to an admissible claim, it is an empirical-bridge annotation, (vii). If it requires non-finite or unbookkept semantics not amenable to honest bookkeeping, it falls outside the covered scope and is an outside-domain structure, (viii). If it is bookkeeping or annotation data already accounted for by the audit discipline of FATCD, it is already-covered bookkeeping or annotation data, (ix). These are the dispositions available to a recognized feature.

The three high-priority threats. The classes most apt to demand a seventh role are recognized unconditionally. A probability or uncertainty feature of one of the recognized forms is, by Proposition 6.3, a $P2$, $P4$, or $P6$ projection, an activation-threshold refinement, an empirical-bridge annotation, or a composite of these — all in (i)–(ix). A residual-causality feature of a recognized form is, by Proposition 6.4, a $P4$, $P3$, or $P6$ projection, a $P4/P6$ composite, or an outside-domain structure, with $P5$ excluded for want of quotient packaging. A residual-agency feature of a recognized form is, by Proposition 6.5, a $P2/P4/P6$ composite, an empirical-bridge annotation, or an outside-domain structure, with $P1$ excluded for want of a descent obstruction. The three closures thus establish the recognition hypothesis for the three highest-priority threat classes without assuming it.

Conclusion. Under the recognition hypothesis (Definition 6.7), every residual feature r of D is recognized and hence receives a label in (i)–(ix): its content supplies a role witness (cases (i)–(ii)), instantiates an annotation kind (cases (iii)–(vii)), or is already-covered substrate (case (ix)), and only non-finite or unbookkept content is sent to the outside-domain category (viii). Since every residual feature receives a label in (i)–(ix) by the foregoing, category (x), the unresolved-threat category, is empty; and since the recognition hypothesis assumes that no finite audited in-domain feature has content irreducible to the six roles, category (xi), the genuine-seventh-role category, is empty. The emptiness of (xi) is the content of the recognition hypothesis, not a consequence of the FATCD definitions: by Remark 6.8 a finite audited conservation law or an emergent composite obstruction would populate (xi), and whether any such feature exists is the open recognition-completeness question. The theorem holds exactly to the strength of its hypothesis. $\square$

Remark 6.10 (Scope of the theorem). Theorem 6.9 is conditional on the recognition hypothesis (Definition 6.7) and scoped to FATCD. It does not derive that no finite audited feature requires a seventh role; it assumes this, and discharges the assumption unconditionally only for the three threat classes of Propositions 6.3–6.5. It does not claim that no seventh role can arise outside FATCD, that the alignment rules are exhaustive for richer domains, or that the recognition hypothesis holds in general. Nor does it prove exact-six on its own: the decomposition/exhaustion theorem of Section 7 remains necessary. The open recognition-completeness question and broader questions are deferred to Sections 11.3 and 12.

7. Decomposition and exhaustion

The lower-bound theorem of Section 5 and the upper-bound classification of Section 6 reduce the scoped exact-six question to a decomposition-existence problem: given $D \in \text{FATCD}$, does D admit a well-formed record combining the six role channels, their statuses, and the classification of any raw features not absorbed into active projections? This section answers affirmatively and records the channel-versus-activation discipline that the answer requires. The lower bounds show that each of the six channels is materially realizable in the domain; the upper-bound classification shows that, under the recognition hypothesis, no covered seventh role channel is needed.

7.1. Decomposition records and channel statuses

Definition 7.1 (Channel status). For each $i \in \{1, \dots, 6\}$, the *role channel* $Chan_i(D)$ of a description $D \in \text{FATCD}$ is the P_i -labeled slot in the six-channel interface associated to D , regardless of whether that slot is active. A *channel status* for $Chan_i(D)$ is exactly one of the five values

- *active_projection* — an active derived role projection $\pi_i(D)$ in the sense of Definition 4.11 exists, with the full attachment data of Definition 4.9;
- *below_threshold* — candidate raw features are present but the threshold data required by Definition 4.9 is absent;
- *collapsed_boundary_case* — the candidate cluster is a boundary case that collapses under the witness schema W_i ;
- *absent_with_record* — no relevant raw features for role P_i are present in $Raw(D)$, and the absence is itself recorded;
- *inapplicable_outside_scope* — the role- P_i content of D lies outside the covered scope.

We write $Status_i(D)$ for the channel status of $Chan_i(D)$.

Definition 7.2 (Decomposition/exhaustion record). A *decomposition/exhaustion record* for $D \in \text{FATCD}$ is a tuple

$$Dec(D) = \left(Raw(D), (Chan_i(D), Status_i(D))_{i=1}^6, Residual(D), Classify, \mathcal{L}, \mathcal{V}, Audit(D), \mathcal{N} \right)$$

where

- the six role channels together with their statuses form the six-channel status sequence. If $Status_i(D) = \text{active_projection}$, the channel carries an active projection $\pi_i(D)$ with its full attachment data. If $Status_i(D)$ is inactive, the channel carries the corresponding below-threshold, collapsed-boundary, absent-with-record, or outside-scope status record;
- $Residual(D) \subseteq Raw(D)$ is the residual feature set (the raw records and annotations not already accounted for by an active projection or an inactive status record), and $Classify$ assigns to each $r \in Residual(D)$ a label in the scheme of Definition 4.14;
- $\mathcal{L}$ collects the loss and lift records for any level crossings involved in projections or classifications;
- $\mathcal{V}$ collects instrument-visibility and empirical-bridge annotations where these are asserted;
- $Audit(D)$ is the honest bookkeeping subrecord;
- $\mathcal{N}$ is the explicit nonclaim record.

A decomposition/exhaustion record is *well-formed* when each $Status_i(D)$ is drawn from the five values of Definition 7.1, each active $\pi_i(D)$ is admissible in the sense of Definitions 4.11 and 4.9, each inactive status record is honestly supported by the corresponding raw-feature situation, and every raw feature is accounted for at the channel/exhaustion level either by an active projection, by a status record for its channel, or by a residual classification. The records in $\mathcal{L}$ and $\mathcal{V}$ are auxiliary annotations on this coverage, not additional role channels.

7.2. Existence, six-channel exactness, and residual exhaustion

Theorem 7.3 (Decomposition existence). *Every description $D \in \text{FATCD}$ admits a well-formed decomposition/exhaustion record $Dec(D)$ in the sense of Definition 7.2.*¹⁸

Proof. For each $i \in \{1, \dots, 6\}$, inspect $Raw(D)$ for a cluster aligned with the meaning of P_i under Proposition 6.1. If such a cluster carries the threshold, witness, level, and audit data of Definition 4.9, record $Status_i(D) = \text{active_projection}$ and an active projection $\pi_i(D)$ in the sense of Definition 4.11; otherwise record one of the four inactive statuses, determined by which attachment data is missing, whether the cluster is a boundary-case collapse, whether no relevant cluster exists, or whether the role lies outside the covered scope. Six status entries result, one per role. The raw features absorbed by these channel records are removed from the residual stage; the remainder pass to $Residual(D)$. By Definition 4.14 the scheme assigns each residual feature a single label drawn from (i)–(xi), so $Classify$ is total on $Residual(D)$ and the record is well-formed irrespective of which labels occur; under the recognition hypothesis (Definition 6.7), Theorem 6.9 further places every label in (ii)–(ix), with (x) and (xi) empty. The records $\mathcal{L}$ and $\mathcal{V}$ are read from $Raw(D)$ and $Audit(D)$; the bookkeeping subrecord is $Audit(D)$; and $\mathcal{N}$ records the nonclaims carried forward from the witness discipline. The resulting tuple is a decomposition/exhaustion record, well-formed by construction. Non-vacuity of the six channels follows from Theorem 5.1, which supplies, for each channel, a description in which that channel is active. $\square$

Theorem 7.4 (Six-channel exactness and residual exhaustion). *For every $D \in \text{FATCD}$ and every well-formed $Dec(D)$ in the sense of Definition 7.2:*

1. (i) *the channel sequence $(Chan_i(D))_{i=1}^6$ has length exactly six, one per primitive role $P1, \dots, P6$;*
2. (ii) *every raw feature of D is accounted for at the channel/exhaustion level by an active projection $\pi_i(D)$, a channel-status record for an inactive channel, or a residual classification under the scheme of Definition 4.14; the records in $\mathcal{L}$ and $\mathcal{V}$ are auxiliary annotations on this coverage;*
3. (iii) *under the recognition hypothesis (Definition 6.7) and within the covered scope, no raw feature is classified as an unresolved seventh-role threat or as a genuine seventh irreducible typed closure role (Theorem 6.9); the three principal threat classes are closed unconditionally.*

19

Proof. Exactness of the channel sequence is built into the definition of $Dec(D)$; there are six role channels by Definition 7.2. Coverage of raw features is the well-formedness condition of Definition 7.2 together with totality of $Classify$ on $Residual(D)$, which holds by Definition 4.14. Emptiness of categories (x) and (xi), under the recognition hypothesis (Definition 6.7) and within the covered scope, is Theorem 6.9 directly. $\square$

7.3. Channel-versus-activation discipline

Theorem 7.4 asserts six channels in every decomposition/exhaustion record. It does not assert six active projections in every record, nor does it assert that a well-formed record is unique. The distinction between channel presence and channel activation prevents six-channel exactness from being read as an all-six-active claim.

Proposition 7.5 (Channel-versus-activation discipline). *Let $D \in \text{FATCD}$ and let $Dec(D)$ be a well-formed decomposition/exhaustion record. The following are distinct statements:*

- i. *Six channels: the channel sequence $(Chan_i(D))_{i=1}^6$ has length six with the fixed role names $P1, \dots, P6$.*
- ii. *Six active projections: for each i , $Status_i(D) = \text{active_projection}$ and the channel carries an active $\pi_i(D)$.*

Theorem 7.4 establishes (i) for every $D \in \text{FATCD}$. It does not establish (ii). In particular, an admissible $D \in \text{FATCD}$ may carry any number of active projections between zero and six; inactive channels are recorded with status *below_threshold*, *collapsed_boundary_case*, *absent_with_record*, or *inapplicable_outside_scope* and are not active witnesses for the corresponding roles.²⁰

Proposition 7.6 (Decomposition records are not unique). A description $D \in \text{FATCD}$ may admit more than one well-formed decomposition/exhaustion record in the sense of Definition 7.2. Distinct well-formed records may differ in level assignments, choice of witness schema, residual classification labels for features admitting more than one admissible category under the alignment rules, or placement of loss/lift and visibility records. Theorems 7.3 and 7.4 assert existence, not uniqueness.²¹

Remark 7.7 (What the decomposition theorems do and do not prove). Theorems 7.3 and 7.4 establish the structural side of the scoped exact-six claim: every $D \in \text{FATCD}$ admits a six-channel decomposition/exhaustion record, with residual exhaustion in the covered scope under the recognition hypothesis (Definition 6.7). Together with the lower-bound theorem (Section 5) and the upper-bound classification (Section 6), they prepare the scoped exact-six theorem but do not replace the integrated-stress check that follows. They do not assert that every description activates all six roles, that the decomposition is canonical or unique, or that the scoped exact-six program extends beyond FATCD; those remain nonclaims.

8. Integrated stress

With the theorem chain of Sections 5, 6, and 7 in place, we record six stress families that act as consistency checks against overclaims. These are not proofs of the scoped exact-six target; they verify that the theorem chain does not silently admit overclaims. Each family is a guardrail that restates a nonclaim already secured by an earlier result, so each is recorded as a remark. Their joint satisfaction is the last requirement before the terminal statement in Section 9.

8.1. Dependency stress

Remark 8.1 (Dependency stress). The scoped exact-six target is dependency-coherent only when the theorem chain is assembled in the order

$\text{FATCD domain} \rightarrow \text{lower bounds} \rightarrow \text{upper-bound classification} \rightarrow \text{decomposition/exhaustion}.$

Accordingly, the following are inadmissible:

- asserting the target without the FATCD domain of Section 4;
- asserting the target without the lower-bound theorem of Section 5;
- asserting the target without the upper-bound classification theorem of Section 6;
- asserting the target without the decomposition and exhaustion theorems of Section 7;
- promoting the assembled target before the integrated stress check of Proposition 8.7.

8.2. Circularity stress

Remark 8.2 (Circularity stress). Three circularity patterns are inadmissible under the theorem chain:

- redefining FATCD so that it already contains exactly six role slots by construction;

- declaring the absence of a seventh role by definition rather than by classification under Definition 4.14;
- justifying decomposition existence by assuming scoped exact-six.

The non-circularity schema of Proposition 4.16 and the constructive proof of Theorem 7.3 block each pattern.

8.3. Channel-versus-activation stress

Remark 8.3 (Channel-versus-activation stress). The following conflations are inadmissible:

- requiring that every $D \in \text{FATCD}$ activates all six role channels;
- treating an inactive channel of status `below_threshold` as an active witness for its role;
- treating an inactive channel of status `collapsed_boundary_case` as an active witness;
- treating an inactive channel of status `absent_with_record` as an active witness.

Proposition 7.5 enforces each distinction. The same inactive-status guardrail also applies to `inapplicable_outside_scope`.

8.4. Role-alignment stress

Remark 8.4 (Role-alignment stress). The three forbidden assignments of Remark 6.2 are inadmissible: generic transition data does not project to $P3$, generic path or derivation data does not project to $P5$, and generic semantic or interpretation data does not project to $P1$. The alignment rules of Proposition 6.1 enforce each prohibition.

8.5. Upper-bound residual stress

Remark 8.5 (Upper-bound residual stress). Under Theorem 6.9 and within the covered scope, the following regressions are inadmissible:

- probability or uncertainty features reopening as unresolved threats;
- residual-causality features reopening as unresolved threats;
- residual-agency features reopening as unresolved threats;
- any covered residual feature remaining unclassified under Definition 4.14;
- any covered residual feature requiring a genuine seventh irreducible typed closure role.

Propositions 6.3–6.5 settle the three threat closures, and Theorem 6.9, under the recognition hypothesis, delivers the emptiness of categories (x) and (xi).

8.6. Overclaim stress

Remark 8.6 (Overclaim stress). The following restatements of the scoped exact-six target are inadmissible:

- restatement as a full six-primitives algebra;
- restatement as a global primitive-independence theorem;
- restatement as unique or canonical decomposition;
- restatement as empirical realization;
- restatement as a claim about all emergence phenomena, or any domain broader than FATCD.

Remark 2.9, Proposition 3.12, and Proposition 7.6 carry the corresponding nonclaims forward through the theorem chain. In particular, nonclaim closure survives the exact-six assembly.

8.7. The integrated stress certification

Proposition 8.7 (Integrated stress certification). *The theorem chain of Sections 4, 5, 6, and 7 satisfies each of the six stress families recorded in the stress remarks of this section. Dependency, circularity, channel-versus-activation, role-alignment, upper-bound residual, and overclaim regressions are all blocked, and the assembled target is ready for the terminal statement of Section 9.*²²

Proof. We verify that each of the six guardrails is discharged by the earlier result it rests on, so that none of the listed regressions can arise in the assembled chain.

Dependency. The components named in Remark 8.1 are present and ordered: the FATCD domain is fixed in Section 4, the lower-bound theorem in Section 5 (Theorem 5.1), the upper-bound classification in Section 6 (Theorem 6.9), and the decomposition and exhaustion theorems in Section 7 (Theorems 7.3 and 7.4). Each later statement is stated over the output of the previous one, so the target cannot be asserted while skipping any link, and no link circularly consumes its own output.

Circularity. The three circularity patterns of Remark 8.2 are blocked by the non-circularity schema of Proposition 4.16, which defines FATCD without reference to a six-slot count, and by the constructive proof of Theorem 7.3, which produces a decomposition without assuming scoped exact-six. The absence of a seventh role is decided by classification under Definition 4.14 rather than by fiat.

Channel versus activation. The conflation of Remark 8.3 are excluded by Proposition 7.5 together with the five-status taxonomy of Definition 7.1: a channel carrying status `below_threshold`, `collapsed_boundary_case`, `absent_with_record`, or `inapplicable_outside_scope` is an inactive channel and is never read as an active projection, so no description is required to activate all six channels.

Role alignment. The three forbidden assignments of Remark 8.4 are ruled out by the alignment rules of Proposition 6.1, which refuse generic transition data as *P3*, generic path or derivation data as *P5*, and generic semantic or interpretation data as *P1*, in accordance with the corrected alignments of Remark 6.2.

Upper-bound residual. The regressions of Remark 8.5 are closed within the covered scope by Propositions 6.3–6.5, which settle the probability, causality, and agency threat closures, and by Theorem 6.9, which — under the recognition hypothesis (Definition 6.7) — classifies every covered residual feature and delivers the emptiness of categories (x) and (xi); hence, under that hypothesis, no covered residual reopens as an unresolved threat or forces a genuine seventh role, while the three principal threat classes are closed unconditionally.

Overclaim. The inadmissible restatements of Remark 8.6 are blocked because the nonclaims of Remark 2.9 are carried forward by Proposition 3.12, and non-uniqueness of the decomposition is recorded by Proposition 7.6; thus nonclaim closure survives the exact-six assembly, and no full-algebra, global-independence, unique-decomposition, empirical-realization, or broadened-domain reading is admitted.

Each of the six guardrails is therefore satisfied. Dependencies cohere, the channel-versus-activation distinction holds, nonclaim closure holds, and the chain admits none of the listed regressions, so the assembled target is eligible for the terminal statement of Section 9. □

9. The scoped exact-six theorem

The theorem chain of Sections 4–7 and the integrated stress certification of Section 8 together permit the terminal statement. The scoped exact-six theorem combines the definition of FATCD with the lower bounds, the upper-bound classification, and the decomposition/exhaustion theorems into a single quantified claim over the domain.

Theorem 9.1 (Scoped exact-six). *Let $D \in \text{FATCD}$ satisfy the recognition hypothesis (Definition 6.7). Then D admits a well-formed decomposition/exhaustion record $\text{Dec}(D)$ with exactly six typed closure-role channels:*

- *P1: update-vs-quotient compatibility / descent obstruction;*
- *P2: macro-specification representability;*
- *P3: enriched route/coherence mismatch;*
- *P4: stage/order/transition/refinement structure;*
- *P5: quotient packaging / quotienting;*
- *P6: audit/bookkeeping/provenance/replay/checkability.*

Each channel carries exactly one recorded status from $\{\text{active_projection}, \text{below_threshold}, \text{collapsed_boundary_case}, \text{absent_with_record}, \text{inapplicable_outside_scope}\}$.

Every active projection is supported by threshold data, a witness schema, a level assignment, and an honest bookkeeping record. Every residual raw feature of D is classified, under the upper-bound classification of Section 6, as one of:

- *composite of existing role channels;*
- *presentation variant;*
- *level shift;*
- *activation-threshold refinement;*
- *instrument-visibility artifact;*
- *empirical-bridge annotation;*
- *outside-domain structure;*
- *already-covered bookkeeping or annotation data.*

Under that hypothesis, no residual feature remains as an unresolved candidate seventh role or genuine seventh irreducible typed closure role; for the three principal threat classes — probability, causality, and agency — this holds unconditionally (Propositions 6.3–6.5). We denote this conclusion by

$$\text{ScopedExactSix}(\text{FATCD}).$$

23

Proof. Fix $D \in \text{FATCD}$ satisfying the recognition hypothesis. By Theorem 7.3, D admits a well-formed decomposition/exhaustion record $\text{Dec}(D)$ in the sense of Definition 7.2. By Theorem 7.4, the channel sequence of $\text{Dec}(D)$ has length six, each channel carries one of the five statuses of Definition 7.1, and every raw feature of D is accounted for by an active projection, by a status record, by a residual classification, or by an auxiliary loss/lift or visibility annotation. By Theorem 6.9, which applies because D satisfies the recognition hypothesis, residual features classify under categories (ii)–

(ix) of Definition 4.14, and categories (x) and (xi) are empty. Theorem 5.1 supplies, for each channel, a $D' \in \text{FATCD}$ at which that channel is active. The integrated stress certification of Proposition 8.7 discharges the dependency, circularity, channel-versus-activation, role-alignment, upper-bound residual, and overclaim families. Hence $\text{ScopedExactSix}(\text{FATCD})$. $\square$

9.1. What scoped exact-six means and does not mean

Theorem 9.1 has a reading that should be preserved throughout what follows.

- *Six channels, not six active witnesses.* The theorem asserts six role channels in $\text{Dec}(D)$ for every $D \in \text{FATCD}$. It does not assert that every D activates all six roles, that every D carries six active projections $\pi_1(D), \dots, \pi_6(D)$, or that a single D realizes all six witness schemas of Section 5.
- *Existential, not unique, decomposition.* The theorem asserts existence of $\text{Dec}(D)$; it does not assert uniqueness or canonicity. Distinct well-formed records of the same D may exist (Proposition 7.6).
- *Conditional residual exhaustion.* The residual exhaustion clause is conditional on the recognition hypothesis (Definition 6.7) and restricted to the covered scope: the three principal threat classes are closed unconditionally, while whether every finite audited feature is recognized is the open recognition-completeness question (Remark 6.8). Features requiring non-finite or unbookkept semantics are classified as outside-domain structures under category (viii) of Definition 4.14, not absorbed into primitive roles.
- *No broader claim.* Theorem 9.1 is a statement about FATCD. It does not assert a full six-primitives algebra, global primitive independence, empirical realization, instrument-family completeness, or coverage of all emergence phenomena. See Section 11.

10. Discussion

The terminal theorem is stated at scoped strength over FATCD. This section discusses what the statement earns and what it leaves unstated, how the three parts of Sections 2, 3, and 4–8 combine into one theorem stack, and how the result sits next to earlier Six Birds literature.

10.1. The shape of the scoped claim

Each of the three scope features — domain restriction, existence rather than uniqueness, residual exhaustion within a covered scope — is structural rather than cosmetic. An unrestricted exact-six theorem would require a domain-free argument for which the present paper does not claim to have the instruments. A uniqueness or canonicity claim would require a theory of translation records identifying distinct well-formed decompositions, which lies outside the present admissibility regime. Scope-free residual exhaustion would require scope-free closure of the three high-priority threats addressed in Section 6. Stating the theorem at the scope the proof supports keeps the endogenous-instrument meta-limit (Section 11.3) separate from the current claim instead of folded into it. The posture aligns with the broader view that scientific laws form a patchwork of domain-local regularities rather than a single global theory.^[22]

10.2. The three parts as one theorem stack

The three parts introduced in Section 1.1 play complementary roles. The primitive role architecture (Section 2) fixes the vocabulary: six role meanings, the behavioral, verification, and structural-categorical level pattern, and the local boundary distinctions. This vocabulary alone does not certify any claim, since it does not specify which role talk is admissible. The admissibility and meta-theory regime (Section 3) supplies that certification through honest bookkeeping, typed non-collapse, the forgetting and selected-lift architecture, activation thresholds, instrument-relative visibility, and empirical-bridge admissibility. Admissibility alone, however, yields no terminal claim; it only regulates what may be asserted. The exact-six theorem chain (Sections 4–8) supplies the quantification: the FATCD domain, the six lower-bound witnesses, the upper-bound classification, the decomposition/exhaustion theorem, and the integrated stress certification together deliver Theorem 9.1. Vocabulary without admissibility would be undisciplined; admissibility without a theorem chain would be inconclusive; a theorem chain without the first two would lack a subject.

10.3. Channel versus activation

The channel-versus-activation discipline of Proposition 7.5, enforced throughout Sections 4–9, separates two statements that are easily conflated. The structural statement is that every admissible description carries six typed closure-role slots. The activation statement is that every admissible description realizes all six slots with active projections. The activation statement is too strong: it fails on descriptions that carry no route grammar and therefore have no active $P3$ projection, and on descriptions whose raw material carries no $P6$ -aligned event/provenance/replay/check witness cluster and therefore have no active $P6$ projection. The structural statement, which is what we prove, is compatible with any activation profile. Inactive channels carry one of the statuses `below_threshold`, `collapsed_boundary_case`, `absent_with_record`, or `inapplicable_outside_scope`, describing the way a role fails to activate. Separating the two statements is what lets the terminal theorem hold of every $D \in \text{FATCD}$ without requiring activation of all six roles in every description; it also lets the lower-bound theorem of Section 5 be stated existentially over the domain rather than uniformly. An informal reading that elides the distinction silently recovers the stronger activation claim, which we do not prove.

10.4. Instrument-relative methodology

A methodological choice underlying Section 3 deserves explicit statement as a contribution. Instruments are treated as mathematical probes rather than as neutral tools, in the spirit of ontic structural realism [23] and the sheaf-theoretic treatment of contextuality [24]. The technical form is Definition 3.9; the associated nonclaim is that no result here rests on an instrument-free truth.

The motivation is reflexive. The subject matter — finite audited typed closure descriptions and the six typed closure roles that project them — is the same descriptonal machinery from which any instrument for detecting closure structure must itself be built. Operator-presentation, quotient/descent, coherence/route, constructive/proof-theoretic, and information/regime instruments are not vantage points outside the domain: each is a closure structure probing another closure structure. A claim asserted under one such instrument while silently suppressing what another would have revealed would privilege that instrument without bookkeeping for the privilege — the overread that Definition 3.9 blocks.

The consequence is epistemic. Every admissible claim carries an explicit visibility record: its visible content, its suppressed content, the instruments under which visibility and suppression are asserted, and the resulting instrument-relative claim strength (Definition 3.9). Triangulating a role projection across several instruments is a disciplined way to strengthen an assertion; reading a projection off a single instrument and declaring its content canonical is not admissible. The discipline is one of epistemic hygiene, in the spirit of Bohr’s complementarity doctrine ^[25]: we do not claim that instruments are incompatible in principle, only that they are not interchangeable without explicit visibility bookkeeping and, where contexts shift, explicit translation records (Definitions 3.9 and 3.5).

This posture also explains why the endogenous-instrument meta-limit (Sections 11.3 and 12.1) is deferred to future work. Taking instrument-relativity seriously from the outset lets the meta-limit be stated as its own question — whether the accepted internal instrument family can certify its own soundness and completeness — a question made statable, but not settled, by the stance adopted here.

10.5. Relation to earlier Six Birds results

Six Birds Foundations I ^[1] introduced the theory-package formalism $\mathcal{T} = (Z, f, \Sigma_f, E, \mathcal{A})$ and derived the $P1$ – $P6$ vocabulary under minimal hypotheses of composability with limited interface access. The present paper does not re-derive that vocabulary; it takes the six roles as given and asks a different question: whether, within a specific admissibility regime and over a specific class of admissible descriptions, the six roles are exactly enough. The Foundations theorem gives the canonical unavoidability side of the picture; the terminal theorem here gives the scoped exhaustion side. The two statements address different questions and address them at different levels; combining them does not yield a stronger claim, and neither theorem follows from the other.

The theorem chain also draws on a nearby result. ^[26], complementing the no-go analysis of ^[27], established that protocol holonomy is not a directionality certificate: noncommutativity of update sequences can be nonzero while path-reversal asymmetry remains zero, so route mismatch is a diagnostic, not an arrow of time. That observation is load-bearing for the role-projection alignment rules of Section 6, where $P3$ route/coherence mismatch is kept structurally distinct from the affinity-carrying face of $P6$. The next section records the scope boundaries and nonclaims of the present theorem stack.

11. Limitations and nonclaims

This section collects the scope restrictions of Theorem 9.1 and the global nonclaims preserved throughout.

11.1. Scope of the terminal theorem

Theorem 9.1 quantifies over $D \in \text{FATCD}$, where FATCD is the finite audited typed closure description domain of Definition 4.7. The theorem asserts:

- existence of a decomposition/exhaustion record $\text{Dec}(D)$ with six role channels and one status per channel;
- residual-feature exhaustion under the upper-bound classification, conditional on the recognition hypothesis (Definition 6.7) and within the covered scope, so that no seventh irreducible role is required there (the three principal threat classes are closed unconditionally);

- activation of any $\pi_i(D)$ only under the threshold, witness, level, and audit data of Definition 4.9.

It does not assert:

- decomposition uniqueness or canonicity;
- six active projections, or six active nontrivial witnesses, in every description;
- scope-free residual exhaustion, or any claim about features outside the covered scope;
- that the six roles are recognition-complete for FATCD, i.e. that every finite audited feature reduces to them;
- unrestricted exact-six, or coverage of all emergence phenomena;
- that broader domains require no further work.

11.2. Global nonclaims

Each nonclaim below is preserved throughout the paper and is enforced locally by the theorem-chain result that would otherwise overread it.

- *No full six-primitives algebra.* Scoped exact-six is proved over FATCD; it is not claimed that $P1-P6$ generate a global algebra of closure operations (Remark 2.9, Remark 8.6).
- *No global primitive-independence theorem.* Typed non-collapse (Proposition 3.3) is strictly weaker than global independence and is the sole claim at that level.
- *No empirical realization.* Empirical-bridge admissibility (Definition 3.11) admits threshold-dependent, level-indexed, instrument-visible substrate evidence; it does not assert that any $D \in \text{FATCD}$ is itself a substrate observation.
- *No unique selected lift.* Every lift is recorded as selected and non-unique (Definition 3.6).
- *No lower-to-higher determinacy.* Behavioral or verification content does not determine structural-categorical content.
- *No instrument-free formulation.* Admissible claims carry an instrument-relative visibility record (Definition 3.9).
- *No internal instrument-family completeness.* The accepted instrument family is not asserted to be sufficient for every exact-six judgment formulable in the same regime.
- *No unrestricted all-emergence coverage.* Theorem 9.1 is scoped to FATCD.
- *No surrogate terminal evidence.* The proof chain does not invoke any surrogate for the terminal scoped exact-six theorem.
- *No uniqueness of primitive presentations.* Distinct admissible presentations of the same role may coexist.
- *Cross-level comparison requires translation.* Direct comparison of claims at different levels is inadmissible without a translation record (Definition 3.5).

11.3. The endogenous-instrument meta-limit

An adjacent meta-theoretic question, deferred to future work, is whether the accepted internal instrument family can itself certify both its soundness and its completeness for exact-six judgments formulable inside the same regime. A plausible sharp formulation is that no internal argument should be expected to accomplish all three of the following at once:

1. establish the scoped exact-six claim;
2. certify the soundness of the accepted instrument family;
3. certify the completeness of that family for exact-six judgments expressible inside the same regime.

Such self-certification limits are of a piece with the classical Gödel–Löb tradition on internal consistency and provability statements [28], [29]. This question is not part of the proof of Theorem 9.1 and does not weaken or retract the theorem’s scoped conclusion. It is taken up as a follow-on direction in Section 12.

12. Future work

Theorem 9.1 and the limitations of Section 11 suggest several follow-on research directions. Each subsection below states a direction, not a commitment.

12.1. *Endogenous-instrument meta-limit*

A formalization of the meta-limit described in Section 11.3 would target a three-way impossibility statement: no internal argument can simultaneously establish the scoped exact-six claim, certify the soundness of the accepted instrument family, and certify its completeness for exact-six judgments formulable inside the same regime. Theorem 9.1 would remain intact under such a result; the additional content would be a precise account of self-certification boundaries. A closely related open question is *recognition-completeness*: whether every finite audited feature of a description in FATCD reduces to the six role meanings — the recognition hypothesis on which Theorem 6.9 is conditional (Remark 6.8). Candidate witnesses against it — finite conservation laws and emergent composite obstructions — would, if irreducible, populate the genuine-seventh-role category; establishing or refuting their reducibility is left to future work.

12.2. *Broader-domain exact-six*

Candidate extensions beyond FATCD include admissibility regimes that weaken finiteness, closure, or audit (Section 4), and regimes that admit annotation kinds beyond the six Ann_* categories of Definition 4.2. Any such extension would require fresh upper-bound and decomposition/exhaustion analysis, since the alignment rules and threat closures are not stated at scope-free strength.

12.3. *Uniqueness or canonical decomposition*

Theorem 9.1 is existential. A uniqueness study would investigate when distinct well-formed decomposition/exhaustion records of the same $D \in \text{FATCD}$ are equivalent under explicit translation records, or when a canonical representative can be selected. The channel-versus-activation discipline (Proposition 7.5) is compatible with such a study; no uniqueness claim is part of the present theorem.

12.4. *Full six-primitives algebra*

The generator-and-relation structure of $P1$ – $P6$ beyond typed non-collapse remains open. A full algebra would require scope-free independence data, not just the local boundary distinctions of Table 2, and would likely require a categorical apparatus richer than the raw grammar of Section 4.1.

12.5. Empirical substrate mapping

The empirical-bridge admissibility regime (Definition 3.11) can be connected to concrete substrate-level observables. The present paper treats empirical bridges as threshold-dependent, level-indexed, and instrument-visible annotations; later work can couple this to substrate-specific measurement theory while preserving the nonclaim of empirical realization.

12.6. Mechanization

A Lean 4 development accompanies the present paper; Appendix B records, claim by claim, how far each definition and theorem is currently mechanized. Strengthening that mechanization toward full proofs — in particular a Lean account of the boundary distinctions, the decomposition-existence construction, and the recognition hypothesis underlying the upper-bound classification — is a natural direction, building on the closure-related infrastructure already mechanized in the Foundations program ^[1], ^[30].

12.7. Cross-instrument comparison

Decomposition/exhaustion records of the same $D \in \text{FATCD}$ can be compared under different admissible instrument assignments. Definition 3.9 admits visibility records that suppress different portions of a claim under different instruments; a cross-instrument study could make the instrument-dependence of $\text{Dec}(D)$ explicit and sharpen the notion of instrument-relative claim strength.

Appendix A. Provenance and evidence hygiene

The development of the theorem chain of Sections 4–8 was assisted by automated drafting and proof-search tooling, the Aristotle system ^[31]. We record this for transparency and to fix the division of evidential labor precisely. Such tooling produces *candidate* definitions, statements, and arguments; it carries no evidential weight of its own, and nothing in this paper rests on it.

The mathematical content is established solely by the proofs given in the body. The machine-*checked* component of the work — which declarations are realized in Lean, and at what fidelity — is documented separately in Appendix B; the automated tooling named above is distinct from that mechanization and assisted drafting only. No raw output of that tooling is offered as mathematical evidence anywhere in the manuscript: where a section draws on machine assistance, the load-bearing content is the section's own proof, and the recorded mechanization is exactly as described in Appendix B.

A.1. Code availability

The LaTeX source of this manuscript and the accompanying Lean mechanization artifacts are available at <https://github.com/ioannist/six-birds-foundations-ii>.

Appendix B. Mechanization: declarations, coverage, and trust base

The theorem chain of this paper is accompanied by a Lean 4 development, the integrated SixBirds project, available at the code repository of Appendix A.1. This appendix is the canonical disclosure venue for that development. For every numbered

definition and claim it records the Lean declaration that realizes it and the *fidelity* of that realization, so the per-theorem footnotes in the body can be read against a single source of record. The mapping is maintained against a statements-of-record file in the repository and reproduced here; coverage is audited per declaration against that file.

B.1. What the footnotes mean

Each body theorem and proposition carries exactly one footnote naming its Lean declaration and its fidelity. We distinguish five fidelities:

- *verified* — a faithful Lean derivation of the statement from non-trivial hypotheses;
- *narrowed / parametric* — a genuine Lean proof, but over a single concrete witness, a parametric record, or a narrowed input space rather than the full generality of the paper statement;
- *schema (projection)* — the Lean declaration projects a field of, or restates, a proof-carrying record; it repackages rather than derives, and is therefore *not* described as a verified proof;
- *trivial* — the Lean statement reduces to `True` by construction, or ranges over a hand-built total classifier; it records a bookkeeping fact, not a derivation;
- *typed mirror* — for definitions, a Lean structure or inductive mirroring the paper definition.

We state plainly that **no claim in the present development is a faithful derivation in the first sense**. The Lean track is a typed harness: it fixes the objects as typed records, exposes the channel, status, and residual machinery, and checks that malformed and misaligned cases are rejected (the semantic regression tests of the `SixBirds.SemanticTests` module); but the load-bearing classification and closure statements are projections, schemas, or trivially-true bookkeeping over those records. The footnotes are worded accordingly. A reader who wants the mathematics should read the body proofs; the Lean track certifies typed well-formedness and case rejection, not the theorems.

B.2. The recognition hypothesis and the upper-bound classifier

The honest status of the central no-seventh-role result (Theorem 6.9) is legible in the Lean. The classifier `SixBirds.classifyResidual` is a total function on the finite raw-record kinds, every branch of which lands in a covered class; the declarations `upper_bound_classification` and the three closure lemmas are `True` by construction over it. This is exactly why the body states the result conditionally, under the recognition hypothesis (Definition 6.7): the Lean classifier *is* that hypothesis, encoded as a conditional schema rather than discharged. The paper-side recognition hypothesis has no independent Lean correspondent, and the open recognition-completeness question (Remark 6.8) is precisely the gap between the classifier's by-construction coverage and a derived exhaustiveness.

B.3. Declaration and coverage table

The table maps each paper label to its Lean declaration in the `SixBirds` namespace and its audited coverage. All declarations are sorry-free and axiom-clean relative to the trust base of §B.4; the coverage column records fidelity in the sense of §B.1, not mere presence.

§	paper label	Lean declaration	coverage
§2	def:primitive-roles	SixBirds.Role	typed mirror (def.)
§2	def:level-trichotomy	SixBirds.Level	typed mirror (def.)
§2	def:selected-lift-atlas	SixBirds.SelectedLiftAtlas	typed mirror (def.)
§2	prop:boundary-distinctions	SixBirds.boundary_distinctions	schema (projection)
§3	prop:honest-bookkeeping	SixBirds.honest_bookkeeping_admissibility	schema (projection)
§3	prop:typed-non-collapse	SixBirds.typed_non_collapse	schema (projection)
§3	prop:level-profile	SixBirds.level_profile	trivial
§3	prop:forgetting-lifts	SixBirds.forgetting_lifts	schema (projection)
§3	prop:activation-thresholds	SixBirds.activation_thresholds	schema (projection)
§3	prop:visibility	SixBirds.visibility	schema (projection)
§3	prop:empirical-bridge	SixBirds.empirical_bridge	schema (projection)
§3	prop:nonclaim-closure	SixBirds.nonclaim_closure	schema (projection)
§4	def:raw-record	SixBirds.RawRecord	typed mirror (def.)
§4	def:annotations	SixBirds.AnnotationKind	typed mirror (def.)
§4	def:finite	SixBirds.FiniteDescription	typed mirror (def.)
§4	def:closed	SixBirds.ClosedDescription	typed mirror (def.)
§4	def:audited	SixBirds.AuditedDescription	typed mirror (def.)
§4	def:fatcd	SixBirds.FATCD	typed mirror (def.)
§4	def:threshold-attachment	SixBirds.ThresholdAttachment	typed mirror (def.)
§4	def:role-projection	SixBirds.DerivedRoleProjection	typed mirror (def.)
§4	def:residual-scheme	SixBirds.ResidualClassificationScheme	typed mirror (def.)
§4	def:scoped-exact-six-question	SixBirds.ScopedExactSixQuestion	typed mirror (def.)
§4	prop:non-circularity	SixBirds.non_circularity	schema (projection)
§4	prop:non-overbreadth	SixBirds.non_overbreadth	schema (projection)
§5	def:common-witness-data	SixBirds.CommonWitnessData	typed mirror (def.)
§5	thm:six-role-lower-bounds	SixBirds.six_role_lower_bounds	schema (projection)
§5	prop:p1-witness	SixBirds.p1_witness	narrowed / parametric
§5	prop:p2-witness	SixBirds.p2_witness	narrowed / parametric
§5	prop:p3-witness	SixBirds.p3_witness	narrowed / parametric
§5	prop:p4-witness	SixBirds.p4_witness	narrowed / parametric

§	paper label	Lean declaration	coverage
§5	prop:p5-witness	SixBirds.p5_witness	narrowed / parametric
§5	prop:p6-witness	SixBirds.p6_witness	narrowed / parametric
§6	prop:alignment-rules	SixBirds.alignment_rules	schema (projection)
§6	prop:probability-closure	SixBirds.probability_closure	trivial
§6	prop:causality-closure	SixBirds.causality_closure	trivial
§6	prop:agency-closure	SixBirds.agency_closure	trivial
§6	thm:upper-bound-classification	SixBirds.upper_bound_classification	trivial
§7	def:channel-status	SixBirds.ChannelStatus	typed mirror (def.)
§7	def:decomposition-record	SixBirds.DecompositionRecord	typed mirror (def.)
§7	thm:decomposition-existence	SixBirds.decomposition_existence	schema (projection)
§7	thm:six-channel-exactness	SixBirds.six_channel_exactness	schema (projection)
§7	prop:channel-vs-activation	SixBirds.channel_vs_activation	schema (projection)
§7	prop:non-uniqueness	SixBirds.non_uniqueness	narrowed / parametric
§8	prop:dependency-stress	SixBirds.dependency_stress	schema (projection)
§8	prop:circularity-stress	SixBirds.circularity_stress	schema (projection)
§8	prop:channel-activation-stress	SixBirds.channel_activation_stress	narrowed / parametric
§8	prop:role-alignment-stress	SixBirds.role_alignment_stress	narrowed / parametric
§8	prop:residual-stress	SixBirds.residual_stress	trivial
§8	prop:overclaim-stress	SixBirds.overclaim_stress	schema (projection)
§8	thm:integrated-stress	SixBirds.integrated_stress	schema (projection)
§9	thm:scoped-exact-six	SixBirds.scoped_exact_six	schema (projection)

Table 1.

B.4. Trust base and axiom audit

The dependency audit accepts only the standard Lean and Mathlib trust base: `Classical.choice`, `propext`, `Quot.sound`, and `funext` (together with `choice`). No project-local axiom is admitted for any declaration recorded as a theorem. The validation harness (`scripts/validate_lean.sh`) rejects project-local axioms and sorry, and runs a semantic audit that rejects `:= True` and `:= False` definition stubs and all-purpose residual classifiers. The trust-base check certifies that the typed harness is axiom-clean relative to that base; it does not, and is not intended to, upgrade a schema or a trivially-true statement to a derivation. The fidelity column of §B.3 is the governing record of what each declaration establishes.

pair	distinction rule	permitted interaction	forbidden collapse
<i>P1</i> / <i>P5</i>	<i>P5</i> supplies quotient packaging; <i>P1</i> tests selected update compatibility with quotient packaging	<i>P1</i> uses <i>P5</i> when checking descent or obstruction	packaging alone is not <i>P1</i>
<i>P1</i> / <i>P4</i>	<i>P4</i> supplies stage/refinement structure; <i>P1</i> may vary by stage	stage-indexed descent or obstruction	stage index alone is not <i>P1</i>
<i>P1</i> / <i>P2</i>	<i>P1</i> is update compatibility; <i>P2</i> is macro-specification representability	representability may be studied over quotient structures affected by descent	representability failure is not automatically descent obstruction
<i>P1</i> / <i>P3</i>	<i>P1</i> is failure or success of a descended route; <i>P3</i> compares already-defined routes	<i>P1</i> may block a route before <i>P3</i> comparison	route undefinedness is not route mismatch
<i>P1</i> / <i>P6</i>	<i>P6</i> audits a <i>P1</i> event; <i>P1</i> is the event content	audited descent success or obstruction	audit record is not descent obstruction
<i>P2</i> / <i>P5</i>	<i>P5</i> is quotient packaging; <i>P2</i> requires macro-specification representability beyond quotient equality	<i>P2</i> can collapse to <i>P5</i> when specifications are just quotient classes; nontrivial <i>P2</i> uses <i>P5</i>	nontrivial <i>P2</i> is not mere packaging
<i>P2</i> / <i>P4</i>	<i>P4</i> supplies stages; <i>P2</i> representability may vary by stage	stage-indexed representability	stage variation alone is not <i>P2</i>
<i>P2</i> / <i>P3</i>	<i>P2</i> concerns representability of specifications; <i>P3</i> concerns route disagreement	route grammar may be built over represented macro structure	non-representability is not route mismatch
<i>P2</i> / <i>P6</i>	<i>P6</i> audits representability claims; <i>P2</i> is the representability content	audited representability or non-representability	audit trail is not representability
<i>P3</i> / <i>P4</i>	<i>P4</i> supplies stage/regime structure; <i>P3</i> route mismatch may be stage-indexed	route systems vary over valid stages	stage-indexed route mismatch is not <i>P4</i> itself
<i>P3</i> / <i>P5</i>	<i>P5</i> supplies quotient objects; <i>P3</i> compares defined routes over objects	routes over quotient objects	quotient object is not route mismatch
<i>P3</i> / <i>P6</i>	<i>P6</i> audits route comparisons; <i>P3</i> is the route mismatch content	audited route mismatch	audit record is not route mismatch
<i>P4</i> / <i>P5</i>	<i>P4</i> supplies stage/refinement structure; <i>P5</i> supplies packaging at a stage	stage-indexed packaging	packaging is not stage structure
<i>P4</i> / <i>P6</i>	<i>P4</i> supplies staged structure; <i>P6</i> audits staged histories or transitions	audited stage transitions	audit record is not staging
<i>P5</i> / <i>P6</i>	<i>P5</i> supplies packaging; <i>P6</i> audits packaging events	audited packaging	audit record is not packaging map

Table 2. Local boundary distinctions. Each row records the distinction rule between two roles, a permitted interaction, and the collapse that is forbidden under scoped bookkeeping.

Footnotes

¹ Tracked in the formalization harness as a conditional schema (`SixBirds.boundary_distinctions`); see Appendix B.

² Tracked in the formalization harness as a conditional schema (`SixBirds.typed_non_collapse`); see Appendix B.

³ Tracked in the formalization harness as a conditional schema (`SixBirds.nonclaim_closure`); see Appendix B.

⁴ Tracked in the formalization harness as a conditional schema (`SixBirds.non_circularity`); see Appendix B.

⁵ Tracked in the formalization harness as a conditional schema (`SixBirds.non_overbreadth`); see Appendix B. The Lean declaration is identical in statement and proof to `SixBirds.non_circularity` and re-extracts the admissibility fields of the record rather than establishing the structural-restriction content of the proposition.

⁶ Tracked in the formalization harness as a conditional schema (`SixBirds.six_role_lower_bounds`); see Appendix B.

⁷ The Lean development establishes this for one concrete description `witnessDescription Role.P1` under parametric hypotheses (alignment, threshold attachment, and finite/closed/audited verifications of that fixed witness); the narrowing is recorded in Appendix B (`SixBirds.p1_witness`).

⁸ The Lean development establishes this for one concrete description `witnessDescription Role.P2` under parametric hypotheses (alignment, threshold attachment, and finite/closed/audited verifications of that fixed witness); the narrowing is recorded in Appendix B (`SixBirds.p2_witness`).

⁹ The Lean development establishes this for one concrete description `witnessDescription Role.P3` under parametric hypotheses (alignment, threshold attachment, and finite/closed/audited verifications of that fixed witness); the narrowing is recorded in Appendix B (`SixBirds.p3_witness`).

¹⁰ The Lean development establishes this for one concrete description `witnessDescription Role.P4` under parametric hypotheses (alignment, threshold attachment, and finite/closed/audited verifications of that fixed witness; the attached level is fixed to the verification level rather than the structural-categorical level of the prose); the narrowing is recorded in Appendix B (`SixBirds.p4_witness`).

¹¹ The Lean development establishes this for one concrete description `witnessDescription Role.P5` under parametric hypotheses (alignment over the same raw-kind cluster used for P1, threshold attachment, and finite/closed/audited verifications of that fixed witness); the narrowing is recorded in Appendix B (`SixBirds.p5_witness`).

¹² The Lean development establishes this for one concrete description `witnessDescription Role.P6` under parametric hypotheses (alignment over event/provenance/replay/check records, threshold attachment, and finite/closed/audited verifications of that fixed witness); the narrowing is recorded in Appendix B (`SixBirds.p6_witness`).

¹³ Tracked in the formalization harness as a conditional schema (`SixBirds.alignment_rules`); see Appendix B.

¹⁴ The Lean declaration is a trivial single-record harness check: `classifyResidual` maps `RawKind.probability` to `composite`, whose `CoveredResidualClass` branch is `True`; it does not establish the paper’s universal probability/uncertainty closure. See Appendix B (`SixBirds.probability_closure`).

¹⁵ The Lean declaration is a trivial single-record harness check: `classifyResidual` maps `RawKind.causality` to `composite`, whose `CoveredResidualClass` branch is `True`; the $P5$ -inadmissibility content of the proposition is absent and no universal residual-causality closure is established. See Appendix B (`SixBirds.causality_closure`).

¹⁶ The Lean declaration is a trivial single-record harness check: `classifyResidual` maps `RawKind.agency` to `empiricalBridgeAnnotation`, whose `CoveredResidualClass` branch is `True`; the $P1$ -inadmissibility content is absent and no universal residual-agency closure is established. See Appendix B (`SixBirds.agency_closure`).

¹⁷ The Lean declaration is a classifier-by-construction check: every `RawKind` branch of `classifyResidual` lands in a class whose `CoveredResidualClass` branch is `True`, and the uncovered constructors `unresolvedThreat/genuineSeventhRole` are never produced; “no seventh role” holds by construction rather than by derivation from feature content. See Appendix B (`SixBirds.upper_bound_classification`).

¹⁸ Tracked in the formalization harness as a conditional schema (`SixBirds.decomposition_existence`); see Appendix B.

¹⁹ Tracked in the formalization harness as a conditional schema (`SixBirds.six_channel_exactness`); see Appendix B.

²⁰ Tracked in the formalization harness as a conditional schema (`SixBirds.channel_vs_activation`); see Appendix B.

²¹ The Lean development exhibits two concrete decomposition records of a single description that differ only in their $P1$ channel status; the Lean statement does not assert that either record is well-formed, and the further sources of non-uniqueness in the prose (level assignments, witness-schema choice, residual labels, loss/lift placement) are not mechanized. The narrowing is recorded in Appendix B (`SixBirds.non_uniqueness`).

²² Tracked in the formalization harness as a conditional schema (`SixBirds.integrated_stress`) bundling the six stress conjuncts; see Appendix B.

²³ Tracked in the formalization harness as a conditional schema (`SixBirds.scoped_exact_six`); see Appendix B.

References

1. ^a ^b ^c ^d ^e ^f Tsikos I (2026). "Six Birds: Foundations of Emergence Calculus." arXiv. doi:[10.48550/arXiv.2602.00134](https://doi.org/10.48550/arXiv.2602.00134).
2. ^Δ Anderson P (1972). "More Is Different." *Science*. 177(4047):393–396. doi:[10.1126/science.1774047393](https://doi.org/10.1126/science.1774047393).
3. ^Δ Davey B, Priestley H (2002). *Introduction to Lattices and Order*. Cambridge: Cambridge University Press. doi:[10.1017/CBO9780511809088](https://doi.org/10.1017/CBO9780511809088).
4. ^Δ Grothendieck A (1971). *Revêtements Étales Et Groupe Fondamental (SGA 1) [Étale Coverings and Fundamental Group (SGA 1)]*. Berlin: Springer. doi:[10.1007/BFb0058656](https://doi.org/10.1007/BFb0058656).
5. ^Δ Janelidze G, Tholen W (1994). "Facets of Descent, I." *Appl Categ Struct*. 2(3):245–281. doi:[10.1007/BF00878100](https://doi.org/10.1007/BF00878100).
6. ^Δ Mac Lane S (1998). *Categories for the Working Mathematician*. New York: Springer. doi:[10.1007/978-1-4757-4721-8](https://doi.org/10.1007/978-1-4757-4721-8).

7. [△]Baader F, Nipkow T (1998). *Term Rewriting and All That*. Cambridge: Cambridge University Press. doi:[10.1017/CBO9781139172752](https://doi.org/10.1017/CBO9781139172752).
8. [△]Nakahara M (2003). *Geometry, Topology and Physics*. Bristol: Institute of Physics Publishing.
9. [△]Milner R (1989). *Communication and Concurrency*. Englewood Cliffs: Prentice Hall.
10. [△]Back R, von Wright J (1998). *Refinement Calculus: A Systematic Introduction*. New York: Springer. doi:[10.1007/978-1-4612-1674-2](https://doi.org/10.1007/978-1-4612-1674-2).
11. [△]Kemeny J, Snell J (1960). *Finite Markov Chains*. Princeton: Van Nostrand.
12. [△]Green T, Karvounarakis G, Tannen V (2007). "Provenance Semirings." ACM. doi:[10.1145/1265530.1265535](https://doi.org/10.1145/1265530.1265535).
13. [△]Merkle R (1988). "A Digital Signature Based on a Conventional Encryption Function." Berlin: Springer. doi:[10.1007/3-540-48184-2_32](https://doi.org/10.1007/3-540-48184-2_32).
14. [△]Seifert U (2012). "Stochastic Thermodynamics, Fluctuation Theorems and Molecular Machines." *Rep Prog Phys*. 75(12):126001. doi:[10.1088/0034-4885/75/12/126001](https://doi.org/10.1088/0034-4885/75/12/126001).
15. [△]Martin-Löf P (1984). *Intuitionistic Type Theory*. Napoli: Bibliopolis.
16. [△]Hoare C (1969). "An Axiomatic Basis for Computer Programming." *Commun ACM*. 12(10):576–580. doi:[10.1145/363235.363259](https://doi.org/10.1145/363235.363259).
17. [△]van Fraassen B (1980). *The Scientific Image*. Oxford: Oxford University Press. doi:[10.1093/0198244274.001.0001](https://doi.org/10.1093/0198244274.001.0001).
18. [△]Libkin L (2004). *Elements of Finite Model Theory*. Berlin: Springer. doi:[10.1007/978-3-662-07003-1](https://doi.org/10.1007/978-3-662-07003-1).
19. [△]Cardelli L, Mitchell J (1991). "Operations on Records." *Math Struct Comput Sci*. 1(1):3–48. doi:[10.1017/S0960129500000049](https://doi.org/10.1017/S0960129500000049).
20. [△]Pearl J (2009). *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press. doi:[10.1017/CBO9780511803161](https://doi.org/10.1017/CBO9780511803161).
21. [△]Bratman M (1987). *Intention, Plans, and Practical Reason*. Cambridge: Harvard University Press.
22. [△]Cartwright N (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press. doi:[10.1017/CBO9781139167093](https://doi.org/10.1017/CBO9781139167093).
23. [△]Ladyman J, Ross D (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press. doi:[10.1093/acprof:oso/9780199276196.001.0001](https://doi.org/10.1093/acprof:oso/9780199276196.001.0001).
24. [△]Abramsky S, Brandenburger A (2011). "The Sheaf-Theoretic Structure of Non-Locality and Contextuality." *New J Phys*. 13(11):113036. doi:[10.1088/1367-2630/13/11/113036](https://doi.org/10.1088/1367-2630/13/11/113036).
25. [△]Bohr N (1958). *Atomic Physics and Human Knowledge*. New York: John Wiley & Sons.
26. [△]Tsiokos I (2026). "Six Birds Protocol Trap: Holonomy Without Entropy Production." Zenodo. doi:[10.5281/zenodo.18906902](https://doi.org/10.5281/zenodo.18906902).
27. [△]Tsiokos I (2026). "Six Birds: No-Go Theorems for Audited Emergence." Zenodo. doi:[10.5281/zenodo.19187777](https://doi.org/10.5281/zenodo.19187777).
28. [△]Gödel K (1931). "Über Formal Unentscheidbare Sätze Der Principia Mathematica Und Verwandter Systeme I" [On Formally Undecidable Propositions of Principia Mathematica and Related Systems I]. *Monatsh Math Phys [Journal of Mathematics and Physics]*. 38:173–198. doi:[10.1007/BF01700692](https://doi.org/10.1007/BF01700692).
29. [△]Löb M (1955). "Solution of a Problem of Leon Henkin." *J Symb Log*. 20(2):115–118. doi:[10.2307/2266895](https://doi.org/10.2307/2266895).
30. [△]de Moura L, Ullrich S (2021). *The Lean 4 Theorem Prover and Programming Language*. Cham: Springer International Publishing. doi:[10.1007/978-3-030-79876-5_37](https://doi.org/10.1007/978-3-030-79876-5_37).
31. [△]Achim T, Best A, Bietti A, Federico M, Halpern-Leistner D, Kudryashov Y, Michelsen M, Rodriguez E, Shanker V, Sicca V, Swope A, Tamas M, Thomm J, Williams H (2025). "Aristotle: IMO-Level Automated Theorem Proving." arXiv. doi:[10.48550/arXiv.2510.01346](https://doi.org/10.48550/arXiv.2510.01346).

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.