To Classify a Stone with Six Birds: Primitive Activation and Observable Separation of Layer-Birth Classes

Ioannis Tsiokos

March 16, 2026

Abstract

We ask not merely whether layer birth occurs, but how many primitive kinds of birth are experimentally accessible. Within the Six-Birds framework, birth events can be organized by activation of packaging ($P5$), directed circulation ($P6_{\mathrm{drive}}$), and staging anomaly ($P4$). We show that these primitives support a small canonical taxonomy of four layer-birth classes and that the corresponding classes can be operationally instantiated and separated in a common measured observable space. Class-I activates packaging without drive; Class-II activates packaging together with directed circulation; Class-III activates packaging together with a staging anomaly but without drive; and Class-IV activates all three. The strongest empirical separation is between Class-I and Class-II, where affinity at structural birth robustly discriminates drive-free from driven births and explicit no-fake-arrow controls rule out coarse-graining artifacts. We then place all four classes on a shared observable map using structural boundary location, affinity at boundary, and a strict-rule-aligned staging proxy. Capacity/depth behaves as a post-birth field rather than a primary classifier, and at least one upstream/PICA-style system can be embedded in the same coordinate system. We do not claim stabilized universality exponents. Instead, we provide a primitive taxonomy and observable separation of layer-birth classes that future universality work must respect.

1 Introduction

Layer birth is the moment when a coarse description stops being merely convenient and becomes a genuine theory object: a description with stable macroscopic states, reproducible distinctions, and, in some cases, its own directed bookkeeping. For a physical theory of layered emergence, existence is not enough. One also wants a taxonomy: how many fundamentally different kinds of birth are there, and what observables separate them?

Earlier Six-Birds work developed the underlying vocabulary for that question. It introduced a finite theory-package viewpoint, identified the primitive operations needed to describe emergence and directionality, and showed how minimal stochastic substrates can be audited for structural and thermodynamic regimes [12, 14]. It also established an important constraint: holonomy or route mismatch alone is not a certificate of directed time-asymmetry; any genuine arrow claim must be supported by an audit channel carrying drive [13, 6, 11]. What remained open was not whether layer birth can occur, but how to classify its experimentally accessible forms.

In the present paper we use the Six-Birds primitives as class variables. The full primitive list is: operator rewrite ($P1$), gating/support restriction ($P2$), holonomy/route mismatch ($P3$), sectors and

Automorph Inc., Wilmington, DE, USA ioannis@automorph.io Zenodo DOI: 10.5281/zenodo.19047665

staging structure ($P4$), packaging ($P5$), and audit/accounting ($P6$), with $P6_{\text{drive}}$ denoting genuine directed circulation. Only three primitives act as class-defining axes at birth in this paper: $P5$, $P6_{\text{drive}}$, and $P4$. $P1$ and $P2$ parameterize approaches to birth, while $P3$ is diagnostically important but not a class axis on its own.

This yields a small canonical taxonomy. Class-I births activate packaging without drive: $P5$ only. Class-II births activate packaging together with directed circulation: $P5 + P6_{\text{drive}}$. Class-III births activate packaging together with a staging anomaly but without drive: $P5 + P4$. Class-IV births activate all three: $P5 + P4 + P6_{\text{drive}}$. The paper's main empirical claim is that these four canonical classes can be operationally instantiated and separated in a shared measured observable space.

The observables used to separate them are deliberately operational rather than asymptotic. Structural birth is tracked by a pair of boundary proxies built from closure error and objecthood order. Drive is tracked by affinity at and after structural birth. Staging anomaly is tracked by a dual-shift proxy that asks whether the structurally relevant boundary moves across depth or timescale. These observables are sufficient to separate the confirmed classes, even though they do not yet support stable critical-exponent claims.

That scope condition matters. We explicitly do not present this paper as the final universality-exponent analysis of layer birth. Our size-stability audits show that exponent extraction is not yet stable enough to carry the main narrative. The contribution here comes earlier in the scientific sequence: identify the class-defining primitive activations, show that the low-complexity classes are real, and construct the observable map that any later universality analysis will have to respect.

Within that scope, the strongest result is the Class-I / Class-II separation. Both classes exhibit structural birth, but only Class-II carries a robust directed-circulation signal. The affinity contrast at structural birth is therefore not a side measurement; it is the primary robust discriminant between the no-drive and drive-active classes. That interpretation is reinforced by explicit no-fake-arrow controls, which show that reversible coarse-grained systems are not spuriously labeled as driven under the default audit.

The P4-active classes, Class-III and Class-IV, require more care. They are confirmed in the canonical manual-lens channel and admit usable phenomenology panels, but they are more fragile under non-manual lens swaps than Class-I and Class-II. We therefore treat them as confirmed classes with quantified robustness caveats, not as equally mature robustness objects across all analysis channels. This asymmetry is part of the result rather than an embarrassment: it shows that staging-anomalous births are real, but methodologically more delicate.

The paper proceeds as follows. Section 2 introduces the primitive framework and the canonical class rubric. Section 3 defines the operational observables and evidence discipline. Section 4 presents the Class-I / Class-II separation and its no-fake-arrow validation. Section 5 establishes the P4-active classes. Section 6 places all four classes in a common observable-space map and summarizes the supporting capacity and upstream-bridge results. Section 7 discusses scope, limitations, and what remains open. The main conclusion is that layer birth is not one phenomenon but a small family of primitive-activation events, and that the experimentally accessible low-complexity cases already support a four-class observable taxonomy.

2 Primitive framework and canonical class rubric

The organizing claim of this paper is that layer births are best classified by *primitive activation*. The Six-Birds framework supplies the primitive vocabulary; the present paper asks which of those primitives actually function as class-defining axes at birth. Our answer is deliberately narrow.

Table 1. Six-Birds primitive vocabulary and the role of each primitive in this paper.

Primitive	Short name	Role in this paper
$P1$	operator rewrite	parameterizes approach to birth
$P2$	gating/support restriction	parameterizes approach to birth
$P3$	holonomy/route mismatch	diagnostic only; not a class axis on its own
$P4$	sectors/staging anomaly	class-defining axis when staging becomes canonically active
$P5$	packaging	class-defining structural birth axis
$P6$	audit/accounting	with $P6_{\text{drive}}$ the class-defining drive axis

Table 2. Canonical layer-birth classes used in this paper.

Class	$P5$	$P6_{\text{drive}}$	$P4$
Class-I	yes	no	no
Class-II	yes	yes	no
Class-III	yes	no	yes
Class-IV	yes	yes	yes

Although the full framework contains six primitives, only three act as class-defining axes in the current taxonomy: packaging ($P5$), directed circulation ($P6_{\text{drive}}$), and staging anomaly ($P4$). The remaining primitives still matter, but in different ways [9]. Operator rewrite ($P1$) and gating/support restriction ($P2$) parameterize approaches to birth rather than the class of the birth itself, while holonomy/route mismatch ($P3$) is diagnostically important but not a class axis on its own [12, 13].

Table 1 lists the six primitives in the compact form used in this paper. The rightmost column does not restate the full framework; it states the role each primitive plays *here*. This is important because the paper is not trying to recapitulate the entire Six-Birds program. It is using the primitive system as a controlled grammar for a birth taxonomy.

The central restriction is that $P3$ is not allowed to certify a class by itself. Earlier work showed that route mismatch or holonomy alone is not an arrow-of-time certificate; any genuine directional claim must be supported by an audit channel carrying drive [13]. That constraint is reflected directly in the taxonomy: there is no "holonomy-only" class in the canonical table below. Likewise, $P1$ and $P2$ are essential in how one moves through parameter space toward a birth event, but they are not treated here as independent class labels.

This leaves three class-defining axes. $P5$ records whether packaging has activated strongly enough that the theory now recognizes stable objects. $P6_{\text{drive}}$ records whether genuine directed circulation is present at birth. $P4$ records whether staging structure—timescale or depth structure that is invisible in one channel and visible in another—has become class-defining rather than merely detectable. With those axes fixed, the canonical low-complexity birth classes are exactly those shown in Table 2.

Table 2 gives the compact class-by-primitive view used throughout the paper.

Two clarifications are crucial for the rest of the paper. First, the present taxonomy is *canonical*, not exhaustive in the logical sense. It captures the experimentally accessible primitive combinations that are both theoretically motivated and operationally measurable in the current framework. Second, the paper distinguishes between $P4$-*like signal* and *canonical* $P4$ *activation*. A system may show measurable tau/depth sensitivity in structural boundary proxies without that sensitivity

being strong enough, stable enough, or symmetric enough across the relevant channels to count as a class-defining staging anomaly. This distinction is not cosmetic. It is what allows the reference Class-I and Class-II systems to exhibit measurable staging sensitivity while remaining canonically non-$P4$ classes.

The canonical class rubric is therefore a two-level object. At the primitive-signal level, one may detect traces of $P4$-like behavior or weak route dependence. At the class level, however, the birth is labeled only by which primitive axes are sufficiently activated to define a stable experimental category. The four classes studied in this paper are the four lowest-complexity such categories: packaging alone (Class-I), packaging plus drive (Class-II), packaging plus staging anomaly (Class-III), and packaging plus both staging anomaly and drive (Class-IV). The rest of the manuscript operationalizes those labels and shows that the corresponding classes can be separated in measured observable space.

3 Operational observables and evidence discipline

The present paper classifies births operationally rather than asymptotically. The class-defining quantities are therefore not final critical exponents but a small set of measurable observables that track the activation of $P5$, $P6_{\mathrm{drive}}$, and $P4$. The point of this section is to make that measurement discipline explicit. The paper does not claim that these observables exhaust every possible diagnosis of birth; it claims that they are sufficient to separate the experimentally confirmed classes in the current framework.

Structural birth is tracked through two boundary proxies. The first is a closure-based proxy built from closure error, using the reference threshold $\mathrm{CE} = 0.025$. The second is an objecthood-based proxy built from the objecthood order parameter, using the reference threshold $M_{\mathrm{obj}} = 0.90$ [8, 16]. Whenever both are available, we define the structural reference boundary as the more conservative of the two,

$$\lambda_{\mathrm{struct,ref}} = \max\big(\lambda_{\mathrm{CE}=0.025},\ \lambda_{M_{\mathrm{obj}}=0.90}\big).$$

This convention matters because the paper is not interested in merely detectable structure, but in the onset of *formed* structure. The structural channel is therefore intentionally conservative.

Directed circulation is tracked by affinity measured at and after structural birth [6, 11]. The primary scalar is the affinity at the structural reference boundary, denoted here as $\mathrm{Aff}_{\mathrm{ref}}$. We also record onset locations for fixed affinity thresholds, typically 10^{-3}, 10^{-2}, and 5×10^{-2}. These onset thresholds are not themselves class labels; they are operational markers of when a drive-carrying audit channel becomes appreciable. In practice they are what make the Class-I / Class-II separation so sharp: both classes activate packaging, but only Class-II carries affinity at structural birth at a level far above the effectively zero values seen in the drive-free classes.

Staging anomaly is tracked by tau/depth sensitivity in the structural boundary proxies [5, 3]. The crucial distinction is between $P4$-*like signal* and *canonical* $P4$ *activation*. A system exhibits $P4$-like signal if at least one structural boundary proxy moves materially across the tau comparison, or if a boundary proxy appears in one channel and disappears in another. Canonical $P4$ activation is stronger. In the present rubric it requires either a presence/absence shift or simultaneous material movement in both structural channels. To make that distinction measurable, the paper uses a strict-rule-aligned staging proxy

$$P4_{\mathrm{dual}} = \min\big(\Delta\lambda_{\mathrm{CE}},\ \Delta\lambda_{M_{\mathrm{obj}}}\big),$$

Table 3. Operational observables used in the present paper and their role in the class rubric.

Observable family	Operational proxy	Role in this paper
Structural birth	CE $= 0.025$, $M_{\mathrm{obj}} = 0.90$, and $\lambda_{\mathrm{struct,ref}} = \max(\lambda_{\mathrm{CE}}, \lambda_{M_{\mathrm{obj}}})$	primary structural birth channel ($P5$)
Directed circulation	$\mathrm{Aff}_{\mathrm{ref}}$ and affinity onset thresholds	primary drive channel ($P6_{\mathrm{drive}}$)
Staging anomaly	$\Delta\lambda_{\mathrm{CE}}$, $\Delta\lambda_{M_{\mathrm{obj}}}$, and $P4_{\mathrm{dual}} = \min(\Delta\lambda_{\mathrm{CE}}, \Delta\lambda_{M_{\mathrm{obj}}})$	primary staging channel ($P4$)
No-fake-arrow controls	reversible false-positive checks and protocol-trap handling	evidential validation of the drive channel
Capacity/depth	post-birth depth/headroom proxies	supporting, not class-defining

where $\Delta\lambda_{\mathrm{CE}}$ and $\Delta\lambda_{M_{\mathrm{obj}}}$ are the tau-separated shifts in the two structural boundary proxies. This is the observable that enters the four-class map. It is intentionally conservative: weak or asymmetric staging sensitivity can still count as $P4$-like without yet counting as class-defining $P4$ activation.

The no-fake-arrow controls are equally important. The paper does not interpret nonzero affinity as a class axis merely because it is convenient. Earlier Six-Birds work established that route mismatch or protocol holonomy alone is not enough to certify directionality [13]. The empirical controls used here are designed to test exactly that point. Reversible coarse-grained controls must not be labeled as driven under the default audit, while protocol-trap constructions with hidden scheduling must separate from phase-aware handling. This gives the affinity channel an evidential status that is stronger than a mere auxiliary measurement: it is the audited indicator that the drive axis is genuinely active.

These observables are summarized in Table 3. The table separates *class-defining* observables from *supporting* ones. This distinction should be kept in mind throughout the paper. Capacity/depth quantities, for example, are scientifically useful and later shown to behave as post-birth fields, but they are not class-defining observables in the present taxonomy. Likewise, susceptibility and Binder-like cumulants are useful phenomenology assets where available, but they are not what fixes the class rubric [4, 2].

A second discipline concerns analysis channels. For Class-I and Class-II, the signal survives standard lens perturbations well enough that one can discuss lens robustness directly in the main text. For Class-III and Class-IV, however, the situation is more delicate [7, 10]. The confirmed P4-active classes are established in the canonical manual-lens channel, and that is the channel in which their primary class claims should be read [15]. Non-manual lens behavior is measured and reported as a robustness caveat rather than folded into the class definition. This asymmetry is not a defect of presentation. It reflects the fact that staging-anomalous classes are empirically more fragile than the two lower-complexity classes.

The final discipline is one of scope. The paper does not treat exponent extraction as a settled result. Size-stability and identifiability audits showed that the current finite-size-scaling fits are too grid-sensitive and too boundary-attached to support strong universality-exponent claims. Exponent language is therefore intentionally softened throughout. The contribution here is earlier in the scientific pipeline: identify the class-defining observables, show that the corresponding classes are experimentally real, and build the observable map that any later exponent-level universality story would have to respect [12, 14].

Table 4. Reference structural and drive observables for the confirmed Class-I and Class-II systems.

Class	$P5$	$P6_{\mathrm{drive}}$	$P4$	λ_{CE}	$\lambda_{M_{\mathrm{obj}}}$	$\lambda_{\mathrm{struct,ref}}$	$\mathrm{Aff}_{\mathrm{ref}}$
Class-I	active	inactive	inactive	0.874	0.763	0.874	≈ 0
Class-II	active	active	inactive	0.556	0.500	0.556	3.45×10^{-2}

4 Class-I and Class-II: structural birth with and without drive

The first two confirmed classes are the simplest nontrivial birth types in the canonical rubric. Both activate packaging and therefore exhibit a genuine structural birth. What separates them is not the presence or absence of structure, but whether that structural birth is accompanied by genuine directed circulation. In the language of Section 2, both classes activate $P5$, while only Class-II activates $P6_{\mathrm{drive}}$. This is the cleanest and most robust empirical contrast in the paper.

Table 4 summarizes the reference quantities used in the comparison. Class-I is structurally conservative: its closure-based boundary occurs at $\lambda_{\mathrm{CE}=0.025} \approx 0.8743$, its objecthood-based boundary at $\lambda_{M_{\mathrm{obj}}=0.90} \approx 0.7625$, and its structural reference boundary therefore at $\lambda_{\mathrm{struct,ref}} \approx 0.8743$. Class-II is structurally earlier: $\lambda_{\mathrm{CE}=0.025} \approx 0.5556$, $\lambda_{M_{\mathrm{obj}}=0.90} = 0.5000$, and hence $\lambda_{\mathrm{struct,ref}} \approx 0.5556$. Structural birth is therefore present in both classes, but it occurs at different locations in control space.

What makes the classes sharply separable is the drive channel. At the structural reference boundary, Class-I has effectively zero affinity, $\mathrm{Aff}_{\mathrm{ref}} \approx -5.7 \times 10^{-19}$, while Class-II has a clearly nonzero value, $\mathrm{Aff}_{\mathrm{ref}} \approx 3.45 \times 10^{-2}$. The ratio is formally enormous because the Class-I reference value is numerically at machine zero; the stable statement is the absolute separation. Class-II also exhibits a finite affinity onset within the scanned lambda range, with the 10^{-3} threshold reached at approximately $\lambda = 0.40$, whereas Class-I has no nontrivial affinity onset at the configured thresholds. In other words, the two classes do not differ because one has structure and the other does not. They differ because one carries a directed circulation channel at structural birth and the other does not.

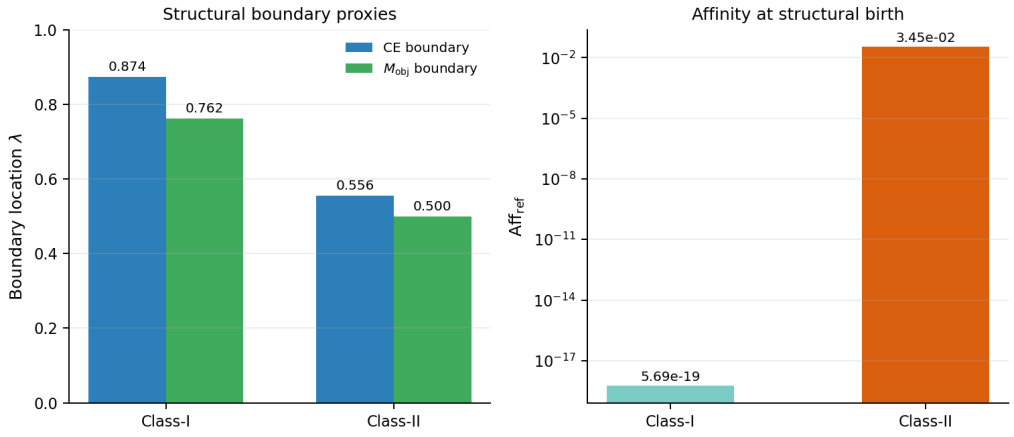

Figure 1. Consolidated Class-I / Class-II comparison showing structural boundary proxies and affinity separation. Both classes are structurally active, but only Class-II carries a robust drive signal at structural birth.

Figure 1 makes this contrast visible in one view. The structural proxies confirm that Class-I and Class-II both belong on the packaging side of the taxonomy. The affinity panel then separates them cleanly. This is the key reason the paper foregrounds affinity rather than exponent fits as the

Table 5. No-fake-arrow control summary used to validate the affinity channel.

Control family / case	Affinity	Driven candidate?
Hidden-schedule protocol trap effective kernel	0.0576	yes
Phase-aware protocol kernels (pair01/pair12/pair20)	0.0	no
Reversible coarse-grained controls	0.0	no

primary Class-I / Class-II discriminant. Earlier audits showed that exponent extraction remains too fragile to support a strong universality claim, but the affinity separation is already large, stable, and mechanistically interpretable.

That interpretation is strengthened by the no-fake-arrow controls. Earlier Six-Birds work showed that route mismatch by itself is not enough to certify directionality; any arrow claim must be supported by a true audit channel carrying drive [13, 6, 11]. The empirical controls in this project enforce that discipline directly. Table 5 summarizes the key cases. The hidden-schedule protocol-trap effective kernel is correctly identified as driven-positive. By contrast, the phase-aware kernels and the explicitly reversible controls remain drive-negative under the default audit. The important point is not merely that the controls "look reasonable." It is that the affinity channel survives an empirical falsification test: reversible coarse-grained systems are not spuriously promoted into the driven class.

Class-I and Class-II also differ in how tau sensitivity enters the interpretation. Both reference systems exhibit measurable staging sensitivity at the signal level, but under the present rubric neither is canonically $P4$-active. Class-I shows moderate shifts in both structural proxies without crossing the conservative $P4$ activation rule. Class-II shows a larger closure-based shift but no corresponding objecthood shift. This is precisely why the paper distinguishes between $P4$-like signal and canonical $P4$ activation. The existence of some tau sensitivity does not collapse the Class-I / Class-II distinction into the staging-anomalous classes.

Equally important, this separation is not confined to a single partition choice. For these two lower-complexity classes, lens perturbations preserve the core story well enough that the main text can speak about a robust empirical contrast rather than only a manual-channel contrast. The measured structural/drive signatures survive the standard spectral and diffusion lens perturbations strongly enough for the consolidated dashboard to classify the separation as stable. This is one reason the Class-I / Class-II result is the most mature result in the paper.

The right claim, then, is not that Class-I and Class-II already exhibit fully stabilized and distinct critical exponents. The right claim is that they are two experimentally real low-complexity birth classes with a robust primitive difference: both activate packaging, only one activates drive, and that difference is visible in a validated observable channel. In the current evidence base, this is the strongest non-exponent result in the paper.

5 Class-III and Class-IV: staging-anomalous births

The next two confirmed classes in the canonical rubric are the staging-anomalous births. They differ from Class-I and Class-II not by losing structural birth, but by making staging anomaly class-defining. In the language of Section 2, these are the births for which $P4$ is canonically active rather than merely detectable as weak tau sensitivity. Class-III activates packaging together with staging anomaly but without directed circulation. Class-IV activates packaging together with both staging anomaly and directed circulation. These classes are more methodologically delicate than

Table 6. Reference observables for the confirmed Class-III and Class-IV systems in the canonical manual channel. Both have $P5$ and $P4$ active.

Class	$P6_{\text{drive}}$	λ_{CE}	$\lambda_{M_{\text{obj}}}$	$\lambda_{\text{struct,ref}}$	Aff_{ref}	$P4_{\text{dual}}$
Class-III	inactive	0.562	0.123	0.562	≈ 0	0.202
Class-IV	active	0.533	0.067	0.533	2.99×10^{-3}	0.216

Representatives: Class-III = `replicated_portal_sp4`; Class-IV = `replicated_portal_sp4_drive_f`.

Class-I and Class-II, but they are nevertheless experimentally real in the canonical manual-lens channel [12].

Table 6 summarizes the representative systems used in the rest of the paper. For Class-III, the strongest confirmed representative is the replicated-portal system `replicated_portal_sp4`, which remains strict-confirmed across the checked size panel $\{16, 32, 64, 128\}$ in the manual channel. At the representative size used in the observable map, its structural reference boundary lies at $\lambda_{\text{struct,ref}} \approx 0.5617$, its affinity at boundary is effectively zero ($\text{Aff}_{\text{ref}} \approx 1.13 \times 10^{-17}$), and its strict-rule-aligned staging proxy is large ($P4_{\text{dual}} \approx 0.202$). This is exactly the signature required for a canonical Class-III birth: structural activation, no drive, and class-defining staging anomaly.

Class-IV is the driven counterpart. The confirmed representative `replicated_portal_sp4_drive_f` remains Class-IV across the checked size panel in the canonical manual channel. At the representative size used in the observable map, its structural reference boundary is $\lambda_{\text{struct,ref}} \approx 0.5333$, its affinity at boundary is positive ($\text{Aff}_{\text{ref}} \approx 2.99 \times 10^{-3}$), and its strict-rule-aligned staging proxy remains large ($P4_{\text{dual}} \approx 0.216$). The resulting signature is the expected one for Class-IV: structural birth, directed circulation, and staging anomaly all active together.

The conceptual importance of these two classes is that they show why the distinction between $P4$-like signal and canonical $P4$ activation matters. Earlier in the paper, the Class-I and Class-II references already exhibited measurable staging sensitivity at the signal level. That by itself was not enough to make them staging-anomalous classes. By contrast, the Class-III and Class-IV representatives clear the stricter dual-shift threshold, so staging anomaly is no longer just detectable—it becomes part of the class identity. In this sense the P4-active classes are not merely "more tau-sensitive" versions of the lower-complexity classes; they are births in which depth/timescale structure becomes class-defining [5, 3]. This also makes them natural places to look for later balance or edge-of-chaos-type phenomena, because they are the first classes in the present taxonomy for which competing depth/timescale structure enters at the class-defining level rather than as a weak side signal. That interpretive possibility is one motivation for taking the P4-active classes seriously, but it is not itself established by the present evidence.

The important point in both cases is persistence: the class labels survive across the checked size panels in the canonical manual analysis channel. This is the experimental core of the Class-III / Class-IV claim.

The phenomenology layer is weaker than for Class-I and Class-II, but still usable. Class-III admits a manual-channel shadow ensemble with seven admissible class-preserving perturbations and a usable susceptibility/Binder panel. Class-IV admits a corresponding panel with fifteen admissible perturbations. In both cases, the primary phenomenology observable is the tau-separated objecthood difference $\Delta M_{\text{obj}} = M_{\text{obj}}(\tau = 2) - M_{\text{obj}}(\tau = 1)$. The resulting susceptibility scales are modest (peaks on the order of 10^{-6}), and the Binder-like cumulants include partially flat regions, but the grouped

Table 7. Non-manual lens outcomes for the canonical Class-III and Class-IV representatives.

Case	Manual label	Non-manual outcome	$P5/P6_{\text{drive}}$	$P4_{\text{class}}$
Class-III / spectral	Class-III	Class-I	act. / inact.	false
Class-III / diffusion	Class-III	Class-I	act. / inact.	false
Class-IV / spectral	Class-IV	uncl. / Class-II	act. / unk.–act.	false
Class-IV / diffusion	Class-IV	uncl. / Class-II	act. / unk.–act.	false

act. = active; inact. = inactive; unk. = unknown; uncl. = unclassified.

panels are coherent enough to function as phenomenology assets rather than merely qualitative sketches. Figures 2 and 3 show the corresponding susceptibility panels.

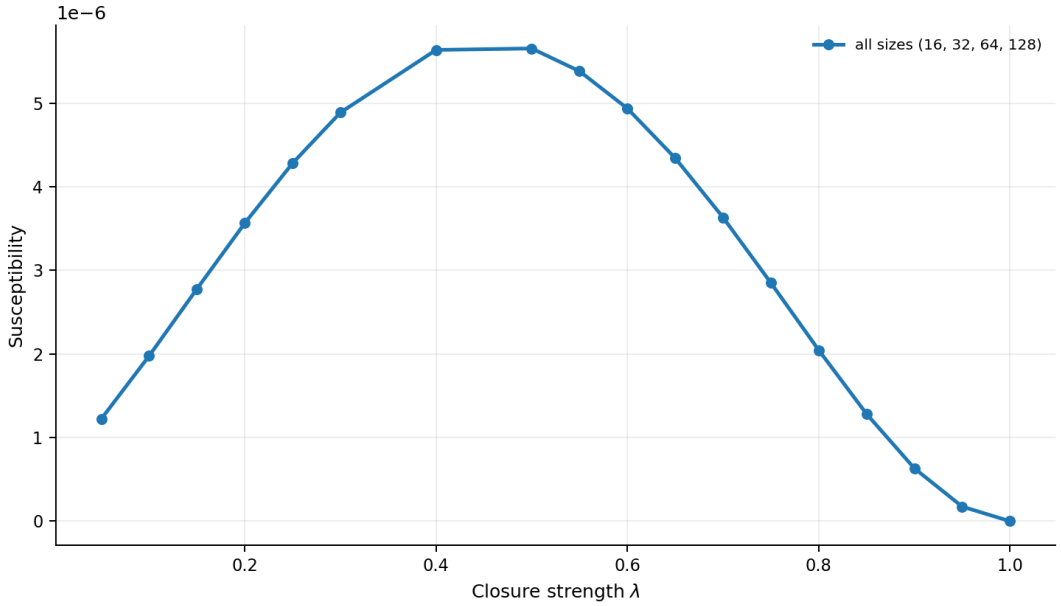

Figure 2. Manual-channel susceptibility panel for the Class-III shadow ensemble, using the tau-separated objecthood difference as the primary phenomenology observable.

This is also the place where the paper must be explicit about robustness limits. Under non-manual lens swaps, neither Class-III nor Class-IV reproduces its full canonical label cleanly. For Class-III, both the spectral and diffusion lenses collapse the representative back toward the Class-I side while still preserving the basic structural/no-drive signature. For Class-IV, the same lens swaps degrade the result further, yielding unclassified or Class-II-like outcomes as the P4 channel and, in some cases, the drive-confidence channel weaken. Table 7 summarizes this caveat. The point is not that the Class-III / Class-IV results disappear under perturbation; the point is that the confirmed claim belongs specifically to the canonical manual channel, and that this limitation should be carried honestly into the paper.

That caveat does not erase the result. What the present evidence supports is the following narrower but still strong statement: the Six-Birds primitive taxonomy has experimentally confirmed $P4$-active classes in the canonical manual analysis channel, and those classes can be separated both from one another and from the lower-complexity classes in measured observable space. What the

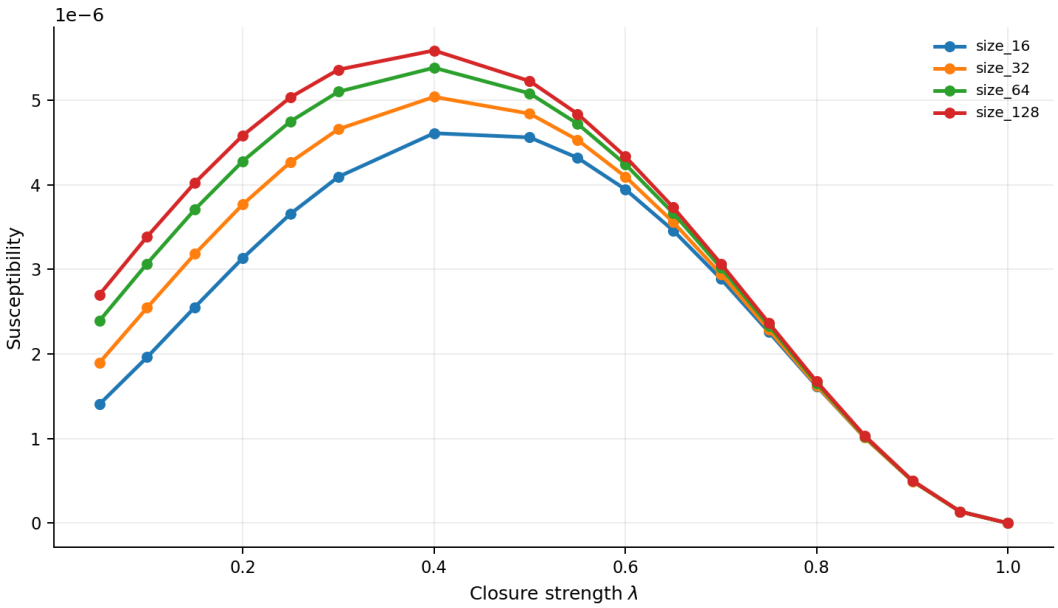

Figure 3. Manual-channel susceptibility panel for the Class-IV shadow ensemble. The panel is used as a phenomenology asset rather than as a stabilized exponent claim.

evidence does not yet support is the stronger claim that Class-III and Class-IV have the same level of lens robustness maturity as Class-I and Class-II. The section is therefore best read as a confirmation of the P4-active classes together with a quantified robustness boundary.

6 Four-class observable map and supporting assets

The central synthesis asset of the paper is the four-class observable map, shown in Figure 4. The purpose of this figure is not to present a final universality-exponent phase diagram. It is to show that the four confirmed classes are already visibly and quantitatively separable in a common measured observable space. The coordinates are deliberately operational: the horizontal axis is the structural reference boundary

$$\lambda_{\mathrm{struct,ref}} = \max\bigl(\lambda_{\mathrm{CE}=0.025}, \lambda_{M_{\mathrm{obj}}=0.90}\bigr),$$

the vertical axis is the clipped logarithm of affinity at that boundary, and the color channel is the strict-rule-aligned staging proxy

$$P4_{\mathrm{dual}} = \min\bigl(\Delta\lambda_{\mathrm{CE}}, \Delta\lambda_{M_{\mathrm{obj}}}\bigr).$$

Read in this way, the figure is a measured observable-space map of the primitive taxonomy.

Table 8 lists the corresponding class coordinates. Class-I sits at late structural birth with effectively zero affinity and low $P4_{\mathrm{dual}}$. Class-II shifts structurally earlier and acquires a robust nonzero affinity while remaining $P4$-inactive. Class-III remains effectively drive-free but moves sharply upward in the staging channel, with $P4_{\mathrm{dual}} \approx 0.202$. Class-IV combines both effects: it sits on the driven side of the affinity axis and on the staging-anomalous side of the $P4$ channel. In this sense the observable map is not merely a visualization convenience. It is the empirical object that makes the primitive class taxonomy visible in a single coordinate system.

The map also clarifies what is and is not doing the classificatory work. The structural boundary coordinate is needed because all four classes are still births of packaged structure. The affinity

Table 8. Representative observable coordinates for the four confirmed classes.

Class	Representative	λ_{CE}	$\lambda_{M_{\mathrm{obj}}}$	$\lambda_{\mathrm{struct,ref}}$	$\mathrm{Aff}_{\mathrm{ref}}$	$P4_{\mathrm{dual}}$
I	`class_i_reference`	0.874	0.763	0.874	≈ 0	0.062
II	`class_ii_reference`	0.556	0.500	0.556	3.45×10^{-2}	0.000
III	`replicated_portal_sp4`	0.562	0.123	0.562	≈ 0	0.202
IV	`repl._portal_sp4_drive_f`	0.533	0.067	0.533	2.99×10^{-3}	0.216

coordinate separates the drive-inactive classes from the drive-active ones. The staging coordinate separates the lower-complexity classes from the staging-anomalous ones. No one coordinate is sufficient on its own. The taxonomy works because the three channels together separate the four corners of the experimentally confirmed low-complexity space. Figure 5 gives the same point axis by axis, stripped of the two-dimensional map geometry.

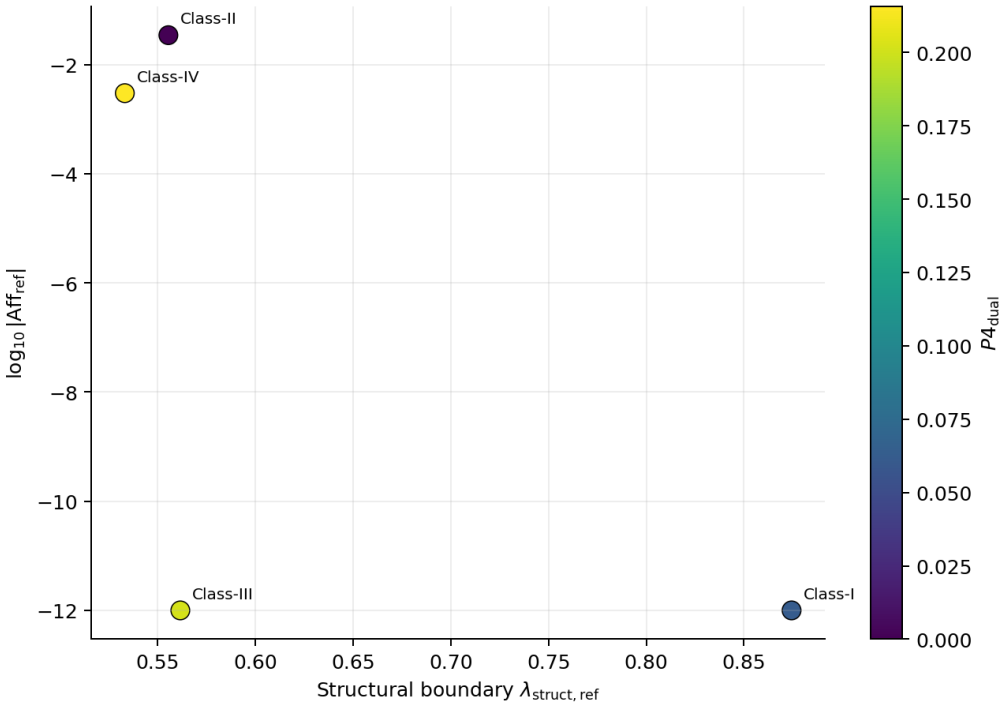

Figure 4. Measured observable-space map of the four confirmed classes. The horizontal axis is the structural reference boundary, the vertical axis is clipped logarithmic affinity at boundary, and the color channel is the strict-rule-aligned staging proxy.

This observable-space synthesis is stronger than the current exponent story. The finite-size analyses performed earlier in the project showed that exponent extraction remains too grid-sensitive to carry the main narrative. By contrast, the present map does not depend on unstable collapse minima. It is built directly from the measured structural, drive, and staging observables that define the class rubric itself. That is why this paper presents the map as the main synthetic payoff and treats exponent-level universality as future work rather than as a completed result.

Capacity and depth-growth then enter naturally as supporting assets rather than as class-defining coordinates. The capacity dashboard, summarized in Figure 6, shows that all four classes admit a meaningful post-birth comparison in terms of lambda headroom, depth survival, and a composite

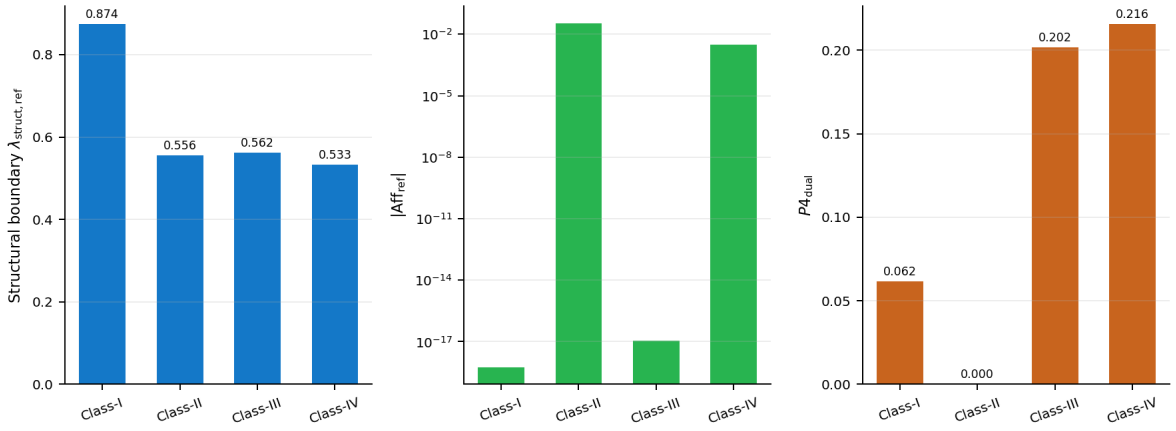

Figure 5. Axis-by-axis view of the structural, drive, and staging coordinates underlying the four-class observable map.

capacity-depth index. But the same dashboard also shows why capacity is not promoted to a primary class axis here. The source verdict is mixed: the signal does not align cleanly enough with either the $P4$ axis or the $P6_{\text{drive}}$ axis to function as a canonical classifier. That is a useful scientific result. It says that capacity/depth is best treated as a post-birth field—important for what happens after birth, but not for deciding which birth class one is in.

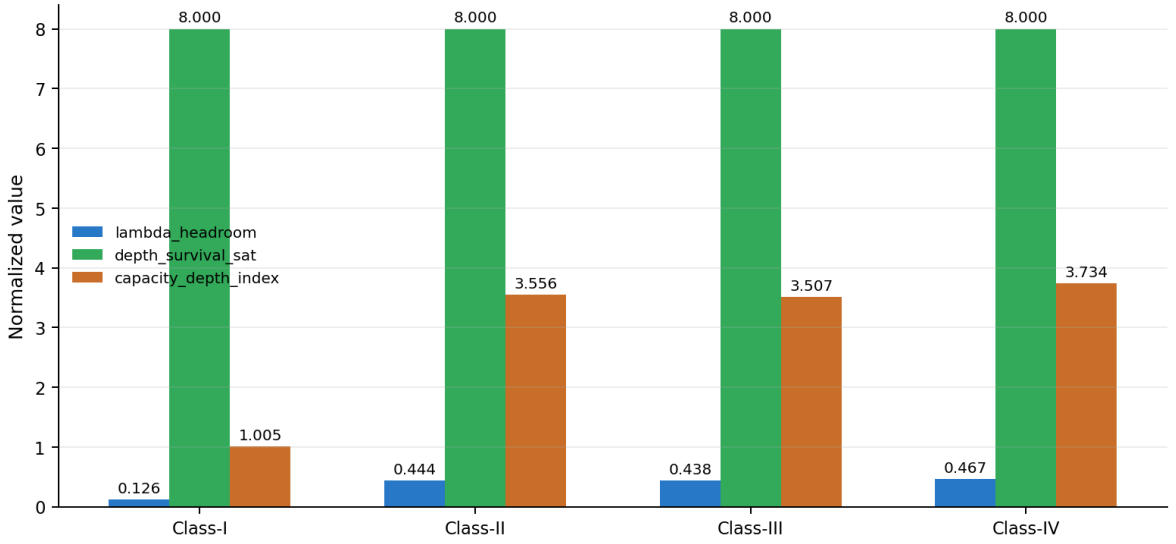

Figure 6. Representative post-birth capacity/depth comparison across the four confirmed classes. The signal is useful as a post-birth field but does not behave like a primary classifier.

The upstream bridge plays a different supporting role. Figure 7 places the best upstream/PICA-style candidate on the same observable map as the four synthetic classes. In the present bridge protocol, the successfully evaluated upstream candidates all land on the Class-I side with high operational confidence. That does not weaken the taxonomy. On the contrary, it shows that the class coordinate system is meaningful outside the closed world of synthetic benchmark families. The upstream bridge is therefore an external-validity asset: it demonstrates that at least one real

vendored upstream-style system can be measured and placed in the same observable language, even when it lands only on one side of the current taxonomy.

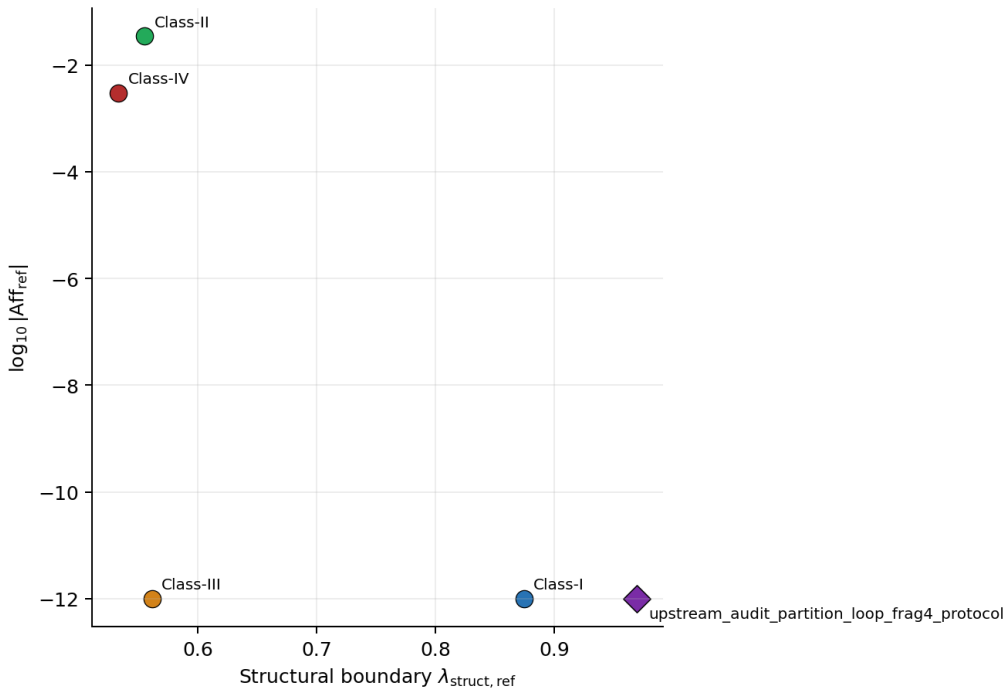

Figure 7. Best upstream/PICA-style bridge candidate placed on the same observable map as the four synthetic classes. The bridge provides external-validity support without requiring that all classes appear upstream.

Taken together, these supporting assets sharpen the scope of the paper. The four-class observable map is the main synthetic result. Capacity/depth shows how the confirmed classes compare as post-birth objects. The upstream bridge shows that the observable coordinates are not internal bookkeeping for synthetic families alone. None of these assets depends on the paper claiming final universality exponents. Their value is precisely that they make the primitive taxonomy empirically legible before that harder asymptotic step has been settled.

7 Discussion

The main result of this paper is not a final exponent-level universality classification. It is a primitive taxonomy and observable separation of layer-birth classes. Within the Six-Birds framework, the experimentally accessible low-complexity births are organized by activation of $P5$, $P6_{\mathrm{drive}}$, and $P4$, and the four corresponding canonical classes can be instantiated and separated in a common measured observable space. That is the level at which the evidence is now strongest, and it is the level at which the paper should be read.

The most mature part of that result is the Class-I / Class-II contrast. Both classes exhibit structural birth, but only Class-II carries a robust directed-circulation signal at structural birth. The empirical separation is therefore not merely taxonomic by declaration. It is enforced by a validated observable channel: affinity. The no-fake-arrow controls make that point especially important. Reversible controls do not get spuriously promoted into the driven class, while hidden-schedule protocol-trap constructions do produce the expected driven-positive signal under naive handling.

This is precisely the kind of evidential discipline that a primitive taxonomy needs if it is to be more than nomenclature.

The Class-III / Class-IV results add a second class-defining axis: staging anomaly. Here the paper had to become more careful. The P4-active classes are experimentally real and are confirmed in the canonical manual channel, but they are not yet as mature under analysis perturbations as Class-I and Class-II. In particular, the lens-robustness suite shows that non-manual lenses can degrade the full canonical labels for the P4-active representatives. We therefore treat those classes as confirmed but more fragile objects. That is a quantified robustness boundary, not an excuse. It tells us something substantive about the empirical delicacy of staging-anomalous births.

One of the most important conceptual clarifications produced by the project is the distinction between *$P4$-like signal* and *canonical $P4$ activation*. The lower-complexity classes already show measurable tau sensitivity at the signal level, but that does not collapse them into the staging-anomalous classes. Conversely, the confirmed Class-III and Class-IV systems clear the stricter dual-channel criterion and therefore justify a class-level staging label. This distinction is not just a technicality in the metric layer. It is what keeps the taxonomy from becoming either too coarse or too trigger-happy.

The negative results matter as much as the positive ones. The holonomy program is the clearest example. The project explored whether route mismatch could function as a crossover field and found that it could not be cleanly isolated as such. That is not a failure of experimentation. It is an empirical confirmation of a structural constraint already implicit in the primitive framework: $P3$ by itself is not a directionality certificate [13, 6]. Whenever the holonomy route looked promising, it did so by coupling into the drive channel. The correct lesson is not that the experiments came up short. The correct lesson is that the primitive theory was right to keep holonomy diagnostic rather than class-defining.

The finite-size scaling work produced a similarly useful negative constraint. Collapse fits can be computed, but the current exponent estimates are too grid-sensitive and too boundary-attached to support a strong universality-exponent claim [4, 2, 1]. In that sense the project learned what the paper is *not*. It is not yet the paper that settles whether the four birth classes are distinct universality classes in the strict asymptotic sense. The stability audits argue against making that claim now. What they do support is a more disciplined sequencing of the science: first identify the class-defining primitive activations and the observable map that separates them; only then attempt a stronger asymptotic universality analysis.

That limitation does not weaken the present paper. In some respects it strengthens it. The observable-space map is a more reliable synthetic asset than the current exponent fits because it is built from the measured class-defining observables themselves rather than from unstable collapse minima. Likewise, the affinity result is stronger than any present exponent claim. It is large, mechanically interpretable, empirically validated, and orthogonal to the projector-dominated structural channel in the baseline analysis. This makes the paper's main message cleaner rather than narrower: layer births are already classifiable before the final asymptotic story is settled.

The supporting assets sharpen that interpretation rather than distracting from it. Capacity/depth behaves as a post-birth field rather than a primary classifier, which is exactly where it belongs at this stage. The upstream bridge shows that the observable coordinates are meaningful outside the closed world of synthetic benchmark families, even though the current bridged candidates land on the Class-I side. And the phenomenology panels for Class-III and Class-IV show that the P4-active classes admit more than a yes/no structural confirmation in the manual channel, even if those panels are not yet the basis for a strong universality claim. One reason those P4-active classes remain especially interesting is that they may be the first place in the present taxonomy where conditions associated with balance-at-the-edge behavior can emerge in a class-defining way: once staging

structure is part of the birth identity itself, competing effective timescales can in principle become part of the post-birth organization rather than a small correction. That possibility is suggestive rather than demonstrated here, but it helps explain why the P4-active classes matter scientifically even before a final exponent-level story exists.

What remains open is therefore clear. First, the lens robustness of the P4-active classes should be improved beyond the canonical manual channel. Second, the asymptotic status of the classes—whether they correspond to distinct stabilized universality classes, or to a smaller number of exponent-level classes decorated by secondary fields—remains unresolved. Third, broader upstream coverage is needed if the taxonomy is to be claimed as more than a controlled synthetic-and-bridged program. These are real open problems, but they are now posed against a cleaner background than before. The paper leaves the field with a controlled rubric, a confirmed set of low-complexity classes, and an observable map that future work will have to respect.

8 Conclusion

This paper asked a classification question: how many primitive kinds of layer birth are experimentally accessible before exponent-level universality is settled? Within the Six-Birds framework, the answer supported by the present evidence is four. The confirmed low-complexity births are organized by activation of packaging ($P5$), directed circulation ($P6_{\mathrm{drive}}$), and staging anomaly ($P4$), yielding canonical Class-I, Class-II, Class-III, and Class-IV births.

The strongest empirical separation is between Class-I and Class-II. Both activate structural packaging, but only Class-II carries a robust drive channel at structural birth, and explicit no-fake-arrow controls validate that interpretation. Class-III and Class-IV extend the taxonomy by showing that staging anomaly can become class-defining, even though those P4-active classes remain methodologically more fragile under non-manual lens changes than the two lower-complexity classes. Taken together, the four classes are visibly and quantitatively separable in a common measured observable space.

The paper therefore does not conclude with a final exponent claim. It concludes with a primitive taxonomy and an observable map. Capacity/depth is shown to behave as a post-birth field rather than a primary classifier, and an upstream/PICA-style bridge demonstrates that the coordinate system is meaningful outside the synthetic benchmark families alone. The resulting picture is that layer birth is not one phenomenon but a small family of primitive-activation events, and that the experimentally accessible low-complexity cases are already rich enough to support a four-class observable taxonomy. Establishing the asymptotic universality structure of those classes is the next problem, not the one solved here.

A Metrics and rubric details

This appendix records the operational metric layer used by the paper. Its purpose is not to restate the full implementation-level documentation in the repository, but to make the class rubric and measurement thresholds explicit enough that the main text can remain conceptually focused. The paper's class taxonomy is defined by primitive activation, but the primitive activations are themselves inferred through a small set of observables and thresholds.

Structural birth is operationalized through two boundary proxies. The first uses closure error, with the reference threshold $\mathrm{CE} = 0.025$. The second uses the objecthood order parameter, with the reference threshold $M_{\mathrm{obj}} = 0.90$. Whenever both are available, the paper uses the conservative

Table 9. Operational thresholds used in the present class rubric.

Observable layer	Quantity	Threshold / rule
Structural birth	closure boundary	$\mathrm{CE} = 0.025$
Structural birth	objecthood boundary	$M_{\mathrm{obj}} = 0.90$
Structural reference	boundary location	$\max(\lambda_{\mathrm{CE}}, \lambda_{M_{\mathrm{obj}}})$
Drive activation	inactive threshold	$\mathrm{Aff}_{\mathrm{ref}} \leq 10^{-6}$
Drive activation	active threshold	$\mathrm{Aff}_{\mathrm{ref}} \geq 10^{-3}$
Staging signal	$P4$-like signal	any material shift or presence change
Canonical $P4$	strict activation	presence shift, or strong dual-channel shift
Observable-map staging proxy	$P4_{\mathrm{dual}}$	$\min(\Delta\lambda_{\mathrm{CE}}, \Delta\lambda_{M_{\mathrm{obj}}})$

structural reference boundary

$$\lambda_{\mathrm{struct,ref}} = \max\big(\lambda_{\mathrm{CE}=0.025}, \lambda_{M_{\mathrm{obj}}=0.90}\big).$$

This is the structural coordinate used throughout the main text and in the four-class observable map.

Directed circulation is operationalized by affinity at structural birth and by fixed affinity-onset thresholds. The main scalar is $\mathrm{Aff}_{\mathrm{ref}}$, the affinity evaluated at the structural reference boundary. Supporting onset markers are then defined at fixed thresholds, typically 10^{-3}, 10^{-2}, and 5×10^{-2}. These markers are used only as operational diagnostics. They do not themselves replace the class rubric.

Staging anomaly is tracked through tau/depth sensitivity in the structural boundary proxies [5, 3]. The paper distinguishes carefully between *$P4$-like signal* and *canonical $P4$ activation*. A measurable change in one structural proxy across tau is enough to count as $P4$-like signal. Canonical $P4$ activation is stricter. In the current rubric it requires either a presence/absence change in a structural boundary proxy or sufficiently strong coordinated movement in both structural channels. The staging quantity used in the observable map is

$$P4_{\mathrm{dual}} = \min\big(\Delta\lambda_{\mathrm{CE}}, \Delta\lambda_{M_{\mathrm{obj}}}\big),$$

which is intentionally aligned with the conservative class rule rather than with a more permissive notion of staging sensitivity.

Table 9 summarizes the key operational thresholds used in the paper. They should be read as the thresholds for the present taxonomy, not as universal constants. The paper's empirical claims depend on the reproducible use of these thresholds across the current dataset, not on the claim that no nearby thresholding choice could ever be scientifically meaningful.

A final point is worth making explicit. The paper's lower-complexity reference classes already show measurable staging sensitivity in the weak sense. That does not promote them into the staging-anomalous classes. The metric layer is therefore deliberately two-tiered: weak signal and canonical activation are not the same thing. This distinction is what allows the taxonomy to remain both sensitive and stable.

B Negative results and constraints

Several of the project's most useful outcomes are negative in form. They do not undermine the paper's main claims; they define their scope. This appendix records the most important of those constraints so that the main text does not have to carry them in full detail.

Table 10. Negative results and limiting constraints established during the project.

Question tested	Outcome	Constraint on the paper
Holonomy as crossover field	negative / constrained	$P3$ is not treated as an independent class axis
Stable exponent extraction	not supported	exponent claims are softened throughout
Non-manual lens robustness of Class-III / IV	partial / fragile	P4-active classes are confirmed in the manual channel with explicit robustness caveats
Capacity as primary classifier	not supported	capacity is treated as a post-birth field
Full upstream realization of all classes	not yet supported	upstream bridge is presented as external-validity overlay, not full class realization

The first is the holonomy result. The project tested whether route mismatch could function as a clean crossover field and found that it could not. The initial construction produced affinity contamination, and the refined construction eliminated the contamination without producing a meaningful boundary-shift effect. This should not be read as an unsuccessful side experiment. It is an empirical confirmation of a structural point already built into the primitive framework: $P3$ is diagnostically important, but it is not a directionality certificate on its own. Whenever the route-mismatch channel appeared to generate a class-relevant effect, it did so by coupling into the drive channel rather than by standing independently. The paper therefore treats holonomy as a constrained non-axis rather than a missing success.

The second is the exponent-stability result. The project explored finite-size-scaling fits for the lower-complexity classes and found that the resulting exponent estimates were too grid-sensitive and too boundary-attached to support a strong universality claim. Enlarging the fit grids did not stabilize the inferred parameters; in some cases it increased the degeneracy by allowing the optimizer to drift toward trivial boundary solutions. Adding one more size panel was not enough to repair the problem. The consequence is not that there is no size-dependent signal in the data, but that the current evidence does not justify strong exponent claims. For that reason the paper consistently treats exponent language as softened and secondary.

The third is the robustness asymmetry between the lower-complexity and staging-anomalous classes. Class-I and Class-II survive the standard lens perturbations cleanly enough to support a robust non-exponent class claim in the main text. Class-III and Class-IV do not yet have the same level of non-manual-lens maturity. Their canonical confirmation is therefore restricted to the manual channel, and their non-manual-lens behavior is carried as a quantified caveat. This is a limitation, but it is also a substantive measurement result: staging-anomalous births appear to be empirically more fragile under lens changes than the two lower-complexity classes.

The fourth is a more interpretive constraint. Capacity/depth is scientifically meaningful, but it does not behave as a primary classifier in the current data. Its signal is mixed or orthogonal relative to the $P4$ and $P6_{\mathrm{drive}}$ axes. This is why the paper uses capacity only as a post-birth field. Likewise, the upstream bridge is a successful external-validity asset, but the currently bridged upstream candidates land only on the Class-I side of the observable map. That is still useful, but it is not yet a full external realization of the whole four-class taxonomy.

Table 10 summarizes the main negative or limiting results. The point of collecting them here is not to apologize for what the paper does not do. It is to make clear which stronger claims were explicitly tested and not retained.

These constraints are not detachable from the positive results. They are what make the paper's final framing credible: the primitive taxonomy and observable-space separation are claimed strongly, while the stronger asymptotic and crossover-field claims are deferred.

C Reproducibility and evidence ledger

The project's final implementation state includes an explicit reproducibility freeze bundle. The purpose of that bundle is not only archival. It is also evidential: it records which repo assets support which paper claims, checks that the paper-critical assets are present and rerunnable, and distinguishes blocking issues from nonblocking caveats. The manuscript therefore rests on a frozen evidence ledger rather than on ad hoc references to intermediate experiment outputs. The code and experiment bundles for this manuscript are maintained in the public repository https://github.com/ioannist/six-birds-layer-birth.

The freeze bundle is located in the repository at

```
results/freeze/reproducibility_freeze/
```

and contains, among other files:

- a machine-readable evidence ledger,
- claim-coverage records,
- a reproducibility summary,
- and a ready-for-writing status record.

The key machine-readable files are:

```
analysis/evidence_ledger.json
analysis/claim_coverage.json
analysis/reproducibility_summary.json
analysis/ready_for_writing_status.json
```

The evidence ledger maps each paper-critical claim to one or more concrete bundle roots in the repository. In the final freeze state, the core claims are supported by the following primary assets:

- canonical class rubric and reference classification,
- Class-I / Class-II consolidation dashboard,
- four-class observable map,
- Class-III strict confirmation,
- Class-III / Class-IV phenomenology panels,
- capacity postcritical dashboard,
- and the upstream/PICA bridge.

This means that the manuscript's main experimental claims can be traced to stable bundle roots rather than to transient working notes.

Table 11 gives the compact claim-to-asset map used in the final freeze. The table is not a substitute for the full machine-readable ledger, but it makes the manuscript-level evidence structure visible inside the paper itself.

In the final corrected freeze state, the paper-critical assets are marked reproducible and the repository is marked ready for writing. The remaining caveats are therefore scientific rather than logistical. They include, for example, the partial non-manual-lens fragility of the P4-active classes and the deliberate softening of exponent-level claims. Those caveats matter for interpretation, but they do not block the manuscript's experimental basis.

Table 11. Compact manuscript claim-to-asset map used in the reproducibility freeze.

Claim ID	Primary supporting asset(s)
CLAIM-01	canonical class rubric; P4 anomaly metric layer
CLAIM-02	Class-I / Class-II consolidation dashboard; no-fake-arrow controls
CLAIM-03	four-class observable map
CLAIM-04	Class-III strict confirmation bundle
CLAIM-05	Class-III and Class-IV shadow-panel phenomenology bundles
CLAIM-06	capacity postcritical dashboard
CLAIM-07	upstream/PICA bridge bundle

The practical role of this appendix is simple: it tells the reader that the paper is not resting on hand-selected screenshots or non-reproducible intermediate analyses. The claims in the main text are backed by an explicit, frozen ledger. That is the standard against which the present manuscript should be read.

References

[1] Vincent Ardourel and Sorin Bangu. Finite-size scaling theory: Quantitative and qualitative approaches to critical phenomena. *Studies in History and Philosophy of Science*, 100:99–106, 2023.

[2] Kurt Binder. Critical properties from monte carlo coarse graining and renormalization. *Physical Review Letters*, 47(9):693–696, 1981.

[3] Nicolò Defenu. Metastability and discrete spectrum of long-range systems. *Proceedings of the National Academy of Sciences*, 118(30):e2101785118, 2021.

[4] Michael E. Fisher and Michael N. Barber. Scaling theory for finite-size effects in the critical region. *Physical Review Letters*, 28(23):1516–1519, 1972.

[5] Claudio Landim. Metastable markov chains. *Probability Surveys*, 16:143–227, 2019.

[6] Shiling Liang and Simone Pigolotti. Thermodynamic bounds on time-reversal asymmetry. *Physical Review E*, 108(6):L062101, 2023.

[7] Andreas Mardt and Frank Noé. Progress in deep Markov state modeling: Coarse graining and experimental data restraints. *The Journal of Chemical Physics*, 155(21):214106, 2021.

[8] Alexandria C. Marino and Brian J. Scholl. The role of closure in defining the "objects" of object-based attention. *Perception & Psychophysics*, 67(7):1140–1149, 2005.

[9] J. T. M. Miller. Are all primitives created equal? *The Southern Journal of Philosophy*, 56(2):273–292, 2018.

[10] Nir Nitskansky, Kessem Clein, and Barak Raveh. Building multiscale Markov state models by systematic mapping of temporal communities. *Bioinformatics*, 42(1):btaf585, 2026.

[11] Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, 75(12):126001, 2012.

[12] Ioannis Tsiokos. Six birds: Foundations of emergence calculus, 2026. arXiv:2602.00134 [cs.LO]. DOI: https://doi.org/10.48550/arXiv.2602.00134.

[13] Ioannis Tsiokos. Six birds protocol trap: Holonomy without entropy production, 2026. Zenodo preprint. DOI: https://doi.org/10.5281/zenodo.18906902.

[14] Ioannis Tsiokos. To create a stone with six birds: Emergent geometric and thermodynamic regimes from a minimal stochastic substrate, 2026. Zenodo preprint. DOI: https://doi.org/10.5281/zenodo.18838994.

[15] Ioannis Tsiokos. To lay a stone with six birds. *Preprints.org*, 2026. Preprints.org preprint. DOI: https://doi.org/10.20944/preprints202602.1699.v1.

[16] Song Wang, Jun Wang, and Toshiro Kubota. From fragments to salient closed boundaries: An in-depth study. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 291–298, 2004.