Function of Tzolk'in Calendar structures in the Postclassic Maya Codices

Final Report

Jakub Špoták

June 2026

Project 09I03-03-V04 – Fellowships for excellent researchers R2–R4
Comenius University Bratislava, Faculty of Arts – Centre for Mesoamerican Studies

Calendar Intervals in the Maya Codices: A General Model for Generating Almanac Intervals

I. INTRODUCTION

This report closes the MayaCal project. It summarises what was built, what was found, and what remains open, and it serves as the documentation of the general model for the use of Tzolk'in intervals in the Maya codices — the final objective of the project and the deliverable of Work Package 7.

The project set out from two places in the Paris Codex. Page 21, the so-called Table with the day-signs, has been known since the nineteenth century and read for most of that time as a symbolic display of the 260-day calendar. Pages 23–24, the constellation pages, are among the most reproduced passages in all of Maya studies; the principal investigator had identified, in the eroded lower register, traces of a second calendrical table that had gone unremarked in more than 150 years of research.

Both objectives turned out to depend on the same thing: a complete, queryable record of how the codices actually use the Tzolk'in. That is why the first eight months of the project were spent building a database rather than analysing pages. The order was correct. The results reported below are quantitative statements about a corpus, not readings of an isolated page, and they could not have been made in the other order.

II. METHODOLOGY

Three instruments were used, in sequence.

The relational database. A FileMaker Pro system organised as COD_Codices → pages → registers → structures, with the 260-day calendar implemented in calculation fields and a recalculation script that propagates a series of distance numbers across every level of an almanac. Its design is documented in Stage Reports 1 to 3. The corpus it holds comprises 15,332 transcribed blocks from the Dresden, Madrid and Paris codices — 5,787, 7,890 and 1,655 respectively — resolving into 334 distinct structures, of which 2,041 almanac records were extracted for analysis.

Epigraphic and iconographic analysis. Standard reading of the day-signs, coefficients and distance numbers of the two Paris Codex sections, with attention to the physical division of page 21 by the painted rope and to the eclipse and sky-band iconography of pages 23–24.

Arithmetical and statistical testing. Every structural claim made below was tested against the corpus or verified exhaustively rather than argued from example. Where a claim is a general arithmetical statement, it was checked over the complete space of cases; where it is an empirical claim about scribal practice, it is reported with its counts and its exceptions.

III. RESULTS

III.a The table on page 21

The table has eleven rows. Rows 1–5 and 7–11 carry Tzolk'in day names; row 6 carries a single number per column. Those numbers are Tzolk'in coefficients, which caps the table at thirteen columns, one per coefficient. Eight columns are preserved or reconstructible.

Every one of the ten day names in a column shares that column's coefficient. This is arithmetically necessary rather than a scribal convention: the ten days of a column stand 26 days apart, and since 26 = 2 × 13 the coefficient cannot change across them. The same fact appears in the layout as a uniform rule — each day in the lower half stands six positions further on in the 20-name cycle than the day above it.

Reading a column as two separate five-day circuits is therefore only one of two available readings. Adding 26 days repeatedly from row 1 traverses all ten rows in a fixed zigzag — rows 1, 7, 5, 11, 4, 10, 3, 9, 2, 8 — and returns to the start after 10 × 26 = 260 days. The upper circuit alone, stepping 52 days, closes in five steps. A single column thus encodes both a five-level and a ten-level almanac skeleton.

Across columns, the starting day advances by a triangular sequence over the first six columns: +2, +3, +4, +5, +6, so that the shift from column n to column n+1 is n+1. Re-examination of the two rightmost preserved columns shows that the rule generalises more simply: the day-name shift into a column equals that column's coefficient. For columns 2–6 the two formulations coincide. For column 7, the two surviving glyphs — Men in row 4 and Chuwen in row 5 — identify it as a second coefficient-6 column rather than the coefficient-7 column the triangular sequence would predict, and the general rule then gives the observed shift of +6; applied once more it places column 8 at coefficient 7 with starting day Ix, a reconstruction that satisfies the 26-day rung throughout but awaits a surviving glyph. The coefficient sequence of the preserved table thus reads 1, 2, 3, 4, 5, 6, 6, 7. An arithmetical extension to all thirteen coefficients would cover 130 distinct Tzolk'in dates — exactly half the cycle, with no repetition — but the preserved and reconstructible table comprises eight columns, and the corpus evidence below supports the eight-column reading as the intended extent.

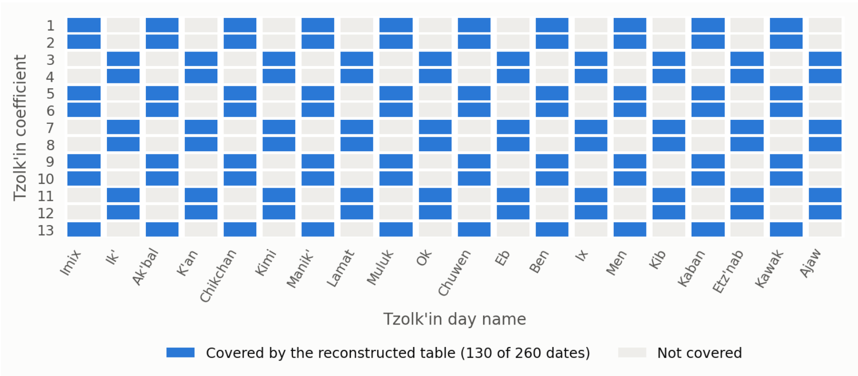

Figure 1 The 130 Tzolk'in dates covered by the reconstructed thirteen-column table, plotted against the full 260-day cycle. The paired-row banding reflects the alternation of day-name groups between successive pairs of columns.

Two columns of the preserved table do not obey the uniform six-position rule. In columns 3 and 4 the shifts read +6, +18, +6, +14, +6 and +6, +6, +18, +6, +14 respectively. A single exchange of Ajaw and Lamat in the lower half of each column restores +6 throughout. That the same pair is displaced, at the same relative positions, in both columns points to a scribal error rather than to a deliberate variant — and the two days concerned stand twelve positions apart in the name cycle, further than any other adjacent pair in the layout, which is the kind of distance at which a copying slip is easy to make.

One observation from the preserved table initially resisted the stated rules and was flagged for re-checking against the original. The internal step within a column is +8 in columns 1 and 6 but +16 in columns 2, 3, 4, 5 and 7, although the day-name grouping that the painted rope marks physically would predict a single step per group. The re-check resolved the matter in an unexpected direction: the step is not a function of the day-name group at all. The two steps are two orderings of the same material — +8 traverses a column's five-day half in real steps of 52 days, +16 traverses an equivalent half in real steps of 104 days — and the corpus analysis in III.d shows that these are precisely the two ladder formats attested among five-level almanacs. The table carries both reading orders because scribal practice used both: columns 1 and 6 lay out the 52-day ladder, the remaining columns the 104-day ladder, and the duplicate coefficient-6 column supplies the same ten days of coefficient 6 in the second order.

III.b The tables on pages 23–24

The upper calendrical table of the constellation pages uses a small recurring set of day names at 28-day intervals, which permits the destroyed portions to be restored with confidence. The second table, in the eroded lower register, was reconstructed from physical traces: a red coefficient fragment and the curve of a round day-sign beneath the bat figure, and a further coefficient with sign fragments at the middle of the preceding page.

The reconstruction was tested arithmetically. Advancing the 28-day interval from 8 Ix gives 10 Ik', 12 Ok, 1 Etz'nab and then 3 Kimi — a date whose coefficient and day-sign both match the surviving fragment. The identification is therefore not a guess constrained only by plausibility; it is the value the sequence requires.

The relationship between the two tables is the more consequential finding. The double page does not begin in the upper table but in the lower one, and the two interlock at 14-day intervals: 11 Ix, 12 Lamat, 13 Ik', 1 Kib, 2 Ok, alternating between registers. Two such steps make the 28-day interval observed within each table, and thirteen 28-day intervals — one per constellation figure — span 364 days, the computing year. The 14-day step is what makes the interval between successive constellation positions calculable in a single operation instead of two.

Interpretatively, a 28-day interval corresponds approximately to the displacement of a constellation over the horizon, and the animals hang from eclipse glyphs, that is, from the sign of the sun's disappearance. The 14-day interlock gives the observer a finer division of the same phenomenon. This supports the archaeoastronomical reading of the double page reported under Work Package 6.

III.c The general model

Every constant appearing in the two Paris Codex tables — 14, 26, 28, 39, 40, 52, 156, 208 — is a special case of a single expression. Because the 260-day cycle is the product of a 13-cycle and a 20-cycle, the Chinese Remainder Theorem gives a closed form for the interval between any two Tzolk'in dates:

$$D = (40 \cdot \Delta\text{coefficient} - 39 \cdot \Delta\text{day-name}) \bmod 260$$

Δcoefficient is the difference of the two coefficients modulo 13 and Δday-name the difference of the two day-name indices modulo 20. The expression was verified exhaustively over the complete space of 67,600 combinations — every one of the 260 starting dates against every one of the 260 possible intervals — with no discrepancy.

The two multipliers have a direct interpretation, and this is what makes the model usable rather than merely correct. 40 is the coefficient key: it advances the coefficient by one and leaves the day name untouched. 39 is the name key: it advances the day name by one and leaves the coefficient untouched. Every structural constant of the page 21 table is this expression evaluated at particular inputs — the upper-to-lower rung (Δcoefficient 0, Δname 6) gives 26; the step within a column of the first day-name group (Δname 8) gives 208; the step within the second group (Δname 16) gives 156; the cascade between adjacent columns (Δname 1) gives 221, that is −39.

A scribe who held these two numbers could compute the distance between any two Tzolk'in dates directly, with no table at all. On this reading the table on page 21 is not a calculating device but a pre-computed extract of that arithmetic, written out for the date range the page required.

The expression is an analytical description, not a claim about Maya method. Nothing in the corpus suggests that scribes multiplied or reduced modulo 260; their attested procedure was additive — small distance numbers, none exceeding 39 in the frame-level dataset, counted forward step by step. The table materialises exactly that arithmetic in navigable form: the rung, the circuit and the column-to-column move are the additive steps from which any interval the page required could be assembled. The formula compresses the structure for the analyst; the table lays it out for the practitioner. The formula is ours; the table was theirs.

III.d Verification against the corpus

The model predicts which almanac formats a table of this construction would serve. That prediction is testable, and the database was built to test it. Of the 2,041 almanac records, 1,201 have a regular interval between levels; the distribution is as follows.

Interval between levels	Almanacs	Share	Present in the p. 21 table
52 days	666	55.5 %	yes — upper circuit, five steps
65 days	275	22.9 %	no
26 days	106	8.8 %	yes — full zigzag, ten steps
irregular	82	6.8 %	—
104 days	23	1.9 %	yes — reading order of the +16 columns
208 days	14	1.2 %	yes — step within group 1
117 days	9	0.7 %	no
156 days	9	0.7 %	yes — step within group 2

Table 1 Interval between successive levels of an almanac, computed for all 2,041 almanac records by means of the general formula.

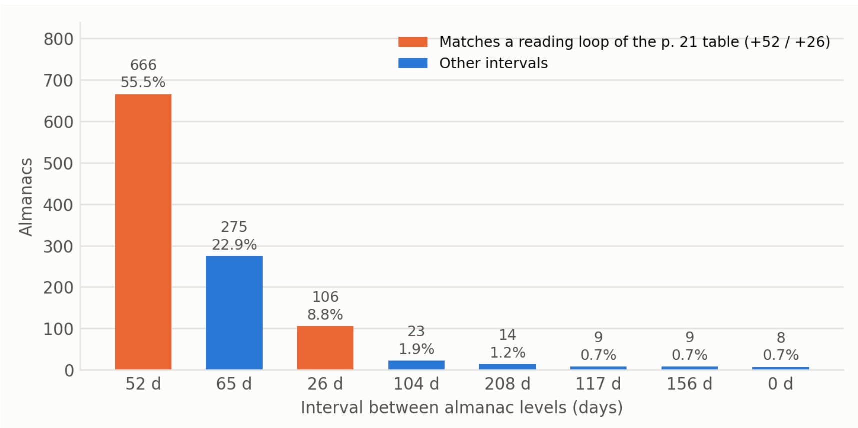

Figure 2 Distribution of level intervals across the almanac corpus. Nearly two thirds of all regular almanacs use one of the two intervals the page 21 table lays out.

772 of 1,201 almanacs — 64.3 % — use an interval of 52 or 26 days, the two circuits the table physically contains. Counting the 104-day ladder, which the +16 columns lay out in their reading order, raises the covered share to 795 of 1,201, 66.2 %. The correspondence also holds in the other direction: 635 of 680 five-level almanacs use 52 days, and 104 of 111 ten-level almanacs use 26 days, which is what 260 divided by the number of levels requires. Two further intervals attested in the corpus, 208 and 156 days, are the internal steps of the table's two day-name groups.

The honest limitation is visible in the same table. The second most common interval is 65 days, used by 275 almanacs — 22.9 % of the corpus — and the page 21 table does not cover it. The table is therefore not a universal aid to almanac construction. It is an instrument tuned to two specific formats, which happen to be the two most frequent.

Individual distance numbers were examined separately, as a check that the level-interval result is not an artefact of how almanacs were extracted. Restricting to blocks classified unambiguously as distance numbers, and excluding period notations, gives 982 values.

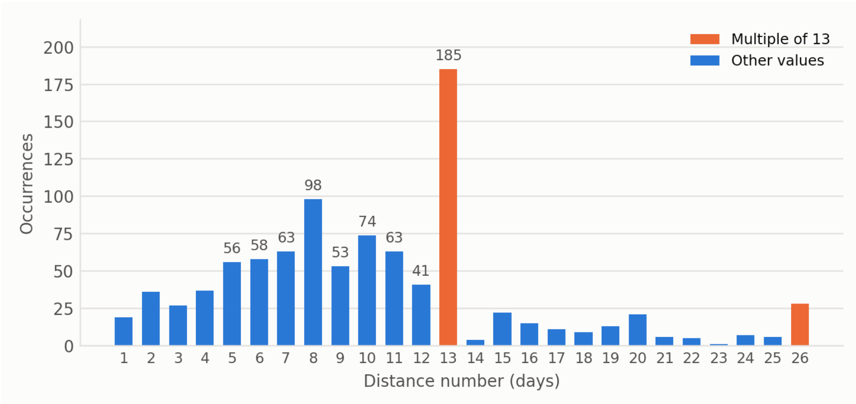

Figure 3 Distribution of the 982 unambiguously classified distance numbers. Multiples of 13 are marked.

The distribution peaks sharply at 13 — 185 occurrences, 18.8 % of the total, with multiples of 13 accounting for 21.9 % — which is consistent with the level intervals: four steps of 13 make 52. 17.5 % of values exceed 13, reaching as high as 39, which confirms that these are intervals and not coefficients. The distance numbers describe the steps taken within an almanac; the level intervals describe the frame those steps fill. The page 21 table addresses the second, not the first, and the two datasets agree.

A finer re-extraction of the corpus, carried out to test the ladder reading directly, records every almanac frame by frame: 434 almanacs across the three codices, 2,549 levels, 9,286 dated frames and 7,596 frame-to-frame distance numbers. Three results follow. First, the structures the table natively encodes — five levels of 52 days, ten levels of 26 days, five levels of 104 days — account for 309 of the 434 almanacs, 71 % of the corpus, and the twenty-one 5 × 104 almanacs match the reading order of the +16 columns in nineteen cases, the two exceptions having erased levels. Second, the specific dates on which almanac levels begin are over-represented among the table's cells: 31.8 % of 2,367 computable level starts fall on one of the table's seventy dates, against a chance expectation of 26.9 % ($z = 5.3$), whereas the dates of frames within levels show no such enrichment (27.6 %). The table's affinity is to the skeleton of an almanac, not to its interior. Third, every one of the 7,596 frame-to-frame distance numbers is 39 or smaller, with the distribution peaking at 13; the table's large inter-column intervals never occur as painted distance numbers, which confirms that the table addressed level structure rather than the small steps within levels. A weaker signal points the same way: among almanacs whose constant coefficient lies in the table's range of 1 to 7, the pairing of coefficient with day-name group matches the table's own pairing in 59–62 % of cases against a chance level of 50 %, and all eight ten-level almanacs with coefficient 7 use the odd day-name group — the pairing the reconstructed eighth column carries. The single Paris Codex almanac in the corpus, on pages 17c–18c, pairs coefficient 6 with the odd group, a combination the table does not contain; the pairings were evidently a preference of practice rather than a rule.

IV. DISCUSSION

Three claims can be made with the evidence now assembled, and they should be distinguished by strength.

The arithmetical claim is settled. The interval between any two Tzolk'in dates is given by a single closed expression built from two keys, 40 and 39, and every constant in both Paris Codex tables is an instance of it. This is a proof over a finite space, not an inference, and it holds independently of any reading of the codices.

The empirical claim is strong. Two thirds of the almanacs in the surviving codices are built on the two intervals that the page 21 table lays out, and the table's remaining structural constants also occur in the corpus. The alternative reading — that the table is a decorative or purely symbolic display of the calendar — does not explain why its two circuits should coincide with the two dominant almanac formats, nor why its internal steps should be attested as level intervals as well. The frame-level re-extraction sharpens the point: the table's three native structures cover 71 % of the corpus, almanac level starts are measurably drawn from its cells, and its two internal steps are the two ladder orders of scribal practice.

The interpretative claim is plausible and no more. That Maya priests used this particular page as a working aid, in the manner of a notebook, remains an inference from function to practice. What the project has established is that the page has a function of the kind claimed, and that the function fits the documented shape of the corpus. Whether page 21 was consulted, or was a record of arithmetic already internalised, the evidence does not decide — and the model itself suggests the second possibility, since a scribe holding the two keys would not have needed the table.

For pages 23–24 the position is similar. The reconstruction of the lower table rests on physical traces plus an arithmetical constraint that the traces satisfy, which is a strong basis. The 14-day interlock between the two registers follows from the reconstruction and is verifiable. The archaeoastronomical reading — constellations displacing over the western horizon, eclipse glyphs marking the sun's disappearance — is consistent with the calendrical structure and with comparable passages in the Dresden and Madrid codices, but it is an interpretation of iconography and carries the uncertainty proper to that.

V. CONCLUSION

The project delivered a complete relational record of the calendrical content of the four surviving Maya codices, analyses of both target sections of the Paris Codex, and a general model of Tzolk'in interval arithmetic that reduces the constants of those sections to a single verified expression.

The substantive shift in understanding concerns page 21. A page read for a century and a half as a display of the calendar is better understood as a pre-computed extract of interval arithmetic, tuned to the two almanac formats that dominate the corpus. The identification of a second calendrical table on pages 23–24, and of the 14-day interlock between it and the long-known upper table, revises the reading of one of the most frequently reproduced passages in Maya studies.

Of the three items earlier flagged for resolution, two are settled. The internal column-step pattern is reconciled: the two steps are the table's two ladder orders, corresponding to the two five-level almanac formats attested in the corpus, and the re-checked transcription confirms the assignment. The extent of the table is settled at eight columns — the arithmetical thirteen-column extension is an idealisation that the corrected transcription does not support — with the coefficient sequence reading 1, 2, 3, 4, 5, 6, 6, 7 and the eighth column reconstructed at coefficient 7, starting day Ix, as a hypothesis awaiting a surviving glyph. The day-name orthography used across the database and its derived datasets still requires normalisation; it affects presentation, not results.

The database, the derived datasets and the reconstruction tool are published alongside the project reports under open-science principles, so that the corpus statistics reported here can be re-derived independently.