S5 Supplementary analyses

Stellar selection and temporal discrimination under explicit assumptions

Stonehenge | Heritage v1.28 | 27 September 2026

This supplement reports executed retrospective tests of stellar selection, timing scale, compiler feasibility and construction-phase discrimination. With W=(28,72,16) and its timing scale fixed, Cygnus ranks first in the declared 5,280-class panel. New bounded-epoch diagnostics quantify the error penalty in the later sarsen phase and preserve that ranking within the narrower ditch interval. Free words and scale admit alternatives. Sections 3.3 and 8 document the new calculations, their post-result ordering and verification of the publicly archived word and epoch. The earlier S1–S4 analyses retain their original inputs, comparison universes and execution records.

1. Protocol, inputs and selection order

PROTOCOL.md was recorded before the new executions, after the primary findings and reviewer criticisms were known. No prospective discovery claim follows from that ordering. The catalogue is the archived S4 table of 5,041 Hipparcos entries with V≤6. For each of the 88 modern constellation regions, select the five lowest V magnitudes, with HIP as the tie-breaker. There is no prior filter on shape, spatial tuple, archaeological date or visibility. The modern regional partition is an explicit analytical convenience, not an inferred prehistoric star classification. The original protocol is preserved. PROTOCOL_v126_ADDENDUM.md records that the original-compiler controls were formulated after these unrestricted results were known. PROTOCOL_v128_ADDENDUM.md records that the bounded-epoch comparisons were formulated after the fixed-word results and chronology objections were known.

The temporal graph uses five distinct role-ordered stars and three unsigned right-ascension widths: roles 1–2, 2–3 and 4–5. Widths are multiplied by four to express sidereal minutes. Positions at catalogue epoch J1991.25 are propagated using the archived proper motions and long-term precession engine. Primary and six-dimensional astrometry are kept as separate nominated-input sensitivity branches. They are not mixed with the common-catalogue panel when ranking competitors.

Each regional set supplies 120 labelled role assignments. Exchanging roles 4 and 5 leaves their unsigned interval unchanged, so 60 computations per region represent 120 assignments. The 512 ordered words from V={4,16,28,32,40,68,72,100} yield 5,406,720 labelled role-word descriptions, or 2,703,360 classes under that one symmetry. These are dependent descriptions, not independent statistical trials. First/third-role reversal with exchange of the first two word channels remains separately labelled.

2. Objective, search and numerical verification

For each word W and role graph, the absolute objective is $RMS(t,W)=\{\Sigma_i[T_i(t)-W_i]^2/3\}^{1/2}$. Search astronomical years −8000 through +2000, without chronological fitting constraints. Each sampled local minimum is refined with bounded scalar minimization; endpoints and appropriate boundary cells are also evaluated. The primary new run uses a 1-year grid and the verification run a 0.5-year grid. The 3000–2920 BCE initial-enclosure window is inspected only after full-window fitting. Section 3.3 separately imposes explicit chronological bounds and does not replace this original unconstrained search.

Because the library is a Cartesian product, its nearest word at any epoch is found by independently selecting the nearest allowed value in each channel. For this qualifying threshold, no passing word is lost by searching that lower envelope: an RMS≤0.390126 minute bounds each absolute residual by about 0.676 minute, below half the minimum library-value spacing of 4 minutes. A passing word must therefore be componentwise nearest at its passing epoch. For every qualifying graph-word pair the best discovered qualifying minimum is retained. This shortcut is verified against direct evaluation of all 512 words for 1,000 deterministic random width vectors; scores agree exactly.

The two temporal grids return the same 336 qualifying graph-word classes and the same 39 qualifying ratio classes. Across all 5,280 role graphs, the maximum difference in best absolute score is 4.09×10^{-11} minute and in relative-dispersion score 9.99×10^{-10}. A unit-conversion check agrees within 1.12×10^{-15}. Agreement across two grids is numerical evidence, not an analytic proof against every unbracketed minimum. CSVs contain all graph summaries and all qualifying classes; they do not claim to archive every point of every score curve. Scripts and both execution logs are included.

3. Crossed stars, roles and words: results

The uniform-catalogue nominal benchmark has RMS 0.390126064 minute at 2965.425 BCE, with temporal vector (27.789129,72.605469,16.213391) minutes. The primary archived specification remains approximately 0.418975 minute at 2965.2185 BCE. These input conventions produce slightly different scores and are kept separate.

The new panel contains 336 qualifying classes, equivalent to 672 labelled descriptions, in 57 regions. Six classes have their best qualifying word-specific minima within 3000–2920 BCE. They comprise four Cygnus descriptions and two Lacerta descriptions related by first/third reversal. The six entries are not six independent astronomical systems. All appear in the table below; HIP identifiers give the complete role order. Final-pair order has been canonicalized by increasing HIP.

Region / HIP roles 1–5	Word	BCE	RMS (min)
Cyg 102098 / 100453 / 95947 / 97165 / 102488	28, 72, 16	2965.425	0.390126
Cyg 100453 / 102098 / 95947 / 97165 / 102488	28, 100, 16	2967.187	0.285972
Cyg 95947 / 102098 / 100453 / 97165 / 102488	100, 28, 16	2967.187	0.285972
Cyg 95947 / 100453 / 102098 / 97165 / 102488	72, 28, 16	2965.425	0.390126
Lac 111169 / 109754 / 110538 / 109937 / 111022	32, 32, 28	2921.499	0.137723
Lac 110538 / 109754 / 111169 / 109937 / 111022	32, 32, 28	2921.499	0.137723

The closest fit anywhere in this panel is a Scutum graph with HIP roles (92175,91117,91726,90135,90595), W=(16,4,4), RMS 0.018962 minute, near 3237.2 BCE. Reversing the first three roles and the first two word channels gives an equivalent score. This example lies outside the initial-enclosure window. The Lacerta example with W=(32,32,28), RMS 0.137723 minute and epoch near 2921.5 BCE shows directly that a better

timing fit plus inclusion in that window is not exclusive to Cygnus within this panel. Its second word channel, 32, cannot equal 3q for integer q. The corresponding Scutum channel, 4, fails for the same reason. These words therefore do not arise from any integer inventory under the original compiler.

This does not show that Lacerta or Scutum reproduces the full material model. The panel releases words without requiring each word to arise from a feasible alternative count compiler, and does not require the spatial tuple or a horizon test. Its question is the specificity of timing and chronology under joint star-word freedom. It cannot replace an equal-freedom comparison of complete archaeological pipelines. The earlier S4 censuses cover more star subsets within regions but retain spatial filters; their conditional survivors and these new alternatives are compatible results for different questions. Section 3.1 explicitly quantifies that compiler distinction using constraints already present in S1, rather than treating every free-word pass as a complete competing model.

3.1 Original word and positive-integer compiler controls

For an arbitrary positive integer word (d,m,r), the unchanged map gives q=m/3, h=(d−r)/2, B=q+h, D=t, C=t+h and A=d+q+t. A positive integer solution exists exactly when q and h are integers and B>0; every integer t>max(0,−h) then supplies one. This proof defines the test implemented in check_original_constraints.py. It is feasibility under the original map, not satisfaction of its three extra closure equations, and not archaeological feasibility.

Only 55 of the 512 words satisfy these conditions, giving 55×5,280=290,400 classes in the same regional role panel. Of the 336 unrestricted timing passes, 26 are compiler-feasible. The existing complete qualifying lists suffice for this classification because feasibility removes words without changing their timing scores. The arbitrary example inventories in compiler_classification.csv demonstrate existence; they are not proposed inventories of Stonehenge.

The fixed-word control directly minimizes the absolute-RMS objective for W=(28,72,16) on all 5,280 role classes. It reuses the archived astrometry and precession engine but implements that objective independently of audit.py's Cartesian lower envelope. The 1-year grid, refined local minima and endpoint cells rank nominal Cygnus first, with RMS 0.390126064 minute at 2965.425143 BCE. Reticulum is second, with HIP roles (19780,17440,19921,18597,18744), RMS 0.554285025 minute at 2848.633911 BCE. Reticulum has 1.42 times the nominal RMS. The “pass” threshold used in the retained files is Cygnus’s own score; one pass denotes sole first place, not an independently calibrated absolute cutoff. No chronological, spatial-tuple or visibility filter is imposed. This direct control was run on a 1-year grid; the two-grid agreement reported in section 2 concerns the original crossed and ratio searches.

3.2 Sensitivity to the archaeological comparison interval

Among the 26 compiler-feasible passes, only nominal Cygnus has its best qualifying minimum in Darvill et al.'s 3000–2920 BCE initial-enclosure window. Greaney et al.'s revised 2980–2720 cal BCE 95% ditch interval retains three classes: Cygnus and two Vela assignments. The 2940–2800 cal BCE 68% interval retains the two Vela assignments and excludes the unconstrained Cygnus optimum. These are comparisons after full-window fitting, not reruns with a date prior (main references 2 and 18).

Vela's roles are (44816,39953,45941,42913,52727), W=(32,72,68), RMS 0.364695145 minute at 2905.808369 BCE, and (44816,45941,39953,42913,52727), W=(40,72,68), RMS 0.276188419 minute at 2884.338143 BCE. Both fail the full three-closure model because r=68 cannot equal D^2 for integer D. Requiring all three original

closures uniquely restores (56,30,10,4) and hence the original word. This mathematical restriction is distinct from evidence that prehistoric builders used those closures.

3.3 Fixed-model phase contrasts and interval-restricted ranking

These additional diagnostics preserve the word, five roles, minute-per-integer scale, equal channel weights and propagation engine. They minimize absolute RMS separately within 3000–2920, 2980–2720, 2940–2800 and 2620–2480 BCE, using one-year samples, local refinement and both boundaries. The original full search, astronomical years −8000 to +2000, is also repeated for the nominated configuration under primary, six-dimensional and common-catalogue astrometry. None of these new interval-constrained fits is described as chronology-free.

The global optimum belongs to both the initial-enclosure and revised 95% ditch intervals, leaving the minimum unchanged there. The three input variants give the following phase contrast. All Stage 2 minima occur at its oldest boundary, 2620 BCE. Ratios compare deterministic errors within one model; they are not likelihood ratios or independent archaeological confirmations.

Astrometry	Global BCE	Global RMS min	Stage 2 RMS min	RMS ratio
Primary	2965.2185	0.418975	2.081039	4.967
Six-dimensional	2965.8550	0.420359	2.084465	4.959
Common catalogue	2965.4251	0.390126	2.079616	5.331

Restricting the epoch to Greaney et al.'s 68% interval, 2940–2800 BCE, gives a minimum at 2940 BCE for each nominated input set. Primary RMS is 0.444750 minute (+6.15%); six-dimensional RMS is 0.447313 (+6.41%); common-catalogue RMS is 0.418210 (+7.20%). These are best fits at an interval boundary, not replacements for the nominal astronomical epoch.

The fixed-word common-catalogue ranking follows from a simple bound. For every competitor j and restricted interval I, min over I of RMS_j is at least its unrestricted minimum. All 5,279 non-nominal classes have archived unrestricted RMS ≥0.554285025 minute, whereas Cygnus attains 0.418209685 inside I. It therefore remains first. Reticulum's unrestricted optimum is at 2848.633911 BCE, inside I, so its restricted score equals 0.554285025 and it remains second, with 1.325376 times the Cygnus RMS. This numerical ranking proof uses the archived all-class global scores and the new bounded Cygnus curve; it does not claim a new full 5,280-class census.

The count of unrestricted optima inside a window answers a different question. Of the 5,280 fixed-word classes, 22 have global optima in 3000–2920 BCE, 65 in 2980–2720, 39 in 2940–2800 and 20 in 2620–2480. These are dependent panel frequencies, not historical false-positive probabilities. Neither these fractions nor a ratio of chronological window lengths calibrates the entire retrospective selection procedure.

Figure 11 in the main manuscript shows the phase contrast and constrained ranking. temporal/audit_temporal.py reproduces results.json and temporal_profiles.csv from the retained inputs; temporal/verify_temporal.py checks the endpoint scores, the archived ranking bound, phase ratios and exact input identity. The new comparison does not remove the free-word Vela cases, scale-free alternatives or lack of independent evidence for prehistoric use of the compiler.

4. Free temporal scale

Define $z_i(t)=T_i(t)/W_i$ for W=(28,72,16), mean $z=\bar{z}$, and $E(t)=\{\text{mean}[(z_i/\bar{z}-1)^2]\}^{½}$. Multiplying all T by a constant unit conversion leaves E unchanged. Report $\bar{z}$ as the common timing scale in minutes per numerical unit. This is the coefficient of variation of the three channel ratios: it changes weighting relative to absolute RMS and is not a transfer of the original objective or its ranks. The numerical value of the reported scale changes with units, while E and the selected epoch do not.

Nominated input	Minimum (BCE)	E (%)	Scale (min/unit)
Primary archival astrometry	2986.805	0.681862	1.001865
Six-dimensional variant	2987.516	0.681667	1.001907
Uniform catalogue	2985.482	0.649026	1.001529

The primary ratio curve also has a worse local minimum near 6547.1 BCE, with E about 1.588%; the complete primary and 6D minimum lists, including endpoints, are retained in the summaries. The free-scale primary result is therefore near the earlier branch without being the same objective or the same numerical epoch.

In the common-catalogue panel, 39 classes in 25 regions match or improve Cygnus, equivalent to 78 labelled role assignments under final-pair exchange. The best is a Draco graph with HIP roles (83895,94376,87833,80331,85670), E=0.0205813%, scale 2.032299 minutes per unit, near 3052.0 BCE. Its epoch falls outside 3000–2920 BCE. No archaeological filter is imposed on the ratio ranking. Alternative regions also score well; all qualifying graphs are supplied.

The comparison tests a temporal scale convention rather than just acknowledging it. A near-2987 BCE Cygnus branch survives; exclusivity does not. This is not historical justification for reading architectural numbers as angles. It also does not remove the fixed-J2000 absolute spatial tuple or the auxiliary solar degree convention. Ratios cannot identify the absolute size of W: replacing W by any positive multiple leaves E unchanged.

5. Inventory identification and construction phases

The temporal map is $W=M(A,B,C,D)^T$, with rows of M equal to (1,−1,1,−2), (0,3,−3,3) and (1,−1,−1,0). M has rank three and kernel span{(1,0,1,1)}. Consequently W=(28,72,16) admits exactly the affine family (56+t,30,10+t,4+t); positive integer inventories require integer t≥−3. This follows by subtracting any solution from the nominated one and solving the homogeneous system. The script checks the rank and kernel and samples t from −3 to 1,000 as an arithmetic check, not as the proof of the infinite family.

For fixed W, r=16 and the additional closure $r=D^2$ select positive D=4, hence t=0. The original theorem under all three specified closures is preserved. The new result isolates what timing alone cannot infer. It does not argue that the full conditional compiler is underdetermined. Nor does mathematical uniqueness establish that its assumptions were available or used during the earliest archaeological phase.

With k=(10,24,26) and exterior x=(33.435,80.095,86.44) m, $u=(k\cdot x)/(k\cdot k)=3.331412721893491$ m. Applied without refitting to inner k=(4,10,12,16) and x=(13.7,33,39,54) m, the signed percentage residuals are −2.732475, +0.951901, +2.505007 and −1.291475. RMS is 2.019623% overall, 2.046037% for trilithon/sarsen and 1.992858% for Z/Y. These subsets share one scale and selected model roles; they are not independent statistical trials. This calculation uses the original outer-only scale, distinct from the all-channel descriptive fit also printed in S1 Note 2.

The earliest Aubrey circuit is compared with the initial enclosure chronology. Darvill et al. [main reference 2] (Table 3) place the major sarsen settings and Stations in Stage 2, 2620–2480 cal BCE, and Y/Z excavation in 1630–1520 cal BCE. Station dating remains insecure; assignment in that model is not an independent determination for every setting. English Heritage's broader context places central stones around 2500 BCE and late holes at 1800–1500 BCE (main reference 35). Thus the complete inputs cannot securely be assigned to the nominal early epoch. Shared approximate diameters alone do not demonstrate inheritance.

6. Context figure and measurement provenance

Fig 1A reproduces an unaltered historical aerial photograph by Philip Henry Sharpe (Royal Engineers, 1906), public domain through expired Crown copyright according to its Wikimedia Commons record. It illustrates the condition in 1906 and supplies no measurements. The full source URL and acquisition date appear in README.md. Panels B and C plot existing archaeological drawing proxies; D identifies the primary measurement boundaries.

Panel B uses 56 manually read Aubrey symbols and four Station proxies from Cleal Plan 1. The 20 m bar spans 357 pixels, or 0.0560224 m/pixel: sampling, not archaeological uncertainty. Stations 91/93 use projected footprint centres, 92/94 socket-outline proxies, with partial excavation at 94. Panel C separately uses Rees click-plan label or hatch-centroid proxies, transformed using the archived registration constants. It labels entrance 1–30 and complete inner opening pairs 51–52, 53–54 and 57–58; separate fragment labels are retained. Dotted normals describe these proxies, not securely reconstructed original mounting axes. The panels are not mutually registered or independent surveys. Neither supplies full outlines, every missing centre, or measured per-feature uncertainty. The primary dimensional fit continues to use published measurements, not distances newly extracted from this figure.

7. Observations that can contradict the reconstruction

Claim to test	Positive adverse observation	Present status
Specified geometric transfer across phases	Predicted withheld later coordinates disagree beyond a justified survey/construction tolerance under fixed registration; alternatives scored on identical holdouts perform better.	Not yet testable with a homogeneous independent coordinate and uncertainty set.
Specific opening or cord realization	Securely reconstructed original socket or mounting axis is incompatible with the nominated position/direction prediction.	Original-setting evidence needed; existing drawing proxies are not decisive.
Early availability of all material roles	Stratigraphy establishes that required physical roles arose later.	Already constrains contemporaneity; inheritance is a distinct hypothesis requiring evidence.

Claim to test	Positive adverse observation	Present status
Exclusive timing/early-window selection with arbitrary library words	Another admitted star-word description equals or improves the benchmark and has its minimum in the declared window.	Fails in the unrestricted-word panel (Lacerta). Those words fail original compiler integrality. Fixed W and scale retain only Cygnus in the same role panel.
Independent cultural identification of the stellar convention	A documented alternative reference convention or material-role definition defeats the claimed correspondence under rules fixed beforehand.	Current absolute-angle convention lacks independent cultural calibration.
Exclusive fixed-word selection in the stated bright-star role panel	Another role assignment matches the common-catalogue benchmark with W=(28,72,16), fixed scale and identical astrometry.	Not observed in the 5,280-class direct search. Other star subsets and alternative compilers remain outside this control.

Absence of a portable astronomical instrument, incomplete excavation or a missing socket is not automatically a falsifier. The relevant distinction is between an observed contradiction of a specified prediction and insufficient evidence to score it. A successful geometric transfer would support continuity of geometric information; it would not independently identify Cygnus or establish timekeeping intention.

8. Reproducibility and relationship to the original archive

The README specifies commands and file meanings. The v1.25 archive contains both temporal grids and verification. The content audit reran the full S5 1-year panel and verify.py, compared against the archived 0.5-year outputs, reran the original three-closure audit, and executed the new direct fixed-word search. The complete original S2–S4 runners were not rerun. All 5,280 fixed-word results and 336 feasibility classifications are supplied, together with three-window summaries. S1 begins with a revision guide and reference crosswalk and carries notices at the relevant original notes. S2–S4 retain their original analysis records. No construction-date confidence interval, global false-positive probability or probability of intention is produced. The v1.28 addition executes the bounded-epoch diagnostics and their verifier, and adds Figure 11. make_figure_v128.py reconstructs that figure from the numerical profile. The current manifest separately records this execution and the untouched older comparison outputs.

8.1 Public antecedent and limits of the provenance claim

The public record https://zenodo.org/records/22103786 was created and updated on 26 August 2026. Its Giza-titled ZIP contains the subtree S3/repro/CALANAIS/PRECESSIONAL_CLOCK/. The script calanais_chronology_audit.py fixes CLOCK_TARGET_MIN=(28,72,16), the Deneb–Sadr, Sadr–Albireo and Aljanah–Fawaris channels, and the four-minutes-per-degree conversion. observed_results.json records the clock optimum as astronomical year −2964.2184951979593, labelled 2965.2 BCE, with RMS

0.4189747781503438 minute. Its executed vector agrees with the primary Stonehenge result at the reported precision. The downloaded archive matches the MD5 checksum returned by the public record; provenance/provenance_ledger.json supplies the exact paths and SHA-256 hashes of the extracts.

The earlier routine obtains the nominated clock optimum within astronomical years −3400 to −2500; the present Stonehenge analysis tests the full −8000 to +2000 window. The earlier archive therefore documents the prior word, operator and nominated epoch, not prior execution of every Stonehenge control. The public Calanais account of 10 September 2026, https://zenodo.org/records/22691368, also states the word and approximately 2965 BCE result. These records establish documentary anteriority of the ingredients relative to their stated Stonehenge application; they do not establish ignorance of the already published archaeological chronology, a blind discovery protocol, or statistical independence between applications. Any material system yielding the same word under this operator receives the same clock epoch.