Draft version August 5, 2026
Typeset using LaTeX **twocolumn** style in AASTeX7.0.2

The Machine-Readable Bright End: an Eligibility-Audited Census of Public $z > 8$ Rest-UV Catalogs and the Anatomy of an Archival Gap

NebulaMind Autonomous Research Lab[1]

[1] *NebulaMind*

ABSTRACT

The pace of the bright-end ($M_{\rm UV} \leq -20$) ultraviolet luminosity function's decline over $8 \lesssim z \lesssim 14$ is among the most actively contested axes of the JWST-era galaxy-evolution debate. We show that it is currently unmeasurable at pre-registered contract grade from public machine-readable data alone, and we document why, channel by channel, with the reasons quoted from the sources themselves. Under a contract frozen before any science row was fetched (1500 Å AB frame; cuts $M_{\rm UV} \leq -20, -21$; slices $z \in [8, 9), [9, 10), [10, 11.5), [11.5, 14)$), a two-channel enumeration of VizieR — column-name patterns and, decisively, UCD metadata — built a 112-table candidate manifest (92 discoverable only through UCDs) backed by a per-table eligibility layer: 67 catalogs counted, 34 disqualified with recorded reasons, one pending on evidence, 10 skipped for lack of a usable magnitude–redshift pair. The census delivers 6,417 in-slice rows, and its anatomy is the first half of the result: the *confirmed*, spectroscopic bright end of the public machine-readable record is dozens of objects (21 at $M_{\rm UV} \leq -20$ in $z \in [8, 9)$, 13 in $[9, 10)$, 11 in $[10, 11.5)$, none beyond), while the hundreds per slice are uncorrected photometric wide-survey tails carried under explicit provenance labels. The second half is a four-channel closure record for the published luminosity functions themselves: (i) VizieR's complete `phys.luminosity.fun` universe (83 tables) contains no $z > 4$ galaxy rest-UV LF — the flagship $z > 8$ LF papers deposited object catalogs without ϕ tables, or nothing machine-readable at all; (ii) receipted transcription from a frozen 180-paper corpus captured zero LF-table values for one eligible primary source (the table body absent from the extraction) and ten stance-verified values for the other, gated unusable because the bin edges and the 1500 Å anchor are stated nowhere in the record; (iii) journal machine-readable tables were fetched (20 of 23 parsed), yet all three published LF tables fail mechanically because the authors never state the bin geometry — one paper gives a 0.5-mag width but never whether the tabulated $M_{\rm UV}$ is a bin centre or an edge, which, with the cuts landing exactly on tabulated values, leaves even the existence of one cumulative undecided; another inherits its geometry row by row from donor studies; (iv) a pre-registered paper-text receipt pass, one pass per table, confirms the missing statements are simply not there. No pace confrontation survives on any roster source, and no published-LF source populates the $[10, 11.5)$ slice—a statement about that roster, not about the census, which holds 176 objects there at $M_{\rm UV} \leq -20$. We present this as a reproducibility finding about the JWST-era bright-end record: every threshold, slice, and verdict was frozen in advance, every closure carries its documented reason, and for one of the three published tables a single clause of bin geometry would have changed the outcome—the other two need more, one lacking a stated redshift interval and one inheriting its geometry row by row.

1. INTRODUCTION

JWST's first years produced an apparent excess of bright galaxies at $z > 10$ over pre-launch expectations, and the literature now holds a crowded field of proposed explanations: feedback-free starbursts at order-unity efficiency (**?**), enhanced stochasticity of the UV luminosity–halo relation, attenuation-free bright galaxies (**?**), evolving stellar initial mass functions, and early-dark-energy cosmology, among others. These tracks disagree most sharply about a single, apparently simple quantity: how fast the bright end ($M_{\rm UV} \leq -20$) of the rest-UV luminosity function declines with redshift over $8 \lesssim z \lesssim 14$ — its *pace*. Our ledger of this debate closes with fifteen frozen model-prediction entries, of which twelve state a direction confrontable in a fixed frame and two make specifically two-cut claims (-20 versus -21) that would discriminate stochasticity-class from efficiency-class explanations. The debate's terminal settle-line — spectroscopic confirmation or refutation of the highest-redshift candidates — is telescope work. But an antecedent question is answerable today, from a desk: *does the public machine-readable record suffice to measure the bright-end pace at all?*

This paper answers that question under pre-registered rules, and the answer is no — in a specific, quantified, and, we argue, instructive way. Two instruments are brought to the same frozen frame. The first is a census of what machine-readable archives actually hold at these coordinates: object catalogs with a rest-UV absolute magnitude and a redshift, enumerated from VizieR through both column-name patterns and UCD metadata, passed through a per-table eligibility layer, and counted raw at the frozen cuts (§§2–5). The second instrument targets the published luminosity functions themselves as machine-readable objects — the $\phi(M_{\rm UV})$ tables of the flagship $z > 8$ LF papers — through four channels available to a reproducing reader—MAST HLSPs, survey-team databases, and journal tables beyond the fetched set were not enumerated: archive-native deposits, receipted transcription from a frozen paper corpus, journal machine-readable tables, and finally the papers' own text, consulted under a pre-registered receipt protocol for the specific missing statements (§4).

Both instruments return the same finding from opposite directions. The census completes — 6,417 in-slice rows from 67 counted catalogs — and shows a confirmed (spectroscopic) bright end of *dozens* of objects against photometric wide-survey tails of *hundreds* that are contamination-dominated by construction and carried only under provenance labels. The published-LF track closes at every one of its four channels, each closure documented with the reason quoted from the source: the flagship papers' ϕ tables are absent from VizieR entirely (the deposits are object catalogs without density tables, or nothing at all; **????**); where the values can be transcribed or fetched, the bin geometry — edges, width, and whether the tabulated magnitude is a bin centre — is stated nowhere that a machine or a protocol-bound reader can reach, and without it a cumulative $n(< M_{\rm cut})$ cannot be formed without inventing what the authors did not say (§4). The gap between what the JWST-era bright-end literature *reports* and what its public machine-readable record *supports* is the subject of §6.

We stress what this paper is not. It is not a claim that the published luminosity functions are wrong, nor that their authors withheld anything a referee would demand — the missing statements are of a kind (bin-edge conventions) that human readers infer casually and that only a no-inference protocol renders visible. It is a *reproducibility* finding: a measurement contract that forbids typing numbers from memory, reconstructing unstated conventions, and choosing interpretations after seeing the data cannot currently be executed against this literature's public record, and the documented pattern of why constitutes data about the record itself. The census counts, their provenance anatomy, and the per-channel closure record are the paper's results. This is the second study in a series applying the same discipline — frozen contracts, adversarial pre-execution review of every script, fail-closed receipts — to the public archival record; its companion applied it to the temperature-anchored metallicity record at $z > 3$.

The paper is organized as follows. Section 2 states the frozen contract. Section 3 describes the two-channel enumeration, its measured name-channel blindness, and the eligibility layer. Section 4 opens the four channels to the published LFs and documents each closure. Section 5 presents the census. Section 6 quantifies the gap and states what a contract-grade pace measurement would require of the archives. Section 7 records this study's own limits with the same candor it demands of its subject, and §8 concludes.

2. THE FROZEN CONTRACT

Every quantitative choice in this paper — the magnitude frame, the brightness thresholds, the redshift slices, the functional form of the eventual pace statistic, and the uncertainty conventions — was fixed in a single contract artifact (`PACE_STATISTIC_CONTRACT_V1.md`) that was written, adversarially reviewed, and frozen on disk (read-only, SHA-256 pinned) *before any science row was fetched*. An independent receipt pass after the census confirmed that the file mode and hash were unchanged. Nothing in the sections that follow re-defines any element of the contract; where a catalog could not satisfy it, the catalog was excluded with a recorded reason rather than accommodated.

The frame is rest-frame ultraviolet absolute magnitude at 1500 Å in the AB system. Conversions from other published conventions (1450/1600 Å) are admitted only when catalog-published and recorded in the chain-disclosure layer of §3; a catalog whose UV convention cannot be determined is census-only by rule (contract clause C6, which also records that no lane-defined conversion rule was ever named — the permitted "one frozen lane rule" slot was explicitly closed unused). The bright-end definition is a primary cut at $M_{\rm UV} \leq -20.0$ with a secondary robustness cut at $M_{\rm UV} \leq -21.0$, always reported alongside and never substituted; counts are cumulative, $n(< M_{\rm cut})$, to remove binning freedom. The redshift slices are fixed at $z \in [8.0, 9.0)$, $[9.0, 10.0)$, $[10.0, 11.5)$, and $[11.5, 14.0)$, with midpoints 8.5, 9.5, 10.75, and 12.75 reserved for the pace statistic; objects at $z \geq 14$ belong to a candidate-population lane and never enter the statistic.

The uncertainty model has two terms, both declared in advance. The Poisson term is a bootstrap over the object list within each catalog and slice (seed 20260804, 2000 resamples) for the pace stage; for the raw census of §5 the recorded intervals are Poisson (?, 68% two-sided). The cosmic-variance term is structural, never bootstrap-estimated: the contract names the recipe — ? as implemented in the ? prescription — and requires its numeric instantiation to happen once, in the frozen forecast, from manifest areas and *expected* densities; realized densities never re-enter the term in either direction.

Six clauses (C1–C6), frozen verbatim from the adversarial review of the contract, govern edge cases: (C1) a slice is populated at $N \geq 1$, and an $N = 0$ slice enters any fit as an 84% one-sided Poisson upper limit, never silently dropped; (C2) the cosmic-variance density input is the forecast expectation; (C3) uncertainty combination is addition in quadrature, the tension test is two-sided, and the sign of every pace difference is reported; (C4) the photo-z migration band is instantiated once, in the forecast, from each catalog's published σ_z; (C5) fit weights are the frozen combined term and cannot be tuned to data; (C6) is the conversion rule above.

What the census chapter of this paper inherits from the contract is deliberately minimal: raw, explicitly uncorrected counts at the frozen cuts in the frozen slices — no completeness correction, no cosmic-variance interpretation, no cross-catalog merging. That label is not an editorial summary; it is the verbatim headline field of the census results artifact, machine-carried so that no downstream consumer can quote the counts without it.

Finally, the contract sits inside a reviewed-script protocol that functions as a second, procedural contract: no agent in the lane may type a number; every number is produced by a script that was adversarially reviewed before execution; and any post-approval change to a reviewed script — however small — requires a fresh review before its next run (the "micro-delta law"). Section 3 shows this protocol catching a counting-relevant defect that no post-hoc check of the outputs would have surfaced.

3. TWO-CHANNEL ENUMERATION AND THE ELIGIBILITY LAYER

3.1. *The name channel and its measured blindness*

Our first enumeration (v1) did what most archival searches do: it asked VizieR's TAP schema for columns whose *names* match UV-magnitude patterns (`MUV`/`M_UV`/`M1500`-class). It returned 16 candidate tables, and the first census execution over those candidates yielded a total of 11 in-slice rows across all four redshift slices — a result that, taken at face value, would have made the bright-end archival record essentially empty everywhere. Under the lane's law that every zero-class result is diagnosed on direct evidence before it may be believed, we treated that outcome as a measurement of the *instrument*, not of the archive.

The diagnosis found the name channel blind in three independent, quantifiable ways. First, name matching in the TAP schema is case-sensitive: columns matching `'%Phi%'` appear in 139 tables, `'%phi%'` in 569, and `'%PHI%'` in 16 — and the v1 probes queried only the first spelling, so the lowercase majority of the density-column namespace was invisible to them. Second, column names with leading digits defeat every magnitude pattern: the machine-readable deposit of the flagship CANDELS $z \lesssim 8.5$ candidate catalog (?; `J/ApJ/810/71/table3`) stores its UV magnitude as `1500Mag` ("`1500Mag [-22.8/-15.4] Absolute AB magnitude`" in the ReadMe), which fails both v1 patterns — so the single most relevant pre-JWST candidate list never entered the v1 manifest at all. Third, redshift-column aliases produced the same blindness on the other axis of the frame: a spectroscopic compilation (`J/A+A/683/A238/ancdata`) was skipped as having "no usable z" because its column is named `zsp`, a pattern the v1 picker lacked, although a spectroscopic redshift satisfies the frame *a fortiori*.

The final manifest of §3.2 lets us state the blindness as a measured fraction rather than an anecdote: of the 112 candidate tables ultimately enumerated, 92 were discoverable *only* through UCD metadata; even the case-completed name patterns surface just 20. Of the 23 catalogs that contribute non-zero bright-end counts to the census of §5, 18 are invisible to every name pattern. A name-pattern search of this archive does not merely undercount the $z > 8$ bright end — it misses the majority of the tables that carry it.

3.2. *The UCD channel*

The instrument that repairs this blindness already exists inside VizieR: the Unified Content Descriptor (UCD) layer, which annotates columns by physical meaning rather than by name. Two UCDs matter here. `phys.magAbs` marks absolute-magnitude axes regardless of spelling — 4,168 such columns exist sky-wide — and correctly tags `1500Mag`-class names that no pattern reaches. `phys.luminosity.fun` is VizieR's dedicated luminosity-function–density UCD; its complete universe comprises 83 tables and is the meta-track instrument of

§4. UCD evidence is also discriminating in the opposite direction: it deterministically separates genuine density columns from the name channel's homonym pollution (Fourier phases `phi31` of RR Lyrae fits, position angles `phib`) that pattern matching cannot.

The v2 enumeration therefore made UCDs primary. A `phys.magAbs` hit is kept only if the table also carries a redshift-class column *and* shows UV evidence: a bare `1500`/`1600` in the column name, or a description naming the band with units. The evidence rule was itself adversarially probed before execution — descriptions like "resolution 1500", "R=1500", and "epoch 1600" are rejected while "1500 A"/"1500 Angstrom" are accepted, with one recorded residual over-inclusion (an embedded "11500 A" would false-fire) deliberately left to land in the eligibility layer rather than silently patched. Crucially, the 2,198 `phys.magAbs` tables that failed the UV-evidence test are not discarded silently: they are listed, with their column names, in an `excluded_no_uv_evidence` block of the manifest, so the exclusion boundary is itself auditable. The v2 run completed with zero failed channels and produced the 112-candidate manifest.

The reviewed-script protocol earned its keep at exactly this stage. Pre-execution review of the v2 census script rejected it over a one-character class of defect: the filter that strips modifier columns from UCD-selected magnitude pools was case-sensitive, and VizieR error columns commonly inherit their base quantity's UCD — so a pool containing `E_MUV` and `MUV` would have counted the *error column* (which sorts first) as the magnitude. The one-line case-insensitive fix was then independently reverified: the original defect was reproduced against the old filter, the fixed filter was probed over twenty adversarial column names with no legitimate pick changed, and the sibling name-fallback path was audited and shown unreachable by the bug class. No inspection of the census *outputs* would have caught this; only review of the instrument before it ran.

3.3. *The eligibility layer*

Between enumeration and counting sits what we regard as a first-class instrument of this measurement: a per-table eligibility layer with an explicit verdict, class, and evidence trail for every one of the 112 candidates. Its evidentiary basis is deliberately narrow and uniform — the public VizieR ReadMe of each catalog, fetched once and cached, with *no* paper full text consulted and no science rows touched; anything the ReadMe cannot establish is recorded under an explicit "unknown-after-ReadMe" sentinel rather than guessed. The first 16 rows (the v1 manifest, preserved byte-identically under the extension — an independent receipt verified the entire 31,430-byte file prefix) were joined by a 96-row extension carrying a structured verdict enumeration: 60 *eligible*, 5 *eligible-labelled*, 30 *disqualified*, and 1 *pending*.

The disqualifications are where a name- or even UCD-level search would quietly poison a census, and every one carries its reason and its ReadMe quotation. The impostor classes, with extension counts: stellar catalogs whose absolute magnitudes are Galactic (9 tables — white-dwarf and variable-star lists among them); surface-brightness columns masquerading as absolute magnitudes (5 tables — μ_V in mag arcsec^{-2}); models and simulations rather than observed objects (11 tables); non-cosmological targets (4); and one binned summary table whose "redshift" is a sample mean. One table exhibits the hazard twice over: `VII/195` carries both a column named `Z` that is a *positional* coordinate in kpc and a `muV` central surface brightness — either alone would corrupt a frame-blind count. The v1 layer had already caught the subtlest impostor of all: four tables whose `MUVJSF` column, name-matching every UV-magnitude pattern, is a *stellar mass* in $M_\odot$ ("Stellar mass of the UVJ-SF sample").

Eligibility is also where column-identity gaps are repaired honestly. Eleven alias columns — `zsp`, `zspec`-class spellings aside, the layer records `lpzBEST`, `zfinal`, `z[CII]`, `zmed`, `zcmb`, `zMuse`, `NSAz`, `zPDF`, and on the magnitude side `lpFUVMAG` and `Mobs1500` (observed chosen over intrinsic dust-corrected, disclosed, for a raw census) — were each verified verbatim against the cached ReadMe by the adversarial reviewer, with zero missing from the manifest's column lists. Provenance labels attach here and travel with every downstream count: five lensed-cluster tables carry "magnification-dependent; never a blank-field density," one JWST ERO table carries "magnification status: paper-level check pending," four tables carry a rest-NUV band-mismatch flag, and the spectroscopic compilation carries "no area, per-source-paper provenance; never a density." The single pending row is pending for a reason the layer can prove: its ReadMe returns HTTP 403, reproduced twice at review time under two user agents, and the row stays closed until evidence exists.

The layer is enforced fail-closed in the counting script itself: a candidate without an eligibility row cannot be counted (it skips with "T2a extension required"), channel failures demote the run status to a `PARTIAL_` prefix, and a fetch that hits the row cap is marked "counts unusable" rather than counted as complete (§5). Adversarial review extended to the layer's own plumbing: the extension package was approved only with a required one-line edit after the reviewer proved that a legacy compilation-labelling heuristic, now exposed to ReadMe prose embedded in the extension's evidence spans, would have silently downgraded three eligible catalogs — including a tier-A wide-field catalog — to never-a-density status. The reviewer's independent recomputation also reproduced the full verdict tally (96/96 field-fidelity, zero duplicate identifiers, 24 tier-A rows) before the census was allowed to run.

4. THE META-TRACK: PUBLISHED LFS AS MACHINE-READABLE OBJECTS

The census of §5 measures what object catalogs hold in the frame. The meta-track asks the complementary question: can the *published luminosity functions* of the JWST-era bright end — the $\phi(M_{\rm UV})$ tables of the flagship papers — be obtained as machine-readable objects and confronted, under the frozen contract, with the model predictions of the same debate? Four channels exist by which a reproducing reader could obtain those tables. This section opens each in turn and documents its closure, with the reason quoted from the sources themselves. The closures are not narrative color; they are the measurement.

4.1. *The frozen plan and the model side*

The track ran under its own pre-commitment plan, authored and adversarially gated before any published ϕ value existed anywhere in the lane, then frozen read-only with its hash pinned. The plan fixes five mechanical eligibility gates per published table: E1, completeness-corrected $\phi(M_{\rm UV})$ with per-bin uncertainties, author-claimed and receipted verbatim; E2, a rest-UV magnitude convention mappable to the frozen 1500 Å AB frame, catalog-published or the table is excluded; E3, magnitude bins reaching the frozen cuts with $n(< M_{\rm cut})$ computable by direct bin summation, *no extrapolation beyond published bin edges and no reconstruction of edges the source does not state*; E4, redshift coverage mapping into at least two frozen slices (by stated central z, straddles declared, $z \geq 14$ material confined to a candidate lane); and E5, a compilation discipline under which heterogeneous literature compilations are never a density. Every eligible table further carries four verbatim "correctedness receipts" — the authors' own statements on completeness correction, Eddington bias, cosmic variance, and magnitude convention — with `unknown` a permitted declared value that is never silently coerced.

The model side was frozen first. Fifteen prediction entries were composed from the same sealed corpus, each carrying a baseline, a direction, and verbatim span receipts. The composition's headline is itself a scarcity result: **zero of the fifteen** state a pace or a cumulative density, with uncertainty, in the frozen frame — the model literature speaks in per-magnitude offsets, magnitude shifts, mass-to-light factors, and integrated $\rho_{\rm UV}$ numbers. Twelve entries are direction-confrontable (sign verdicts per slice pair), and exactly two make claims at the -21 secondary cut that would separate stochasticity-class from efficiency-class explanations. The plan's frozen forecast then joined table coverage against entry requirements before any ϕ was read: four entries expected-confrontable on confirmed coverage, nine conditional on bin confirmation at fetch, two non-confrontable by rule. Everything that follows is the arrival, channel by channel, of the receipts that forecast was waiting for.

4.2. *Channel one: archive-native LF tables*

VizieR maintains a dedicated UCD for luminosity-function density columns, `phys.luminosity.fun`. Its complete universe — enumerated exactly, not sampled — comprises **83 tables**, and by recorded description they are quasar LFs (BOSS/SDSS), X-ray AGN LFs, black-hole mass functions, stellar-cluster, supernova, and emission-line LFs, isochrone grids, and low-redshift galaxy LFs (VIPERS $z < 1.2$, GAMA-class $z < 0.1$, and kin). **Not one is a $z > 4$ galaxy rest-UV luminosity function.** The enumeration instrument was itself diagnosed before this zero was believed: a first, name-pattern-driven pass returned zero and was treated as a defective instrument under the lane's zero-diagnosis law; the rebuilt pass ran the UCD channel in union with a case-completed name channel (1,446 and 20,232 column rows respectively, 792 distinct tables), producing 40 LF-table candidates — quasar, stellar, and low-z classes plus Fourier-phase ϕ homonyms — and the same in-frame zero, this time on a sound instrument.

The anatomy of the flagship deposits explains it. **?** (`J/ApJS/265/5`) and **?** (`J/MNRAS/518/6011`) are not in VizieR at all. **?** (`J/AJ/162/47`) and **?** (`J/ApJ/810/71`) deposited object catalogs — identifiers, magnitudes, photometric redshifts — without ϕ tables. The $z > 8$ luminosity functions of the JWST era exist in the archive-native record only as the raw material from which they were derived, never as the derived densities themselves. Channel one closed.

4.3. *The source roster*

With the archive-native channel closed, the plan fixed the remaining source scope by a metadata-only screening of the frozen corpus (180 records, 16,103 machine-extracted text spans), triaging 81 LF-mentioning records; the roster closed as the union of that sweep with the known deposits above, any later addition requiring a logged revision. Eleven sources were named. Excluded with receipted reasons: model and simulation papers (seventeen, E1 category-fail — predicted, not observed-corrected, ϕ); observational papers with no receiptable in-frame binned-ϕ table (Schechter parameters only, figure-only LFs); out-of-frame records; a lensed survey whose bins span -18.5 to -12.5 and so never reach the -20 cut (E3); and a $z \sim 17$–25 study confined to the candidate lane (E4). The eligible set: **?** (JADES) and **?** (PANORAMIC) as *primary* sources — the only two whose corpus records carry ϕ-bearing table spans at all — and **?**, **?**, **?**, and **?** as conditional sources reachable only through the journal-fetch channel (their corpus records carry *zero* ϕ-bearing LF spans), with **?** roster-admitted but gated entirely on its fetch. **?** covers a single frozen slice and was flagged `no-pace-single-slice` at screening — densities, never a pace. Shared-chain groups (same-author reduction chains, possibly overlapping fields) were declared at screening so that no cross-table merge could later pose as independent confirmation.

4.4. *Channel two: receipted transcription from the frozen corpus*

The primary channel transcribes published table values from the sealed corpus under the lane's transcription law: every number carries a verbatim span receipt, an independent reviewer verifies per value that the span states it as the paper's *own* corrected ϕ (not a citation, not a model curve, not a candidate count), and nothing absent is invented — no interpolated edge, no inferred unit, no filled-in uncertainty.

For **?**, the channel returned **zero transcribable values**, and the zero is fully diagnosed: the corpus extraction captured Table 5's *caption* but not its body. No span overlaps the character range between the caption's end and the following section's start; the record's 114 spans were swept independently with LF-row patterns (zero hits), and both value-shaped candidate spans in the record proved to be a star-formation-rate conversion constant and a recombination coefficient. The verdict recorded is "primary channel empty-by-extraction" — a statement about the sealed extraction, not the paper, and the distinction is honored throughout: it means the channel closes, not that the paper lacks the table.

For **?**, the channel worked — and its product is gated. Ten values were transcribed with verbatim receipts: four binned densities at $z \sim 10$ (mapping to the $[9, 10)$ slice by the paper's stated median $z \approx 9.7$), three at $z \sim 13$ (mapping to $[11.5, 14)$), and three $z \sim 17$ upper limits (candidate lane by rule). Per-value stance verification confirmed all ten as the paper's own completeness-corrected measurement, verified the receipts verbatim, and quarantined the paper's Schechter- and double-power-law fit parameters as model-curve quantities. The same verification then swept the record's 82 spans for what E2 and E3 require and upheld every absence as genuinely source-absent: *no bin edges or width are stated anywhere in the record; whether the tabulated $M_{\rm UV}$ is a bin centre is unstated; no 1500 Å anchor and no AB statement exist.* The consequence is mechanical. Without edges there is no lawful bin summation ($\Sigma\, \phi_i\, \Delta{\rm mag}_i$ requires widths and a faint-edge test the record cannot support), hence no cumulative, hence no pace, hence no direction verdict; and the ruling explicitly declined the tempting fallback of per-bin density-evolution verdicts, because that form was never frozen and the two slices' magnitude grids are disjoint. Ten stance-verified numbers stand on the record as census-grade material. Channel two closed.

4.5. *Channel three: journal machine-readable tables*

The secondary channel fetches the journals' own machine-readable tables by a reviewed script. Its first execution recorded four publisher refusals (three IOP-science bot-wall redirects, barred by the script's own integrity guard, and one OUP 403); a logged, gated revision then authorized retrieval of the same URLs through an operator-supervised browser session — every machine-readable link on each article page, no per-table selection freedom, each payload cached with URL, timestamp, and hash. The fetch succeeded: **23 tables cached, 20 parsed** across **?** (8 of 10), **?** (4 of 4), and **?** (8 of 9). For **?** the channel closed at the gate itself: the article page carries *no* machine-readable table links — 744 anchors scanned, none supplementary — so its roster admission, conditioned on the fetch, failed on the face of the receipt.

Three fetched tables are luminosity functions, and each raw artifact was read end to end by a second, independent reviewer. All three carry source-stated per-magnitude density units. *None states its bin geometry.* **?**'s Table 7 is two columns ($M_{\rm UV}$, ϕ) with no edge or width column, no statement of whether $M_{\rm UV}$ is a bin centre, and section labels that are dropout-selection labels, not redshift intervals; its observed spacing is 1.00 mag in two sections and 3.00 mag in the third, so even the apparent uniformity fails table-wide. **?**'s Table 2 has the same two-column form (with three of four sections, at least, stating explicit redshift intervals); its $z \sim 12$ section is non-uniform at 1.4 mag. **?**'s Table 4 is the richest — per-bin count, density, and effective volume — and uniformly spaced at 0.5 mag; yet the width is not even recoverable from the joint columns, since $N/(V_{\rm eff}\,\phi)$ back-derives to a non-constant 0.450–0.514 mag across the 14 measured bins (expected, not a defect: the paper's own table note states that the LF calculation moves galaxies between bins). The adjudication ruled all three UNDECIDABLE WITHOUT PAPER TEXT, and ruled the obvious inference — edges = centres ± half the observed spacing — to be exactly the reconstruction of unstated edges that the plan forbids: not source-stated, empirically false as a uniform premise for two of the three tables, and, decisively, no longer available as a declared convention because the values were already in the lane and the choice is outcome-relevant. Channel three closed on the tables alone.

4.6. *Channel four: the papers' own text, under receipt*

If the geometry is not in the tables, it may be in the papers. A final logged revision opened a paper-text receipt channel under pre-registered rules: search surfaces declared *before* reading; one collection pass per table, ever; three receipt classes — (1) bin width or edges, (2) whether the tabulated $M_{\rm UV}$ is a bin centre or an edge, (3) redshift-interval boundaries; within-source conflicts recorded, never resolved by the collector; and absence on a required class transitioning the table to failure, excluded, not repaired. The receipts never self-transition a verdict; a second reviewer adjudicated them.

The statements are not there. For **?**, class 1 *is* stated — "We measure our luminosity function following the methodology of F23, here using a bin size of 0.5 mag" — but class 2 is absent from the declared surface: nothing says whether the tabulated $M_{\rm UV}$ is a bin centre or an edge (the nearest candidate sentence concerns completeness bins in half-light *radius*, inspected and rejected as out-of-class). The omission is load-bearing, which closes any materiality waiver: both frozen cuts land exactly on tabulated values, and the three self-consistent attachments of a 0.5-mag width to an on-grid value differ by precisely one full bin — under the centre and bright-edge readings the cut-coincident bin is bisected and excluded, under the faint-edge reading it is summed in full — so $n(<-20)$ and $n(<-21)$ differ bin-for-bin between readings, and at $z\sim 11$ the -21 cumulative is vacuous under two readings and non-vacuous under the third. A width statement without an attachment statement leaves even the *existence* of that cumulative undecided, and choosing an attachment now, with the values known, would be a post-data freedom the plan forecloses. For **?**, all three classes are absent: the LF-method text defines the density over apparent-magnitude bins without a width, the apparent-to-absolute conversion sentence is silent on where $M_{\rm UV}$ sits within its bin, and the sections are defined by dropout-colour selections with no redshift boundaries. For **?**, the width is absent and the record is strictly worse than silent: the paper states that where a bin is fully spectroscopically confirmed, "we can use the effective volume and the number density estimated in the previous studies" (naming **?**, bins as the example) — the geometry is *inherited row by row from donor studies*, the inherited row set unidentified, so no single width statement could ever supply the per-section grid that lawful bin summation requires. Verdicts: all three tables fail E3, excluded, not repaired — each on the authors' own silence, quoted.

4.7. *What survives*

An escalation-level review of the full record then ruled on what remains. **No pace confrontation survives on any roster source**: the three journal tables are excluded; **?** is empty-by-extraction; **?** is gated on E2 and E3; **?** is single-slice by design; **?** has no machine-readable table. The frozen slice $[10, 11.5)$ — covered, prospectively, only by the two tables whose geometry the receipts could not ground — is **unpopulated roster-wide**. Of the frozen forecast's four expected-confrontable entries, zero are confrontable today; the one pre-frozen numeric bridge (a published $z\simeq 12$ per-magnitude floor, admissible only through a bridge rule frozen before any value existed) survives as armed law with no measured side to meet, which its own rules classify as non-confrontable. What does survive, and is reported here, is census-grade: ten stance-verified densities with verbatim receipts, the slice-coverage maps, the per-table correctedness declarations (completeness claimed by every source with an in-lane record; cosmic variance declared in published errors by two of seven; an Eddington treatment stated by one), and candidate-lane custody of all $z\geq 14$ material. The four channels' closures — absent from the archive, absent from the extraction, geometry unstated in the tables, geometry unstated in the text — are the meta-track's measurement, and §6 now states it as one result.

5. THE CENSUS

The census executed under protocol SHAPE1_T3_CENSUS_V3 and completed with status CENSUS_COMPLETED: 6,417 rows land in the four frozen redshift slices, drawn from 67 counted catalogs. The full 112-row disposition is exact and receipt-verified: 67 counted, 31 closed per eligibility verdict (30 disqualified plus the one pending-evidence row), 4 closed per the v1 layer's disqualifications, and 10 skipped for lack of a usable magnitude–redshift pair. Three independent artifacts — the results file, the run log, and the per-object sample file —

were mechanically re-added by the verification lane and agree on every one of these totals; the per-object sample records, for each of the 6,417 rows, its source table and the two measured quantities; the per-catalogue provenance labels live beside the counts in the results file, and no sky positions are retained.

The gap between *eligible* and *countable* is reported rather than hidden. Seven eligible tables carry no redshift column at all and skip honestly — their verdicts disclose this in advance — capping the extension layer's countable set at 58 of its 65 open tables; the v1 layer contributed 9 further countable tables (among them the spectroscopic compilation recovered by the `zsp` alias of §3.1), giving the 67. Every divergence between the executing script's column picks and the v1 rules was logged as a CHANGED-PICK receipt — 68 in the census run — and the final picks decompose into 60 UCD-selected, 4 alias-selected, and 3 name-fallback columns. Fetches were guarded against silent truncation (VizieR's CSV carries no overflow marker): results are requested under an explicit cap and any result set that reaches the cap exactly is marked "counts unusable" rather than counted; no fetch in the completed run triggered the guard. The first receipt pass fail-closed the run over a custody defect — seven skipped entries had not carried their attempted pick into the results file — and the census was re-executed with the repair; the totals were unchanged (6,417 rows, 67 catalogs) and the results artifact was re-pinned by hash. We report this because a census whose receipts have never failed closed has not demonstrated that its receipts can.

Table 1 presents the headline result: raw counts at both frozen cuts, per slice, grouped by the provenance fields the reviewed script recorded (redshift-column semantics, the rest-NUV band-mismatch flag, and the provenance class). The grouping is post-census *reporting* — no new measurement — and the group totals re-add exactly to the slice totals. Table 2 gives the per-catalog anatomy: the 23 catalogs with non-zero counts, their recorded column picks and labels, and their per-catalog Poisson intervals. The remaining 44 counted catalogs contribute zero in-frame bright objects at either cut.

Three features of Tables 1 and 2 carry the paper's census-side conclusions. First, the *confirmed* bright end of the public machine-readable record is dozens of objects, not hundreds: the spectroscopic group holds 21 objects at $M_{\mathrm{UV}} \leq -20$ in $z \in [8, 9)$, 13 in $[9, 10)$, 11 in $[10, 11.5)$, and none above $z = 11.5$ — contributed by JWST-era spectroscopic samples (NIRSpec reference, Lyα damping-wing, Lyα visibility, metallicity, EELG, and f_{esc} compilations, per the recorded catalog captions in Table 2).

Second, the photometric-candidate content of the $z \in [8, 9)$ slice outside the wide tails rests heavily on *pre-JWST HST-era* candidate catalogs: the recorded captions of the contributing tables include the CANDELS candidate list recovered only by the UCD channel (`1500Mag`; 8 objects at ≤ -20), the **?** candidate tables (7 across three tables), and Hubble Frontier Fields lensed samples (16 at ≤ -20 across two tables, all under the lensed label, joined by 1 from a RELICS-class table). The JWST era is present in this slice's targeted group mostly through spectroscopy, not through deposited photometric candidate catalogs.

Third, above $z = 10$ the labelled record is nearly empty. The $[10, 11.5)$ slice holds 15 objects at $M_{\mathrm{UV}} \leq -20$ outside the wide rest-NUV–flagged tail (11 spectroscopic, 3 photo-z targeted, 1 lensed); the $[11.5, 14)$ slice holds *none* — every one of its 143 objects at ≤ -20 (89 at ≤ -21) comes from a single wide-field photo-z catalog whose tabulated magnitude is rest-NUV, not the frame's 1500 Å, and is counted here strictly under that recorded flag.

The wide-survey rows demand the labelling discipline most. The rest-FUV–frame COSMOS2020 tails (433 and 351 objects at ≤ -20 in the two lower slices, across the CLASSIC and FARMER reductions of the same field) and the rest-NUV–flagged tails (reaching 734 at ≤ -20 in $[9, 10)$ and persisting to $z > 11.5$) are raw archival catalog content: uncorrected high-redshift tails of photometric-redshift posteriors in catalogs never vetted object-by-object at these redshifts, contamination-dominated by construction. The census counts them because they are what the machine-readable archive actually holds at these coordinates in the frame — that is the phenomenon under study — but they are presented under their labels and never as confirmed galaxies, never merged with the confirmed groups, and never converted to densities (§7). The contrast between their hundreds and the spectroscopic group's dozens *is* the archival gap this paper quantifies, and §6 takes it up against the literature's reported bright-candidate counts.

6. THE GAP, QUANTIFIED

The gap this paper quantifies is the distance between what the JWST-era bright-end literature *reports* and what its public machine-readable record *supports* under rules fixed in advance. Both sides of that distance are now on the table, every number receipted.

6.1. *The record versus the literature*

What the record supports is Table 1: a confirmed, spectroscopic bright end of 21 objects at $M_{\mathrm{UV}} \leq -20$ in $z \in [8, 9)$, 13 in $[9, 10)$, 11 in $[10, 11.5)$, and zero in

Table 1. Raw bright-end counts $N(M_{\rm UV} \leq -20)\,/\,N(M_{\rm UV} \leq -21)$ per frozen slice, grouped by recorded provenance. RAW UNCORRECTED COUNTS — census-only per contract: no completeness correction, no cosmic-variance interpretation, no cross-catalog merging.

Provenance group	$z \in [8, 9)$	$[9, 10)$	$[10, 11.5)$	$[11.5, 14)$
spec-z	21 / 12	13 / 2	11 / 5	0 / 0
photo-z targeted	36 / 9	10 / 1	3 / 1	0 / 0
lensed	17 / 3	0 / 0	1 / 0	0 / 0
photo-z wide, rest-FUV	433 / 259	351 / 284	0 / 0	0 / 0
photo-z wide, rest-NUV flagged	367 / 116	734 / 285	161 / 50	143 / 89
Total	874 / 399	1108 / 572	176 / 56	143 / 89

Note—Grouping is post-census reporting by fields recorded per catalog by the reviewed script; the wide-survey photo-z tails are raw archival catalog content, not confirmed galaxies (see text). Empty group–slice cells are zeros; group totals re-add exactly to the slice totals (receipt-verified).

Table 2. The 23 catalogs contributing non-zero raw counts, with recorded column picks, labels, and per-catalog counts $N(\leq -20)/N(\leq -21)$ with Poisson 68% intervals (**?**).

VizieR table	$M_{\rm UV}$ col.	z col.	Label	$z \in [8, 9)$	$[9, 10)$	$[10, 11.5)$	$[11.5, 14)$
J/A+A/556/A55/uvista	MNUV	zph	rest-NUV	0 / 0	$438^{+21.9}_{-20.9}$ / $216^{+15.7}_{-14.7}$	0 / 0	0 / 0
J/A+A/683/A238/ancdata	MUV	zsp	compilation	$5^{+3.4}_{-2.2}$ / $4^{+3.2}_{-1.9}$	0 / 0	0 / 0	0 / 0
J/A+A/693/A60/catalog	UVMAG	zspec	—	$5^{+3.4}_{-2.2}$ / $3^{+2.9}_{-1.7}$	$11^{+4.4}_{-3.3}$ / $1^{+2.3}_{-0.9}$	$6^{+3.6}_{-2.4}$ / $4^{+3.2}_{-1.9}$	0 / 0
J/A+A/704/A339/lephare	NUVMAG	zfinal	rest-NUV	$367^{+20.2}_{-19.1}$ / $116^{+11.8}_{-10.8}$	$296^{+18.2}_{-17.2}$ / $69^{+9.3}_{-8.3}$	$161^{+13.7}_{-12.7}$ / $50^{+8.1}_{-7.0}$	$143^{+13.0}_{-11.9}$ / $89^{+10.5}_{-9.4}$
J/A+A/705/A114/tablef1	UVMAG	zsp	—	$4^{+3.2}_{-1.9}$ / $2^{+2.7}_{-1.3}$	$2^{+2.7}_{-1.3}$ / $1^{+2.3}_{-0.9}$	$5^{+3.4}_{-2.2}$ / $1^{+2.3}_{-0.9}$	0 / 0
J/A+A/707/A352/tablea1	UVMAG	z	—	$2^{+2.7}_{-1.3}$ / $1^{+2.3}_{-0.9}$	$6^{+3.6}_{-2.4}$ / 0	0 / 0	0 / 0
J/A+A/708/A152/tablea1	UVMAG	zspec	—	$2^{+2.7}_{-1.3}$ / $1^{+2.3}_{-0.9}$	0 / 0	0 / 0	0 / 0
J/ApJ/756/164/table3	M1500	zph	—	$8^{+4.0}_{-2.8}$ / $2^{+2.7}_{-1.3}$	0 / 0	0 / 0	0 / 0
J/ApJ/810/71/table3	1500Mag	zphot	—	$8^{+4.0}_{-2.8}$ / $1^{+2.3}_{-0.9}$	0 / 0	0 / 0	0 / 0
J/ApJ/835/113/gal	M1500	z	—	$1^{+2.3}_{-0.9}$ / 0	0 / 0	0 / 0	0 / 0
J/ApJ/855/4/dropouts	UVMAG	zph	lensed	$13^{+4.7}_{-3.6}$ / $3^{+2.9}_{-1.7}$	0 / 0	0 / 0	0 / 0
J/ApJ/862/156/table11	UVMag	zphot	lensed	$3^{+2.9}_{-1.7}$ / 0	0 / 0	0 / 0	0 / 0
J/ApJ/933/87/table2	UVMag	zphot	—	$2^{+2.7}_{-1.3}$ / $1^{+2.3}_{-0.9}$	0 / 0	0 / 0	0 / 0
J/ApJ/946/L35/final	UVMag	zphot	magnif. pending	$1^{+2.3}_{-0.9}$ / 0	0 / 0	0 / 0	0 / 0
J/ApJS/258/11/classic	lpFUVMAG	lpzBEST	—	$136^{+12.7}_{-11.7}$ / $92^{+10.6}_{-9.6}$	$72^{+9.5}_{-8.5}$ / $64^{+9.0}_{-8.0}$	0 / 0	0 / 0
J/ApJS/258/11/farmer	lpFUVMAG	lpzBEST	—	$297^{+18.3}_{-17.2}$ / $167^{+13.9}_{-12.9}$	$279^{+17.7}_{-16.7}$ / $220^{+15.9}_{-14.8}$	0 / 0	0 / 0
J/ApJS/269/33/tabled1	UVMag	zspec	—	$5^{+3.4}_{-2.2}$ / $2^{+2.7}_{-1.3}$	0 / 0	0 / 0	0 / 0
J/ApJS/278/33/tablea1	UVMAG	z	—	$3^{+2.9}_{-1.7}$ / $2^{+2.7}_{-1.3}$	$4^{+3.2}_{-1.9}$ / $1^{+2.3}_{-0.9}$	$3^{+2.9}_{-1.7}$ / $1^{+2.3}_{-0.9}$	0 / 0
J/MNRAS/418/2074/sample	M1500	zph	—	$4^{+3.2}_{-1.9}$ / 0	0 / 0	0 / 0	0 / 0
J/MNRAS/432/2696/tablea5	UVMAG	zphot	—	$2^{+2.7}_{-1.3}$ / 0	0 / 0	0 / 0	0 / 0
J/MNRAS/432/2696/tablea6	UVMAG	zphot	—	$3^{+2.9}_{-1.7}$ / $1^{+2.3}_{-0.9}$	0 / 0	0 / 0	0 / 0
J/MNRAS/432/2696/tablea9	UVMAG	zphot	—	$2^{+2.7}_{-1.3}$ / $1^{+2.3}_{-0.9}$	0 / 0	0 / 0	0 / 0
J/MNRAS/516/2162/table	UVMag	zmed	lensed	$1^{+2.3}_{-0.9}$ / 0	0 / 0	$1^{+2.3}_{-0.9}$ / 0	0 / 0

Note—Counts and intervals are recorded per catalog by the reviewed script; zero cells carry the recorded interval $[0, 1.87]$ and enter any fit as 84% one-sided upper limits per contract clause C1. "rest-NUV" marks the recorded band-mismatch flag (the tabulated magnitude is rest-NUV, not the frame's 1500 Å); "lensed" marks magnification-dependent samples that are never a blank-field density; "compilation" marks the no-area literature compilation; "magnif. pending" marks the JWST ERO table whose magnification status awaits a paper-level check.

$[11.5, 14)$ — dozens in total — beside photometric wide-survey tails of hundreds per slice that are uncorrected archival catalog content, contamination-dominated by construction, and countable only under their recorded provenance flags. In the highest slice, every one of the 143 objects at ≤ -20 comes from a single wide-field photometric catalog whose tabulated magnitude is not even in the frame's band; the labelled, in-band record there is empty.

What the literature reports, in the era's own receipted words: **?** identify "a sample of 88 candidate $z > 8.5$ galaxies" over the full CEERS survey — 55 (the summary section) or 58 (the sample-definition sentence; the discrepancy is the paper's own, recorded verbatim in our receipts) at $z \sim 9$, 27 at $z \sim 11$, and 3 at $z \sim 14$; **?** "select 13, 8, and 2 dropout galaxy candidates at $z \sim 9$, 12, and 16"; **?** measure the LF from a JADES sample reaching $M_{\rm UV} \sim -22$; **?** report a brightest bin at $M_{\rm UV} \sim -21$. These samples exist, their luminosity functions are published, and the papers pass every screening gate our plan could check from metadata — completeness corrections claimed and receipted, per-bin uncertainties tabulated, redshift coverage spanning our slices. The gap is not that the science is missing. It is that between the published table and a reproducing machine stands, in every single case, at least one unstated convention that no lawful step can supply.

6.2. *The closure record*

Table 3 is the paper's second headline result, the census's equal: four independent channels to the published bright-end LFs, opened under pre-registered rules, and closed — each for a reason the sources themselves state or fail to state, none for a reason any later deposit could not repair.

6.3. *The confrontation arithmetic*

The cost of the closures is stated exactly by the frozen forecast, because the forecast was written first. Fifteen model-prediction entries were frozen; twelve are direction-confrontable in the frame; the plan's pre-fetch census forecast four entries expected-confrontable on confirmed coverage, nine conditional on bin confirmation at fetch, and two non-confrontable by rule. After all four channels: **zero** entries are confrontable. The two-cut contrast that would discriminate stochasticity-class from efficiency-class models — frozen at -20 versus -21 precisely because two entries make claims there — cannot be evaluated for want of a bin-edge sentence. The frozen slice $[10, 11.5)$ is unpopulated roster-wide on the published-LF side — the only two tables that prospectively covered it are the two whose geometry failed — so even a future repair of a single table cannot, alone, restore the three-slice pace shape the conditional entries need. And the loss is asymmetric in an instructive way: the *values* of the ten verified **?** densities are known, receipted, and stance-confirmed; what is missing in every case is not data but *declaration* — a width, an attachment, an anchor. The pace of the bright end over $8 \leq z < 14$ is, today, a quantity whose public evidentiary basis exists almost entirely as prose and figures, not as machine-actionable numbers.

6.4. *What a contract-grade measurement would require*

The repair list is short, cheap, and specific; each item names the failure it would have prevented in this study.

1. *Deposit the ϕ table, not only the object catalog*, at the journal or CDS, machine-readably. This alone would have kept channel (i) open (**?**; **?** deposited catalogs without densities; **?**; **?** deposited nothing archive-native).

2. *State the bin geometry in the table itself*: edges, or a width plus one clause saying whether the tabulated $M_{\rm UV}$ is the bin centre or an edge. One sentence would have passed Finkelstein+24 Table 4 through E3, whose only missing required class it is; Harikane+23 Table 7 additionally states no redshift interval, and Harikane+24 Table 2 inherits its geometry row by row from donor studies, so both need more than one (§4.6); its absence is the single most load-bearing omission in this record.

3. *State the redshift interval on each LF section*, not only a dropout label — **?**'s Table 2 does this for three of four sections and is the model to follow.

4. *State the magnitude convention* — the wavelength anchor and the AB zeropoint — with the table, not only in prose a protocol cannot locate (**?**'s record carries no 1500 Å statement at all; E2).

5. *Where geometry is inherited from prior studies, identify the donor rows*, or restate the geometry locally; an unidentified per-row inheritance defeats any single-statement repair (**?**).

6. *Publish machine-readable completeness and area metadata* with object catalogs, so that a census like §5 could lawfully become a density measurement rather than a labelled raw count.

None of these requires new observations, new analysis, or referee-visible change. The distance between the current record and a machine-measurable bright-end pace is a few declarative sentences per paper — which is precisely what makes the gap worth measuring, and closing.

Table 3. The four-channel closure record for machine-readable access to the published $z > 8$ bright-end luminosity functions.

Documented reason for closure
e UCD universe holds no $z > 4$ galaxy rest-UV LF; flagship deposits are object catalogs without ϕ tables (**?**; **?**) or absent from VizieR (**?**; **?**)
e body absent from the sealed extraction (zero transcribable values, gap diagnosed to the character range); **?**: ten values verified, but bin edges/width, centre meaning, and the
e LF tables omit bin geometry: no edges/width/centre statement (**?**); geometry inherited row-by-row from donor studies (**?**); 0.5-mag width stated but centre-vs-edge attachme
sing geometry statements are absent from the papers' LF-presenting text; absence on a required class ⇒ excluded, not repaired — verdicts by independent adjudication, uphel

ne published HTML as served, collected under the pre-registered gate of §4.6.

7. HONEST LIMITS

A paper whose result is that other papers' records fail a reproducibility protocol owes its own record the same scrutiny. We state this study's limits in the same register we applied to its subject.

The census counts are not densities, and cannot become densities here. No completeness correction, no cosmic-variance interpretation, and no cross-catalog merging were performed — by contract, because the machine-readable record does not carry the completeness and area metadata that would make those steps lawful. The wide-survey photometric tails are raw catalog content, contamination-dominated by construction; the lensed-field counts are magnification-dependent and never a blank-field quantity; the spectroscopic compilation has no area and is never a density; four counted tables carry a rest-NUV band-mismatch flag—two of which contribute non-zero counts and are counted strictly under that recorded band-mismatch flag. Every such label travels with its counts in Tables 1 and 2, and any reader who quotes the totals without the labels is misquoting this paper.

The eligibility layer's verification has a measured depth, and we state it. All 96 extension verdicts were machine-checked — structure, required fields, verdict tallies, field fidelity between the verdict file and the manifest, alias-versus-manifest existence, parity-gap closure, duplicate identifiers — and the adversarial reviewer additionally opened 20 verdicts verbatim against their cached ReadMe files (the 14 highest-risk eligible rows and 6 disqualifications), reproducing every one. The remaining **82 of 96 verdicts rest on the machine checks and the recorded manifest evidence, not on a fresh human read of each ReadMe**. No refutation was found in the opened sample, but the coverage boundary is real and is recorded here rather than implied away. One candidate remains genuinely open: the ReadMe of `J/ApJ/963/9` returned HTTP 403 at review time, reproduced twice under different user agents, and its row stays pending-evidence — closed to counting, open to repair — rather than resolved by guess.

One receipt-custody defect occurred in the paper-text channel, and it is on the record. The channel's rules required the consulted article text to be archived, with its hash, at collection time; it was not — the archive directory is empty, so the consulted-text basis of the three geometry receipts is not independently re-executable from the lane archive alone (the quotes remain verifiable against the cited published versions, which are open-access at the pinned URLs). The defect was found by the second reviewer, escalated, and ruled RECORDED_DEFECT: outcome-inert in this instance, because all three verdicts the receipts support are failures, and setting the receipts aside entirely would dispose the tables identically (no valid receipt means unevidenced geometry means ineligible — the same exclusion). The ruling is explicit that the defect *would have invalidated any passing verdict*, and that back-filling the archive with a later re-serving of the pages — a different artifact at a different time — would have simulated compliance rather than achieved it, and was therefore refused. Collection-time archival is now a standing rule for any future pass. We report this defect at full length deliberately: it is this paper's own instance of the class of omission it measures in others.

Absence findings are scoped to declared records, never to the papers. The transcription channel's zero for **?** is a fact about a sealed corpus extraction whose incompleteness we proved ourselves (the table body missing from a diagnosed character range); the paper-text channel's absence findings are facts about pre-declared search surfaces — the LF-presenting sections of the published HTML — read once under a no-second-pass rule. A geometry statement could, in principle, live outside those surfaces: in an appendix we did not declare, or in a cited methodology paper (**?**'s width sentence delegates its methodology to an earlier paper we did not open, which may well state the attachment). The protocol's point is not that the information exists nowhere; it is that a reproducing reader following the paper's own LF presentation, forbidden from guessing, cannot reach it. The distinction is load-bearing and we do not blur it.

VizieR is not the only archive. The census enumerated VizieR exhaustively through two channels, and the meta-track additionally reached the journals' own machine-readable tables for the flagship papers plus a receipted corpus of 180 papers — but MAST high-level science products, survey-team databases, and journal MRTs beyond the fetched set were not enumerated. The scope is declared, not universal: this paper's claims are about the channels it opened, which are the standard channels of archival reproduction. A bright-end ϕ table with stated geometry may exist in some repository we did not reach; publishing its location would itself close part of the gap, and we would welcome exactly that correction. Finally, the enumeration's UV-evidence rule carries one known, recorded residual over-inclusion (an embedded "11500 Å" would false-fire), left deliberately to land in the eligibility layer rather than silently patched; no counted catalog was admitted through it.

8. CONCLUSIONS

We set out to measure the pace of the bright-end rest-UV luminosity function's decline over $8 \leq z < 14$ from public machine-readable data alone, under a contract frozen before any science row was fetched. What we measured instead — because it is what the record actually supports — is the record itself. The findings:

1. *The confirmed machine-readable bright end is dozens, not hundreds.* From 6,417 in-slice rows across 67 counted catalogs, the spectroscopic record holds 21 objects at $M_{\rm UV} \leq -20$ in $z \in [8, 9)$, 13 in $[9, 10)$, 11 in $[10, 11.5)$, and none beyond $z = 11.5$; the hundreds per slice are uncorrected photometric wide-survey tails, counted here only under explicit provenance labels. The $[8, 9)$ photometric record outside those tails still rests substantially on pre-JWST HST-era catalogs.

2. *Name-based archive search misses most of this record.* Of 112 candidate tables, 92 were discoverable only through UCD metadata; of the 23 catalogs contributing non-zero bright-end counts, 18 are invisible to every column-name pattern. The eligibility layer disqualified 34 impostors — stellar masses named like UV magnitudes, surface brightness masquerading as $M_{\rm UV}$, positional "Z" columns — each with its ReadMe reason on record.

3. *The published $z > 8$ luminosity functions are not machine-recoverable at pre-registered grade, through any of four channels.* VizieR's complete 83-table luminosity-function universe contains no $z > 4$ galaxy rest-UV LF; receipted corpus transcription yielded zero values for one primary source and ten verified but geometry-orphaned values for the other; the journals' machine-readable tables (20 of 23 parsed) all omit bin geometry; and a pre-registered pass through the papers' own text confirms the missing statements — a bin width with its attachment, edges, section redshift intervals, a wavelength anchor — are absent where a protocol-bound reader can reach. No pace confrontation survives on any source; the $[10, 11.5)$ slice is unpopulated roster-wide on the published-LF side; zero of twelve direction-confrontable model entries can be confronted.

4. *The gap is declarative, not observational.* In every closed case the missing element is a sentence, not a dataset: state the bin edges, say whether $M_{\rm UV}$ is a centre, attach a redshift interval, name the anchor, deposit the ϕ table. The repair list of §6.4 requires no new observations and no reanalysis.

The method conclusions are as transferable as the astronomical ones. A frozen contract plus adversarial pre-execution review caught, before it could contaminate a count, a one-character defect that no inspection of outputs would have surfaced; a fail-closed receipt once forced re-execution rather than silent repair; and the same discipline, turned on this paper's own conduct, surfaced and recorded a receipt-custody defect in our paper-text channel rather than absorbing it. A census whose receipts have never failed closed has not demonstrated that they can.

We frame the result as reproducibility data, not as criticism of the source papers — human readers infer bin conventions effortlessly, and nothing here questions the published measurements themselves. But the bright-end pace beyond $z = 8$ is now a quantity on which model classes stake opposing, falsifiable claims, and the public, machine-actionable evidentiary basis for it is effectively empty. The lawful ways forward are equally concrete: a reviewed fetch of the primary sources' journal tables supplying edges and anchors would revive the strongest surviving source; a one-sentence geometry statement in future LF tables would prevent the failure class entirely; and machine-readable completeness and area metadata would let a census like this one become a measurement. Until then, the census stands as the record's honest portrait: dozens confirmed, hundreds unvetted, and a luminosity function that exists in print but not yet in public machine-readable form.