

Statistical Structure of the Historical Orderings of the I Ching Hexagrams

Pair Rule, Family Gradient, and the Limits of Demonstrability

Alexis García Hurtado*

July 25, 2026

Abstract

The King Wen sequence orders the 64 hexagrams of the I Ching in an arrangement that has resisted explanation for three millennia. We present a statistical comparison of the three historical orderings (King Wen, Mawangdui, Jing Fang) against the binary ordering and against one another. Our central findings are the following. (1) All three historical orderings sit at the random-expected Kendall distance from the binary order (1013, 1008 and 1008 inversions, where the expectation is $n(n-1)/4 = 1008$): none is rank-correlated with binary counting, and two of the three achieve Kendall tau exactly zero. (2) Yet King Wen and Mawangdui are significantly correlated with each other (759 inversions, $z = -2.89$, Monte Carlo $p = 0.0034$, Bonferroni-surviving), and we identify the mechanism: a shared gradient of upper-trigram families, explicit in Mawangdui's octets and implicit in King Wen. (3) Re-analyzing recent Monte Carlo results (Chan, 2026) under a pair-preserving null shows three of the four reported signatures to be corollaries of the pair rule; we also report a small numerical correction, triangulated against a Lean-verified theorem (Radisic, 2026). (4) For the open question of within-pair orientation, the strongest admissible criterion is the classical doctrine of correct centers (12/16, uncorrected one-tailed $p = 0.038$; two-sided 0.077), and we show that no orientation rule of this strength can reach significance over 28 fixed pairs: a measurable limit of demonstrability. Every numerical claim in this paper is asserted by an automated test suite in a public repository.

1 Introduction

The I Ching (Yijing, or Classic of Changes) is built on 64 hexagrams: figures of six stacked lines, each line either broken (yin) or solid (yang). Mathematically, the hexagrams are precisely the elements of the six-dimensional Boolean cube $\{0,1\}^6$, a fact whose earliest documented Western recognition is Leibniz's reading of Shao Yong's circular diagram, transmitted to him by the Jesuit Joachim Bouvet (Leibniz, 1703; Ryan, 1996). The received text arranges the 64 hexagrams in the King Wen sequence, an ordering that proceeds in 32 pairs: each pair is related by 180-degree rotation, or by complementation when the figure is rotation-symmetric. Its organizing principle, beyond the pair rule itself, has been an open question for roughly three thousand years. The tradition's own answer, the Xugua commentary of the Ten Wings, is narrative rather than structural; modern surveys state that the reasoning behind the sequence, if any, is unknown.

*Independent Researcher, Gatineau, Québec, Canada. Email: research@theoriginaliching.com. ORCID: 0009-0003-4636-8206.

Two further historical orderings survive. The Mawangdui silk manuscript (168 BCE), the oldest physical arrangement of all 64 hexagrams, organizes them into eight octets by upper trigram in a fixed family order (Shaughnessy, 1996, 2014). The Eight Palaces arrangement attributed to Jing Fang (1st century BCE) generates each palace from a doubled trigram through a fixed schedule of line changes (Nielsen, 2003). Together with the binary ordering and, as an engineered reference, the binary-reflected Gray code, these give five total orders on the same 64 objects, and a natural comparative question that, to our knowledge, has not been asked quantitatively: how far apart are these orderings from one another, and what structure does each carry?

The question has become timely. In 2026 two independent preprints initiated a formal-methods literature on the King Wen sequence: a Monte Carlo statistical characterization against unrestricted permutation baselines (Chan, 2026), and a Lean-verified theorem showing that the King Wen pairing is the unique cost-minimal K_4 -equivariant matching on the 6-cube under a simple reverse-priority rule (Radisic, 2026). The present paper is complementary to both: where Chan studies the King Wen sequence against free shuffles and Radisic characterizes the pairing, we study the historical orderings against one another, re-analyze Chan’s signatures under a sharper conditional null, and take up the question that Radisic’s theorem leaves exactly delimited: not which hexagrams are paired (now a theorem), but which member of each pair comes first.

Our contributions:

C1 (exact-randomness tie). The three historical orderings lie at Kendall distance 1013, 1008 and 1008 from the binary ordering, against an expectation of exactly 1008: none is rank-correlated with binary counting, and the qualitative consensus that these arrangements are “not mathematical” in the binary sense acquires a precise quantitative form: two of the three sit exactly at the null expectation, and the third is consistent with zero (Section 3).

C2 (kinship and mechanism). King Wen and Mawangdui are significantly closer to each other than chance (759 inversions, $z = -2.89$), while neither is close to Jing Fang. The correlation is not inherited from the pair structure (the 32 King Wen pairs are never adjacent in Mawangdui, and never share an octet); it is explained almost entirely by a shared gradient of upper-trigram families, which Mawangdui makes explicit and King Wen carries implicitly (Section 4).

C3 (conditional re-analysis and a correction). Under a null that preserves the pair rule, three of the four significant signatures reported by Chan (2026) are corollaries of the pair rule, and the fourth is marginal. We also correct one reported statistic (mean within-pair Hamming distance: 3.75, not 3.56), with the correction triangulated by direct enumeration, independent re-implementation, and the Lean-verified total cost of Radisic (2026), which equals $120 = 32 \times 3.75$ (Section 5).

C4 (orientation and the limits of demonstrability). For the within-pair orientation question, declared open in the survey literature, we test a battery of structural criteria. A single coherent signal survives: the pair tends to open with the member whose fifth line is yang and second line is yin, the classical doctrine of correct centers, at 12/16 (uncorrected one-tailed $p = 0.038$; two-sided 0.077). We show this criterion coincides exactly, on all co-decidable pairs, with an independently proposed nuclear-hexagram rule, and that with 28 pairs fixed forever, no rule of this strength can be demonstrated or refuted (Section 5).

C5 (spectral portraits). A common-signal spectral analysis yields a comparative fingerprint of the five orderings, the triple of Walsh, Fourier and Haar summaries, including a striking characterization of Jing Fang as the anti-linear order: zero energy in first-order Walsh characters and 74.5 percent concentrated in fourth-order interactions (Section 6).

A note on scope: this paper makes no numerological claims. Interpretive traditions that attach the golden ratio, quantum mechanics or market behavior to the hexagrams are outside its subject, and where our findings touch traditional doctrines (trigram families, correct centers), the statistics are ours and the doctrinal framing is cited as tradition, not asserted as cause.

2 Conventions, data and reproducibility

Encoding. Lines are numbered 1 (bottom) to 6 (top); yang = 1, yin = 0. We encode a hexagram as the 6-bit integer whose most significant bit is line 1, giving a bijection with $\{0, \dots, 63\}$ (Kun = 0, Qian = 63). Conventions differ across the literature (Chan, 2026, and Radisic, 2026, use least-significant-bottom encodings); all results reported here are invariant under the choice of encoding except where a real-valued signal is built from hexagram values, in which case the convention is declared at the point of use.

The orderings. King Wen is the received ordering of the standard text. Mawangdui follows the standard reconstruction of the silk manuscript (Shaughnessy, 1996), which builds on the Chinese editio princeps: the transcription published by the Mawangdui editing group (1984) and the early studies of Yu Haoliang (1984) and Han Zhongmin (1992); the lacunae of the physical text and the reconstruction choices they force are matters for philology, and this analysis takes the reconstructed sequence as given. It runs: eight octets by upper trigram in the family order Qian, Gen, Kan, Zhen, Kun, Dui, Li, Xun, each octet headed by the doubled hexagram and continued by the remaining lower trigrams in the fixed order Qian, Kun, Gen, Dui, Kan, Li, Zhen, Xun (Shaughnessy, 1996, p. 17, for the two trigram orders; pp. 28-29 for the sequence itself). Jing Fang follows the Eight Palaces (Nielsen, 2003): palaces in the received bagong palace order (Nielsen, 2003, p. 3, Table 2, entry BA GONG GUA) Qian, Zhen, Kan, Gen, Kun, Xun, Li, Dui; within a palace, the pure hexagram, then five generations flipping lines 1 through 5 cumulatively from the bottom, then the wandering soul (re-flipping line 4) and the returning soul (flipping the entire lower trigram). The binary ordering lists hexagrams by integer value; the Gray ordering is the standard binary-reflected Gray code, included as an engineered benchmark of local smoothness. Both reconstructed orderings have been checked position by position against the printed sources rather than against our own reading of them: Mawangdui against Shaughnessy (1996), and Jing Fang against Nielsen (2003), each in all 64 positions and with no discrepancy. Those two sources are printed, third-party and independent of this work; the verification suite of the replication package is not.

Statistics. The distance between two orderings is the Kendall inversion count: the number of pairs of hexagrams whose relative order disagrees. We write inversion count for the raw number and Kendall distance for its metric name; Kendall tau is its normalization to $[-1, 1]$. For two orderings in uniform random relative position, the count has expectation $n(n-1)/4 = 1008$ and standard deviation $\sqrt{n(n-1)(2n+5)/72} \approx 86.3$ for $n = 64$. Monte Carlo protocols use fixed seeds and sample sizes stated at each use, with 20,000 samples throughout: at that size the standard error of a Monte Carlo percentile stays below 0.4 points even at its worst, near the median, which is finer than the resolutions reported here. Seed 424242 covers the permutation nulls of Sections 4 and 5, and seed 20260722 the percentile protocols of Section 5. These are the values in the replication package,

and `verify_paper.py` is their executable source. p -values are reported with their sidedness, and multiple-comparison corrections are applied and stated wherever a battery of tests is involved.

Reproducibility. Every numerical claim in this paper is asserted by an automated suite of 63 verification sections in a public repository (exit status zero is a publication gate for the accompanying interactive laboratory). Throughout the project each computation was independently re-derived at least twice, by separate implementations with different seeds; this protocol caught and corrected errors in both directions during development, and the correction reported in Section 5 was found the same way.

3 Distance to the binary ordering: an exact-randomness tie

Result 3.1. The Kendall inversion counts between the historical orderings and the binary ordering are: King Wen 1013, Mawangdui 1008, Jing Fang 1008, out of a maximum of 2016. The random expectation is exactly 1008, with standard deviation 86.3. All three counts lie within 0.06 standard deviations of the expectation; Mawangdui and Jing Fang sit on it exactly, giving Kendall tau precisely zero, and King Wen gives $\tau = -0.005$. Table 1 collects these counts.

The reading is sharper than a tie: the three historical orderings are uncorrelated with binary counting. Mawangdui and Jing Fang sit exactly at the null expectation, which is arithmetic on fixed orderings rather than an estimate; King Wen, at 1013, is consistent with zero correlation ($z = +0.06$). This turns two qualitative commonplaces into quantitative ones. First, the historical consensus that the binary reading of the hexagrams is a later mathematical imposition, not a principle of the classical arrangements (Ryan, 1996), acquires a number: the arrangements are as far from binary order as pure chance would place them. Second, the survey judgment that the Mawangdui sequence “probably has no great mathematical significance” is both confirmed in the binary sense (tau exactly zero) and, as Section 4 shows, refuted in another: Mawangdui carries structure, just not binary structure.

A companion metric. Orderings that are equally uncorrelated with binary value can differ greatly in local geometry. The total adjacent Hamming cost (sum of line changes between consecutive hexagrams) separates the five orders cleanly: Gray 63 (the theoretical minimum, one line per step), Jing Fang 93, binary 120, Mawangdui 141, King Wen 211. The King Wen sequence is by far the roughest walk on the cube, a property consistent with the high transition distances reported by Chan (2026) and analyzed conditionally in Section 5.

4 The kinship of the two oldest orderings, and its mechanism

4.1 The kinship

Result 4.1. The pairwise Kendall inversion counts between the historical orderings are: King Wen to Mawangdui 759 ($z = -2.89$), King Wen to Jing Fang 909 ($z = -1.15$), Mawangdui to Jing Fang 872 ($z = -1.58$). Under a Monte Carlo null of uniform relative order (20,000 samples, fixed seed), the King Wen to Mawangdui count gives $p = 0.0034$ (two-sided), which survives a Bonferroni correction for the three comparisons ($0.0102 < 0.05$). The correction reported takes the three comparisons among the historical orderings as the family. Under the more demanding correction that treats all ten pairwise comparisons among the five orderings as the family, the result

Table 1: The five orderings against the binary order and by local geometry. Cells marked “–” are not reported in this study.

Ordering	Inversions vs. binary	Kendall τ	Hamming cost	Dominant principle
Binary	0	1.000	120	binary value
Gray	–	–	63	Gray code
Mawangdui	1008	0.000	141	upper-trigram octets
King Wen	1013	–0.005	211	rotation pair rule
Jing Fang	1008	0.000	93	palace generation

The Gray code is a cyclic construction rather than a linear ranking of the 64 objects, so an inversion count against the binary order would not be commensurable with the others and is not reported; its Hamming cost, which is what the code is built to minimize, is.

still holds: $0.0034 \times 10 = 0.034 < 0.05$. `verify_paper.py` checks both calculations. The other two comparisons are not significant.

The picture is asymmetric in an instructive way. Section 3 showed that none of the three historical orderings is correlated with binary counting; Result 4.1 shows that the two oldest of them are nonetheless correlated with each other, while the algorithmically generated Eight Palaces arrangement stands apart from both. In genealogical terms: King Wen and Mawangdui are relatives, and Jing Fang is the loner.

4.2 Excluding the obvious mechanism

The natural first hypothesis is that Mawangdui inherits the King Wen pair structure. It does not, and the exclusion is total: of the 32 King Wen pairs, zero are adjacent in the Mawangdui ordering, and zero fall within the same Mawangdui octet. The reason is structural rather than accidental. Rotating a hexagram by 180 degrees sends its upper trigram to the reversal of its lower trigram; since Mawangdui groups by upper trigram, pair partners are separated by construction whenever these differ, which is the typical case. The mean positional separation between King Wen pair partners inside Mawangdui is 24.4, slightly above the random expectation of $(n + 1)/3 = 21.7$: the octet organization actively splits the pairs. Whatever makes the two orderings close, it is not the pairs.

4.3 The mechanism: a shared family gradient

Decomposing the 759 inversions by Mawangdui’s octet structure localizes the correlation. Of the 224 hexagram pairs lying within a common octet, 95 are discordant between the two orderings (95/224, expectation 112); of the 1792 pairs spanning two octets, 664 are discordant (664/1792, expectation 896). The deficit of roughly 249 inversions therefore lives overwhelmingly between octets: at the level of upper-trigram families, not inside them.

The family-level structure is visible directly. The Mawangdui octet sequence follows the traditional trigram families: the father and sons (Qian, Gen, Kan, Zhen), then the mother and daughters (Kun, Dui, Li, Xun) (Shaughnessy, 1996; Nielsen, 2003). Computing the mean King Wen position of each Mawangdui octet gives, for the male line, 18.0, 26.6, 31.9 and 43.0: perfectly monotone in the family order (a specific monotone arrangement of four means has probability $1/24 \approx 0.042$ under exchangeability). The mother’s octet averages 20.0, early, and the three daughters average 39.8, 38.5 and 42.2, late. The Spearman correlation between Mawangdui octet index and mean

King Wen position is 0.619 (exact permutation $p = 0.0575$ over the $8!$ octet orders).

Two conditional nulls quantify how much this family gradient explains. Shuffling hexagrams within octets while preserving the octet order yields a mean of 776 inversions (sd 11), and the observed 759 sits at $z = -1.51$ of that null: not significant. Shuffling the octet order while preserving octet interiors yields a mean of 991 (sd 131). Fixing the family order alone thus moves the expectation from 1008 to 776, capturing the bulk of the observed deficit, and the residual within-family alignment is not significant.

Conclusion 4.2. The kinship between the two oldest orderings is explained, almost entirely, by a shared gradient of upper-trigram families: explicit in Mawangdui, which sorts by it, and implicit in King Wen, which carries it as a statistical tendency. What the statistics identify is structure at the level of families; the monotone gradient of octet means is the form in which that structure shows up here. We state the historical reading with care: the trigram-family doctrine is documented tradition (Nielsen, 2003), and our finding is consistent with both orderings drawing on a common family convention; the statistics establish the shared gradient, not its causal direction. This also refines the survey judgment quoted in Section 3: the Mawangdui arrangement is not mathematically arbitrary; it is uncorrelated with binary value yet measurably aligned, at the family level, with the received sequence.

5 The pair structure: conditional re-analysis, a correction, and the limits of demonstrability

5.1 Chan’s signatures under a pair-preserving null

Chan (2026) reports four statistically significant properties of the King Wen sequence against unrestricted permutation baselines: high mean transition distance (98.2nd percentile), negative lag-1 autocorrelation of transition distances ($p = 0.037$), an excess of yang-balanced groups of four ($p = 0.002$), and asymmetric within-pair versus between-pair distances (99.2nd percentile). We first reproduced these results independently under free shuffles: mean transition distance 3.349 at the 97.9th percentile; autocorrelation -0.247 against Chan’s -0.251 ; exactly 7 yang-balanced groups, at the 99.2nd percentile.

The King Wen sequence, however, is not a free permutation: it is constrained by its pair rule. The sharper question is which of these signatures remain once the null respects that rule. We therefore re-evaluated all four statistics under a conditional null that permutes the 32 pairs and flips orientations within them (20,000 samples). This is the maximum-entropy null (the uniform distribution over all pair-preserving rearrangements) that preserves the pair partition as structure: it fixes exactly the information the pair rule asserts, which pairs exist, and randomizes everything the rule leaves free, their order and their orientations. All conditional results are relative to this natural randomization; weighted or further constrained pair-preserving variants could in principle be probed, and we record this as the standard caveat of conditional inference (probed directly in Appendix A). The comparison is therefore against the sequence’s own constraint rather than against no constraint at all:

(Percentile convention: $100 P(\text{statistic} < \text{observed})$; on integer-valued statistics, $P(<)$ versus $P(\leq)$ shifts pair-null percentiles by a fraction of a point, e.g. 29.0 versus 29.4 across our two implementations.)

Table 2: Chan’s four signatures re-evaluated under a pair-preserving null. For the yang-balanced groups, the pair-preserving null raises the expectation from 2.6 to 4.4.

Chan’s statistic	value	free	pair-null	verdict
mean transition distance	3.349	97.9	29.0	corollary of the pair rule
lag-1 autocorrelation of distances	−0.247	4.0	6.4	marginal, not significant
yang-balanced groups of four	7	99.2	89.6	largely induced by the pair structure; residual not significant
within/between-pair asymmetry	3.75 vs. 2.94	99.9	invariant (sd 0)	corollary by construction

Conclusion 5.1. Three of the four reported signatures are corollaries of the pair rule; the fourth, the alternation of transition distances, retains a marginal, non-significant lean under the conditional null (see also Appendix A for its behavior at the top of the conditional ladder). The two analyses are complementary rather than contradictory: Chan establishes that the King Wen sequence is not a free shuffle; the conditional analysis establishes that, given the pair rule, little additional structure remains under the tested statistics.

5.2 A numerical correction, triangulated

One reported figure does not reproduce. The mean Hamming distance within the 32 King Wen pairs, computed by direct enumeration under the natural pairing (consecutive positions $2k - 1, 2k$), is exactly 3.75, that is $120/32$; the between-pair mean is 2.94, agreeing with Chan’s reported 2.94. Chan reports the within-pair mean as 3.56. The correction is triangulated three ways: by direct enumeration, by an independent re-implementation in a separate codebase (the laboratory’s verification suite), and by the formally verified theorem of Radisic (2026), whose Lean 4 proof establishes that the total Hamming cost of the King Wen matching is exactly 120, hence a mean of $120/32 = 3.75$. Hamming distance is invariant under any fixed relabeling of bit positions, so no encoding convention can produce 3.56: the discrepancy can only lie in the pairing or in the metric employed. The most likely explanation is therefore a differing pairing or metric convention in the original implementation; the qualitative conclusion, that pair partners are internally more distant than between-pair neighbors, is identical in both analyses.

5.3 The floor: what is now a theorem, and what remains open

Radisic (2026) proves that the King Wen pairing is isomorphic to the unique cost-minimal K_4 -equivariant perfect matching on $\{0, 1\}^6$ among comp/rev pairings, generated by a uniform reverse-priority rule (pair with the reversal unless palindromic, then with the complement), at cost 120 versus 192 for the complement-only matching. The question of which hexagrams are paired is therefore settled as mathematics. What the theorem leaves exactly delimited is the orientation question: within each pair, which member comes first? The survey literature states that no discernible rule is known, and the question has not, to our knowledge, been asked quantitatively; qualitatively, Hacker (1993) anticipated that correctness bears on within-pair position. This is the question we

Table 3: The three mirror components; the signal is specific to the centers. The centers value is 0.038 one-tailed and 0.077 two-sided.

component	correct-first	p (one-tailed)
extremes (1,6)	10/16	0.227
middles (3,4)	9/16	0.402
centers (2,5)	12/16	0.038

now take up.

5.4 A battery of structural criteria

A first observation removes one candidate by theorem: 180-degree rotation preserves the number of yang lines, so the criterion “the member with more yang comes first” is undecidable on all 28 rotation pairs (28/28 ties). Five structural criteria were then tested on the 28 non-palindromic pairs, each with its exact binomial p -value (one-tailed, toward the observed direction): greater binary value first, 14/28 ($p = 0.575$); bottom line yang first, 16/28 ($p = 0.286$); top line yang first, 12/28 ($p = 0.286$); orientation smooths the incoming transition from the previous hexagram, 7/15 with 13 ties ($p = 0.500$); orientation smooths the outgoing transition, 7/15 with 12 ties ($p = 0.500$). All are compatible with a fair coin. The criteria are structural and fixed by the geometry of the problem, hexagram value, the extreme lines, local smoothing, the nuclear hexagram and the components of the rotation; they are not the product of an open search over the space of rules. The direction of the one-tailed test (correct configuration opens) is documented in the qualitative literature three decades before this analysis (Hacker, 1993).

An independently published proposal holds that the nuclear hexagram (built from lines 2 through 5) decides the orientation (the positional prominence of the nuclear hexagrams in the received text, where 1, 2, 63 and 64 open and close the sequence, was recorded by Hacker, 1993, pp. 88-89 and 109). Its simplest form, greater nuclear value first, scores 8 of 24 decidable pairs (4 ties), that is 16/24 toward smaller-nuclear-first, $p = 0.076$: the strongest lean so far, suggestive, and not a confirmation. As a companion fact, in 16 of the 28 pairs the two members fall into different basins of the nuclear-reduction map, the basin of a hexagram being the attractor, a fixed point or a 2-cycle, to which repeated application of the nuclear operator converges.

5.5 The signal of the centers

Rotation pairs the line positions as (1,6), (3,4) and (2,5). For each mirror component, define the correct configuration as yang in the odd position and yin in the even one (the classical dangwei doctrine of appropriateness), and test whether the pair opens with the correctly configured member on the pairs where the component decides:

The signal is specific to the centers: line 5, the ruler’s place, and line 2, the minister’s. The second member of a pair is the rotation of the first, so “yang in line 5 of the first” and “yin in line 2 of the first” are the same event on decidable pairs. The battery’s one surviving criterion is therefore a single rule, the doctrine of correct centers (zhong zheng). Its canonical status is not our construction: the Stanford Encyclopedia of Philosophy singles out Jiji (hexagram 63) as the best hexagram in terms of yin-yang correspondence (Hon, 2023). Jiji opens its pair. The doctrine is, historically, a doctrine for interpreting hexagrams rather than for ordering them; applying it to the orientation of the sequence is our hypothesis under test, with the doctrine contributing its

Table 4: The nine tested criteria of the battery, with the results and exact one-tailed binomial p -values reported in Sections 5.4 and 5.5. Criterion (vi) is stated as greater-nuclear-first, which scores 8 of the 24 decidable pairs; the lean is toward the opposite direction, 16/24, and the p -value is that of the observed direction. Rows (vii) to (ix) restate Table 3.

	critereon	result	p
(i)	greater binary value first	14/28	0.575
(ii)	bottom line yang first	16/28	0.286
(iii)	top line yang first	12/28	0.286
(iv)	orientation smooths the incoming transition	7/15 (13 ties)	0.500
(v)	orientation smooths the outgoing transition	7/15 (12 ties)	0.500
(vi)	greater nuclear value first	8/24 (4 ties)	0.076
(vii)	correct configuration first, extremes (1,6)	10/16	0.227
(viii)	correct configuration first, middles (3,4)	9/16	0.402
(ix)	correct configuration first, centers (2,5)	12/16	0.038

a-priori standing and not a historical claim about the compilers' intent. The connection is not new as a qualitative observation: Hacker (1993, pp. 105-108) records deterministic partial rules linking line correctness to the odd or even position of a hexagram in the received text, noting that correctness has "some bearing on the textual sequence"; the present criterion is the centers form of that doctrine, evaluated statistically under the conditional null. There is also a purely geometric reading: lines 2 and 5 are the centers of the two trigrams, distinguished by the rotation map itself, so the criterion admits a structural reading that requires no historical intent. Moreover, on all 16 pairs decidable by both, the centers criterion and the smaller-nuclear criterion agree 16 out of 16 (enumerated in the replication package): the published nuclear proposal and the classical doctrine are, on this data, one rule seen through two lenses, as expected given that the nuclear hexagram is built precisely from lines 2 through 5. Both criteria read the interior lines of the figure, so the agreement, while not automatic (their domains of decidability differ), is less surprising than it first appears; it remains a verified enumeration and not a theorem.

5.6 The dual reading, and the limit of demonstrability

The statistical status of the centers rule must be stated both ways. Read as a single a-priori hypothesis, and it is the single most canonical structural doctrine the tradition offers, the result is 12/16 with $p = 0.038$ (one-tailed, uncorrected; two-sided: 0.077). Read as the best outcome of the battery tested across this study, it does not survive correction. The battery comprises nine criteria, listed with their results in Table 4. A tenth candidate, more-yang-first, is excluded by theorem rather than tested, being undecidable on all 28 pairs. With nine tests, $0.038 \times 9 = 0.342$.

There is, finally, a structural fact that no methodology can evade: the King Wen sequence contains exactly 28 rotation pairs, and there will never be more data. A rule that decides 12 of 16 informative cases can be neither demonstrated nor refuted at conventional significance on a sample fixed forever at this size. For scale: a criterion deciding all 28 pairs would need at least 22 agreements to survive the battery correction, and one deciding 16, at least 14. Our answer to the three-thousand-year-old orientation question is therefore exact about its own boundary: if the King Wen sequence has an orientation rule, the best candidate the mathematics can identify is the doctrine of correct centers, coinciding with the nuclear proposal; and the question is now answered up to the precise point where statistics stops being able to decide.

Table 5: Spectral fingerprints of the five orderings.

ordering	Walsh energy by order (%)	DFT k	Haar width
Binary	100 / 0 / 0 / 0 / 0 / 0	1	64
Gray	75 / 25 / 0 / 0 / 0 / 0	1	64
Mawangdui	20.9 / 34.4 / 14.1 / 22.3 / 7.3 / 1.2	24	4
King Wen	21.5 / 16.8 / 28.4 / 25.7 / 5.1 / 2.5	16	4
Jing Fang	0.0 / 15.3 / 5.3 / 74.5 / 3.8 / 1.2	7	32




Figure 1: Walsh energy by interaction order for the five orderings (percent of non-DC energy). Binary is a pure first-order tone; Jing Fang is anti-linear, with 74.5 percent in fourth-order interactions and none in the first.

6 Spectral portraits of the five orderings

The comparisons so far are rank-based. A complementary view treats each ordering as a real-valued signal and asks what structure it carries in three classical bases of 64-point analysis: the Walsh characters of the group $(\mathbb{Z}/2)^6$, the discrete Fourier basis of the cycle $\mathbb{Z}/64$, and the Haar wavelet basis. For cross-order comparability we use one common signal per ordering: position maps to the binary value of the hexagram occupying it, centered.¹

Result 6.1 (fingerprints). All spectral quantities below are exact linear-algebra computations on the fixed orderings: no Monte Carlo is involved and there is no variance to report. Percentages are of non-DC Walsh energy by interaction order 1 through 6; the DFT column gives the dominant harmonic k in $1 \dots 32$; the Haar column gives the width of the single wavelet of largest normalized energy. These are the columns of Table 5. Figure 1 plots the Walsh energies.

Three observations. First, the binary and Gray orderings behave as calibration: the binary signal is exactly linear (all energy in order 1), and Gray is nearly so. Second, the Jing Fang

¹This convention necessarily differs from the natural single-order signal, position maps to King Wen number, used when analyzing the King Wen sequence alone; with that signal, the Walsh spectrum of the King Wen ordering concentrates 77.4 percent of its energy in even interaction orders, against 47.6 percent expected for orders $\{2, 4\}$ under uniformity, a spectral confirmation of the pair rule. Both conventions are declared and asserted separately.

arrangement is the anti-linear order: its signal contains no first-order component at all, and concentrates 74.5 percent of its energy in fourth-order interactions, a signature produced by the palace generation schedule (cumulative flips plus the two soul transformations); the loner of Section 4 also has the strangest spectrum, and the fingerprint shows in what way. Third, the dominant cyclic harmonic of the King Wen signal is $k = 16$, that is period 4: the same scale as the yang-balanced groups of four reported by Chan (2026). Two independent instruments, a time-domain count and a Fourier decomposition, point to the same four-hexagram periodicity in the received sequence; we record the convergence without proposing a mechanism.

7 Related work

Formal analysis of the King Wen sequence begins, in the current wave, with Chan (2026), who characterizes the sequence against 100,000 free permutations and reports the four signatures re-analyzed in Section 5, alongside uniformly negative results on their utility for neural network training; and Radisic (2026), who settles the pairing question with a Lean-verified optimality theorem. Our contribution is orthogonal to both: comparative permutation statistics across the historical orderings, a conditional re-analysis, a mechanism for the inter-order kinship, and the orientation battery.

The comparative study of the orderings themselves has a philological literature: Shaughnessy (1996) reconstructs and translates the Mawangdui manuscript, Shaughnessy (2014) situates it among the recently unearthed Changes manuscripts, Rutt (1996) provides a philological study of the received text as a Bronze Age document, and Han Zhongmin (1992) offers a qualitative comparison of the silk-manuscript hexagram sequence with the received order. The trigram-family doctrine underlying Section 4 and the Eight Palaces construction are documented in Nielsen (2003). The binary reading of the diagrams and its history run through Leibniz (1703) and Ryan (1996). Olsvanger (1948), as documented by Hacker (1993, pp. 109-112), found sum-symmetric sectional structure in the received sequence written over the binary values: symmetric sections each totalling 1008.² Sum symmetries of this kind concern totals, not order, and are compatible with the zero rank correlation of Section 3. Earlier quantitative observations exist outside the academic literature: the even/odd profile of first-order differences was tabulated by McKenna and McKenna (1975/1994), discussed in Hacker (1993, p. 108), within an interpretive framework we do not adopt; an independently published proposal that nuclear hexagrams decide pair orientation is tested, and not confirmed in its simple form, in Section 5; and the cycle type $[52, 10, 2]$ of the King Wen permutation relative to binary order, noted on an independent mathematics website, reproduces under our conventions. A companion preprint (García Hurtado, 2026) gives an order-theoretic reading of the ladder of conditional nulls of Appendix A, and reports a verification of the Mawangdui and Jing Fang constructions against their printed sources; the corrections recorded in the errata of this package were found in the course of that work, and each is established from this package alone. The tradition’s own account of the sequence, the Xugua commentary, is narrative and is not treated here as a structural hypothesis, because it provides no formalizable rule.

8 Reproducibility

The primary artifact accompanying this paper is a self-contained replication package:

<https://github.com/theoriginaliching/kingwen-orderings-replication>

²The recurrence of 1008 is an arithmetic coincidence: the sum of the values 0 to 63 is 2016, and the number of hexagram pairs $\binom{64}{2}$ is also 2016; both halve to 1008.

<https://paper.theoriginaliching.com>

(the repository is the primary artifact; the second address is a human-readable gateway to the same package. Archived: the concept DOI <https://doi.org/10.5281/zenodo.21609653> always resolves to the latest version, and the figures in this paper were checked against the version whose own DOI is <https://doi.org/10.5281/zenodo.21776041>). It contains the manuscript source, the compiled PDF and a single script, `verify_paper.py`, which reproduces every numerical claim made here from first principles, using only the Python standard library, with no network access and no data files. The only external datum is the received King Wen sequence, embedded as a literal table; every other ordering is generated from the construction rules given in Section 2, so the package verifies the constructions as well as the statistics. It prints one line per check, showing the reproduced value beside the value printed in this paper, and exits zero if and only if all of them pass:

```
python3 verify_paper.py    tectonic paper.tex
```

As a secondary resource, the same material is explored in an extended interactive laboratory of 45 experiments at <https://experiments.theoriginaliching.com> (in Spanish; an English version is in progress), whose assertion suite comprises 63 sections executed as a publication gate (exit status zero required); computed statistics are frozen as assertions once verified, so that no figure can drift silently.

Throughout the project, a dual-derivation protocol was followed: each result was implemented at least twice, independently and with different random seeds, before publication. The protocol is not ceremonial; it caught and corrected errors in both directions during development, including an early correction of the palace generation order against the documented bagong sequence, a sampling-length error in one of our own preliminary p -values, and the discrepancy reported in Section 5. Monte Carlo results in this paper state their sample sizes and seeds; percentile conventions on integer statistics are declared where they matter.

9 Scope and limitations

This paper makes claims about orderings as combinatorial objects, and only those. It makes no numerological claims: readings that attach the golden ratio to the King Wen sequence, quantum mechanics to the oracle, or market behavior to the hexagrams fall outside its subject and, in our assessment, outside demonstrability. Three limitations are intrinsic. First, the King Wen sequence is a single fixed object; analyses of its internal structure (Sections 5.4 to 5.6) face a sample size that will never grow, and we have stated the resulting limit of demonstrability explicitly. Second, the interpretation of the family gradient as a shared convention (Section 4.3) is historical framing for a statistical finding; the direction of influence between the orderings, or from a common source, is not established here. Third, signal-based results (Section 6) depend on declared encoding conventions; rank-based results (Sections 3 to 5) do not.

10 Conclusion

Measured against binary counting, the three historical orderings of the hexagrams are uncorrelated with binary counting: two of them exactly at the null expectation, the third consistent with zero. Measured against one another, the two oldest recognize each other: not through their pairs, which Mawangdui's octets actively separate, but through a shared gradient of trigram families that the

manuscript makes explicit and the received sequence carries as a tendency. And measured against its own last open question, the King Wen sequence yields one coherent candidate rule, the doctrine of correct centers, standing exactly at the boundary where twenty-eight pairs of evidence stop being able to decide. Three millennia asked what principle orders the Book of Changes. Part of the answer, it turns out, is a precise description of how far statistics can take it: no criterion of the observed strength can be demonstrated over the 28 fixed pairs.

Acknowledgments

The verification suite, exploratory computations and a first draft of this manuscript were prepared with the assistance of an AI system (Anthropic’s Claude). Every numerical claim was independently re-derived and is asserted by the automated test suite described in Section 8; the author takes full responsibility for the content.

References

- [1] Chan, A. (2026). Statistical properties of the King Wen sequence: An anti-habituating structure that does not improve neural network training [Preprint]. arXiv. <https://arxiv.org/abs/2604.09234>
- [2] García Hurtado, A. (2026). Uninformative rungs: an order-theoretic stopping criterion for nested reference sets [Preprint]. Zenodo. <https://doi.org/10.5281/zenodo.21750029>
- [3] Hacker, E. A. (1993). *The I Ching handbook: A practical guide to personal and logical perspectives from the ancient Chinese Book of Changes*. Paradigm Publications.
- [4] Han Zhongmin. (1992). *Boyi shuolüe* [A brief account of the silk-manuscript Changes]. Beijing Shifan daxue chubanshe.
- [5] Hon, T. (2023). Chinese philosophy of change (Yijing). In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/chinese-change/>
- [6] Leibniz, G. W. (1703). Explication de l’arithmétique binaire. *Mémoires de l’Académie Royale des Sciences*, année 1703.
- [7] Mawangdui Han mu boshu zhengli xiaozu [Mawangdui Silk Manuscripts Editing Group]. (1984). Mawangdui boshu “Liushisi gua” shiwen [Transcription of the silk-manuscript “Sixty-four Hexagrams”]. *Wenwu*, 1984(3), 1–8.
- [8] McKenna, T. K., & McKenna, D. J. (1994). *The invisible landscape: Mind, hallucinogens, and the I Ching* (rev. ed.). HarperCollins. (Original work published 1975, Seabury Press.)
- [9] Nielsen, B. (2003). *A companion to Yi jing numerology and cosmology: Chinese studies of images and numbers from Han (202 BCE–220 CE) to Song (960–1279 CE)*. RoutledgeCurzon.
- [10] Olsvanger, I. (1948). *Fu-Hsi: The sage of ancient China*. Massadah. (Record taken from Hacker, 1993, p. 109, n. 14.)
- [11] Radisic, A. (2026). Optimal equivariant matchings on the 6-cube: With an application to the King Wen sequence [Preprint]. arXiv. <https://arxiv.org/abs/2601.07175>

- [12] Rutt, R. (1996). *The book of changes (Zhouyi): A Bronze Age document*. Curzon Press.
- [13] Ryan, J. A. (1996). Leibniz' binary system and Shao Yong's "Yijing". *Philosophy East and West*, 46(1), 59–90. <https://doi.org/10.2307/1399337>
- [14] Shaughnessy, E. L. (1996). *I Ching: The classic of changes: The first English translation of the newly discovered second-century B.C. Mawangdui texts*. Ballantine Books.
- [15] Shaughnessy, E. L. (2014). *Unearthing the changes: Recently discovered manuscripts of the Yi jing (I ching) and related texts*. Columbia University Press.
- [16] Yu Haoliang. (1984). Boshu Zhouyi [The silk-manuscript Zhouyi]. *Wenwu*, 1984(3), 15–24.

Table A1: Chan’s four signatures across the ladder of conditional nulls, in the column order of Table 2: mean transition distance, lag-1 autocorrelation of distances, yang-balanced groups of four, and within/between-pair asymmetry. A cell marked “constant” is not an extreme percentile: the statistic does not vary under that null, so the figure is arithmetic rather than evidence. P0 and P1 restate the free and pair-null columns of Table 2 from this ladder’s own implementation and seed. Percentiles from this implementation’s 20,000-sample runs (seed 20260722); every entry is asserted in the replication package.

Null	transition	alternation	yang-balanced	within/between
P0 free shuffle	97.5	3.8	99.2	99.9
P1 pairs	29.4	6.2	89.8	constant
P2 pairs within canons	31.9	6.5	87.7	constant
P3 pairs, anchors fixed	33.2	10.8	90.4	constant
P4 pairs across blocks	29.9	3.3	constant = 7	constant
P5 orientations only	17.6	3.5	constant = 7	constant

Appendix A: Robustness under a hierarchy of conditional nulls

Section 5.1 recorded the standard caveat of conditional inference: results are relative to the chosen null. This appendix probes that caveat directly. We re-evaluate the four signatures of Table 2 under a ladder of six nulls ordered by containment, an order that is partial rather than a chain: each rung concedes some part of the sequence’s architecture, not uniformly more than the rung before it, and we ask what remains: (P0) the free shuffle; (P1) the pair-preserving null of Section 5.1; (P2) pairs permuted only within the two received canons (15 and 17 pairs), the received division of the text into hexagrams 1-30 and 31-64; (P3) the pair null with the four palindrome pairs, hexagrams (1,2), (27,28), (29,30) and (61,62), fixed in position and orientation (these are the pairs whose members are invariant under rotation and are therefore matched by complementation; detecting complement pairs by value, $a \text{ XOR } b = 63$, would instead catch eight pairs, since in four rotation pairs reversal and complementation coincide); (P4) the sixteen blocks of four consecutive positions permuted among themselves, so that a pair never leaves its block, with the two pairs inside a block free to exchange and all orientations free; and (P5) the received pair order with only the 32 orientations randomized. An order-theoretic reading of this family, and of when a further rung stops being informative, is given in a companion preprint (García Hurtado, 2026). All rungs are Monte Carlo with 20,000 samples and seed 20260722; percentiles are $100 P(\text{statistic} < \text{observed})$.

Three observations. First, conceding the canon split (P2) or the palindrome anchors (P3) changes essentially nothing: the pair rule already carries the explanatory load, and the conditional conclusions of Section 5 are not artifacts of these further structures. Second, from P4 upward the yang-balanced groups of four cease to be a statistic and become a constant: preserving the block partition preserves each block’s yang sum (in the received sequence the sixteen sums include exactly seven equal to 12), so this signature is a corollary of block membership, in the same sense in which Section 5.1 found the within-pair asymmetry to be a corollary of the pair rule. Third, one suggestive residue survives at the top of the ladder: with everything conceded except orientation, the lag-1 alternation of transition distances sits at the 3.5th percentile, more extreme than under the pair null (6.2), indicating that the received orientations themselves participate in the anti-habituation pattern. The ladder was run as a robustness probe of the four already reported statistics, not as a search; still, this cell is one of the seventeen non-constant entries of Table A1, and the dual reading of Section 5.6 applies here too: as a probe of a single pre-reported signature, 3.5 percent; against

the seventeen cells, $0.035 \times 17 = 0.6$. On either reading it is exploratory and is not claimed as statistically significant: we record it as an observation consistent with that participation, not a finding, tied to the orientation question of Section 5.4. Every figure in Table A1 is asserted in the replication package.