

Forced counts: when a symmetry group determines the discordance of a constructed ordering

Alexis García Hurtado
ORCID 0009-0003-4636-8206

Abstract

We count the **discordant pairs of an ordering against the binary order** of the vertices of the n -cube, in the standard sense of rank correlation, and ask when that count is determined by the subgroup of the hyperoctahedral group B_n that the construction of the ordering respects.

The apparatus is an accounting over the orbits of that subgroup acting on pairs of positions. Each orbit is either **forced**, contributing exactly half its cardinality whatever else happens, or **free**, contributing one of two values. A group can therefore only ever force the midpoint of the range, never any other value; equivalently, it can only force Kendall's tau against the binary order to be exactly zero.

Two theorems remove the apparatus once it has been built. An orbit is forced precisely when it contains as many discordant pairs as concordant ones; and when the subgroup contains all translations, an ordering is forced precisely when every difference class splits exactly in two. A third result cuts the other way: if the subgroup contains the translations by a subspace of dimension at least two, every orbit has even cardinality, so the parity of the count is fixed by the structure, and when that parity differs from the parity of the midpoint the midpoint becomes **impossible** rather than merely unattained. The dimension hypothesis is necessary, by an exhibited witness.

The apparatus and its three theorems are the contribution; what follows is where they came from and what they first settled.

Applied to three constructed orderings of the 6-cube that have come down to us, with 2016 pairs and a midpoint of 1008, the three outcomes all occur. Two of the constructions force the count to 1008, and for one of them the demonstration is complete with no enumerative residue; a mechanism previously claimed for the other is refuted here, since the closure it rests on forces only 484 of the 2016 pairs and survives rearrangements that destroy the result, the midpoint coming out in 3836 of the 40320 of them. The third construction makes 1008 impossible by the parity obstruction: its structure leaves the interval $[957, 1059]$ and exactly 52 compatible totals within it, 1008 is not among them, and the observed count is 1013.

That difference of 5 is decomposed completely over the 19 free orbits and has no culprit: every free orbit deviates, and the deviations nearly cancel. Against a list of candidate structures declared before measuring, one narrows the interval, namely the construction's own pairing involution, which is not affine and so was absent from a group that is exactly its centraliser; none explains the difference. We therefore declare the residue informative relative to that list, and stop.

A landscape of 14 rows over block systems in dimensions three to six populates the three outcomes away from the historical cases and records two refutations of shapes we tried to raise into theorems, one of them of an error of our own.

Every measurement was preceded, in an earlier commit of a public repository, by a written statement of what was going to be measured and, where predictions were made, of what would refute them, and that ordering is checkable.

1. Introduction

1.1 The question

Take an ordering of a finite set and a fixed reference order on the same set, and count the pairs on which the two disagree. If the ordering is not arbitrary but built by a rule, that rule usually has symmetries, and those symmetries constrain the count. The question of this paper is how far the constraint goes:

When is the number of discordant pairs of a constructed ordering determined by the symmetry group that its construction respects?

The answer is not a single dichotomy. A group can leave the count completely free, or pin it down to a single value, or do something in between that is more interesting than either: leave an interval of possible values and, inside that interval, make one particular value **impossible**. Those three outcomes, forced, bounded, and barred, are what the paper measures and, where it can, proves.

1.2 The object

The set is the vertex set of the n -cube, the reference order is the binary order induced by reading a vertex as a number, and the symmetries available are the hyperoctahedral group B_n , which permutes coordinates and complements subsets of them. An ordering is a bijection from positions to vertices; a construction partitions it into blocks; and the group of interest is the subgroup of B_n that sends every block onto a block.

What the paper contributes is the general apparatus: the orbit accounting, the parity obstruction and the two characterisation theorems of sections 3 to 5, which are stated and proved for an arbitrary dimension and an arbitrary subgroup. The three constructed orderings below are the motivation for that apparatus and its first application, not the result it is built to deliver.

Three orderings of the 6-cube carry the concrete weight of the paper. They come from a published replication package, read at a fixed tag, and they are of interest here for a structural reason rather than a historical one: each is built by a documented rule with a visible symmetry, and the three rules are different enough that they populate all three outcomes. Two of them force the count to the midpoint of its range. The third bars that midpoint entirely.

To see whether those results are about the constructions or about the dimension, the same apparatus is run over a family of block systems in dimensions three to six, with a reference ordering that has nothing to do with the historical ones.

1.3 What this paper does not claim

The three orderings named above have a long history and a long literature, and this paper is not a contribution to it. Nothing here is claimed about who built those sequences, when, with what knowledge or with what intent. What is measured is a property of a received list of sixty four objects, and the measurement is silent about its origin.

The caution is not ours to invent, and it is worth quoting from the literature itself. Discussing an alternative arrangement, Moore (2005) writes:

"Shao Yong's 'Fu Xi' order of the hexagrams, being a Song dynasty production, is too late to concern us here; besides, it was never intended to order the actual text of the *Zhouyi*, and the notion of ordering the text in this fashion (and in particular that such a textual-ordering has priority to the King Wen sequence) is a 20th century invention [Moore 2005]."

Moore, *Structural Elements in the King Wen Sequence of Hexagrams*, Oracle Paper No. 1, February 2005, printed page 6.

The warning there is about attributing a modern way of ordering to an ancient source. Our version of the same caution runs in the other direction and is stricter: we attribute nothing at all. A symmetry that a construction respects is a property of the construction as it has reached us, and to say that a count is forced by that symmetry is a statement about the arithmetic of the list, not about anybody's design.

1.4 What the paper contains, section by section

Section 2 fixes the cube, the group, the four bit conventions and the denominator, defines discordant pairs, and separates that term from three neighbours that already own words we use: the sinological *inversion*, the two established senses of *balance*, and the Coxeter length in type B.

Section 3 builds the accounting. Lemma 0 relates the discordance of a pair to that of its image under a group element by two computable bits. Lemma 1 shows that if every orbit is forced then the total is exactly half the number of pairs, with the corollary that a group can only ever force that midpoint and no other value. Lemmas 2 and 3 give sufficient conditions for an orbit to be forced, by a uniform witness and by a matching. The germ of Lemma 1 and the mechanism of Lemma 2 belong to the homomesy literature and are cited where they are used.

Section 4 proves the parity obstruction: if the group contains the translations by a subspace of dimension at least two, every orbit has even cardinality, so the parity of the count is fixed by the structure whatever the free choices do. The dimension hypothesis is necessary, by an exhibited witness, and a normality hypothesis that an earlier formulation carried is proved redundant.

Section 5 removes the apparatus again. Being forced is equivalent to contributing exactly half, and when the group contains all translations it is equivalent to every difference class splitting exactly in two. It also separates, with proof, what the group decides from what the ordering decides.

Section 6 applies all of it to the three orderings, including a mechanism previously claimed for one of them which the measurements refute, the anatomy of the residue left by the third, and the point at which we stop and say so.

Section 7 varies the dimension over a parametrised family of block systems, and reports two refutations that came out of trying to raise measured shapes into theorems, one of them of an error of our own.

Section 8 lists what is open, each item with what was measured and with the specific task that would close it.

Section 9 describes what a reader can verify, and how the order in which things were fixed and measured is itself checkable.

1.5 A note on how this paper was made

Every measurement in this paper was preceded, in a separate and earlier commit of a public repository, by a written statement of what was going to be measured and, where predictions were made, of what would count as a refutation of them. Several of those statements were then refuted by the measurements, including one prediction whose refutation criterion turned out to be unfailable and is reported as a defect rather than cashed as a success. Section 9 explains the arrangement. It is mentioned here because it changes how the results below should be read: where this paper says that something was declared in advance, that is a checkable claim and not a rhetorical one.

2. Preliminaries

2.1 The cube and its vertices

A hexagram is a vector of six binary coordinates. We number the coordinates 1 to 6, calling coordinate 1 the bottom line and coordinate 6 the top line, and we write yang for the value one and yin for the value zero. There are $2^6 = 64$ of them, and they are the vertices of the 6-cube.

Everything below is stated for a general dimension n , with $N = 2^n$ vertices, and specialised to $n = 6$ only where the object under study requires it.

2.2 The group

The maps we consider are those that permute the n coordinates and then complement a subset of them. There are $n! \cdot 2^n$ of them, which for $n = 6$ is 46080. This is the **hyperoctahedral group** B_n , the symmetry group of the n -cube.

Correspondence, in one line: the vertices of the n -cube are the binary vectors, permuting lines is permuting coordinate axes, and complementing a subset of lines is reflecting in the corresponding coordinate hyperplanes.

Every element of B_n is written uniquely as a coordinate permutation followed by a complementation mask, so this parametrisation counts no element twice, and the action preserves adjacency of the cube, that is, the relation of differing in a single coordinate.

2.3 Orderings, reference orders, and the four conventions

An **ordering** is a bijection from positions $0, \dots, N-1$ to the vertices. To compare an ordering with a numerical order we must fix how a vertex is read as a number, and there are two independent binary choices: whether yang counts as one or as zero, and whether the bottom line is the most or the least significant bit. That gives **four conventions**, fixed in advance, and every count in this work is reported under all four rather than under the most favourable one.

The second of those two choices is not ours. Schöter (1998) parametrises the orders of the hexagrams by, among other things, "whether the lower or upper line is the least significant bit", and names the two readings Rising Yang and Sinking Yang. We use that parameter and do not present it as new.

The **reference order** throughout is the binary order induced by the convention in use, and the denominator for every rate is $C(N, 2)$, the number of unordered pairs of distinct positions. For $n = 6$ that is 2016, and its half is 1008.

2.4 Discordant pairs

Given an ordering and the reference order, a pair of distinct positions is **discordant** when the two orders disagree on it: the vertex placed earlier by the ordering is the larger one under the reference order. The statistic of this paper is the number of discordant pairs.

First appearance, and what the term already means. Discordant pairs are counted **between the ordering and a fixed reference order**, in the standard sense of rank correlation, where a pair is concordant when two rankings agree on it and discordant when they disagree. It is the same object as the combinatorial inversion number: Björner and Brenti (2005) define, in their equation (1.25), $\text{inv}(x) = \text{card}\{(i, j) : i < j, x(i) > x(j)\}$, which is exactly the count above with the reference order as the second ranking. It is also the quantity in the numerator of Kendall's tau, whose denominator is the same $C(N, 2)$, so the value $C(N, 2)/2$ that we will call the **tie** is precisely the point where that tau is zero. We do not use tau anywhere; we count the pairs.

2.5 Three names that are taken, and are not this

Three words that appear near this work already have owners, and we separate them here, once, rather than in a closing note.

Inversion. In the sinological literature on the Yijing, *inversion* names the 180 degree turn of a hexagram, *fandui*, which is a symmetry of the figure and not a discordant pair. Gritter (2015) states the two pairing principles as inversion and opposition and gives the Chinese terms; Cook (2006) devotes a chapter to inversion and obversion; Drasny (n.d.) writes of inverses of hexagrams, *zonggua*; and Moore (2005) speaks of invertible pairs. In this paper *inversion* is never used for the statistic.

Balance. The word is occupied twice over in the neighbourhood. Radisic (2026) uses it for Hamming weight, writing that weight preservation "may be viewed as preservation of yin-yang balance" and adding at once that formally it is Hamming weight. In the Gray code literature Mütze (2023) defines a *balanced* Gray code by the condition that the transition counts satisfy $|c_i - 2^{n-1}| < 2$, that is, a condition on how often each coordinate flips. Neither is the quantity studied here, and we therefore call $C(N, 2)/2$ the **tie** and not the balance point.

Length in type B. This is the subtlest of the three, because we name B_n and we count inversions in the same breath. In Coxeter theory the length function of a group of type B is a count of certain inversions of signed permutations, as Björner and Brenti (2005) describe in their sections 8.1 and 8.2. That is not our count. Here B_n is only the group acting on the vertices, and the discordant pairs are those of the ordering against the binary reference order, not the Coxeter length of any element of B_n .

2.6 Block systems and what it means to respect one

A construction partitions the ordering into contiguous **blocks**: eight octets for one of the sequences below, eight palaces for another, thirty two adjacent pairs for the third. We say that a map **respects** the construction when it sends every block onto a block, as a set. This is the notion fixed in advance, and where a stronger one is used, namely that the map also preserves the index within the block, it is said explicitly.

For each construction, the set of elements of B_n that respect it is a subgroup, being the stabiliser of a structure, and we compute it exactly rather than sampling it.

3. The orbit accounting

Let G be a group acting on the vertices, σ an ordering, and let π_g denote the permutation of positions induced by g , that is, σ composed with g and with the inverse of σ . The group acts on the $C(N, 2)$ pairs of positions through π_g . Everything in this section is a theorem about that action.

3.1 The state relation

For a pair p of positions and an element g , write A for the bit that is one when π_g reverses the order of the two positions, and B for the bit that is one when g reverses the binary order of the two vertices. Both are computable without knowing whether p is discordant.

Lemma 0. $\text{state}(g p) = \text{state}(p) \text{ XOR } A \text{ XOR } B$.

Proof. Write x and y for the vertices at the two positions of p , taken in position order, so that $\text{state}(p)$ is the truth value of $v(x) > v(y)$. Since σ of π_g of a position is g of the vertex at that position, the image pair carries the vertices $g(x)$ and $g(y)$. If π_g preserves the position order, state of the image is the truth value of $v(g(x)) > v(g(y))$; if it reverses it, state of the image is its negation. In both cases state of the image equals that truth value XOR A . And by

the definition of B , that truth value equals $\text{state}(p) \text{ XOR } B$. Substituting gives the formula. No value can tie, because x and y are distinct and v is injective. QED

Lemma 0b, and why chains give nothing. $\epsilon(gh, p) = \epsilon(g, h p) \text{ XOR } \epsilon(h, p)$, where $\epsilon(g, p) := A \text{ XOR } B$.

Proof. Apply Lemma 0 to h , to g , and to gh , and equate. QED

A chain of group elements applied to a pair therefore accumulates its epsilon, and that accumulation is exactly the epsilon of the product, which is another element of the group. Chaining reaches no further than a single element. The useful generalisation of what follows is not longer chains; it is weaker hypotheses.

3.2 The tie lemma

Fix a representative in each orbit and let $\text{parity}(p)$ be the sum of the epsilon along any path from the representative to p , which is well defined because the state exists. Let c denote the number of pairs of the orbit with parity one. Call the orbit **forced** when c is half its cardinality.

Lemma 1. If every orbit is forced, the total number of discordant pairs is $C(N, 2)/2$.

Proof. By Lemma 0, $\text{state}(p) = \text{state}(\text{rep}) \text{ XOR } \text{parity}(p)$ throughout the orbit, so the orbit contributes c or its cardinality minus c according to the single free bit. If the orbit is forced the two quantities coincide and equal half the cardinality, whatever the state of the representative. Summing over orbits, and since the orbits partition the set of pairs, the total is half of $C(N, 2)$. QED

Corollary. A group can only ever force the tie. It cannot determine the count and have it come out at some other value.

Prior owner of the germ. The idea that an involution which sends the count to its complement forces the average to one half on every orbit is the founding example of the homomesy literature. Propp and Roby (2015), in section 2.1 of arXiv:1310.5201v6, take the permutations of $\{1, \dots, n\}$, let τ send a permutation to its reversal and f be the number of inversions, and observe that since τ^2 is the identity and $f(\pi) + f(\tau(\pi)) = n(n-1)/2$, the statistic f is c -mesic with $c = n(n-1)/4$. For $n = 64$ those two numbers are 2016 and 1008, which are our denominator and our tie: this is not an analogy, it is the same arithmetic.

The difference is where the involution acts. τ reverses the permutation, so it acts on positions and leaves the labels, and the identity is clean because every pair is an inversion in exactly one of the two permutations. Ours is complementation, which acts on the values and, by inducing a permutation of positions, moves both at once. That is precisely why Lemma 0 needs the bit $\epsilon = A \text{ XOR } B$, and why Lemma 1 is their argument restricted to the class of pairs on which the position permutation preserves the order.

The same germ has a third home, and naming it costs one sentence. Elder et al. (2024), in their Theorem 5.7, list the inversion number among the statistics that are homomesic under **both** the reverse and the complement of a permutation, with the same average $n(n-1)/4$. Their reverse sends the entry at position i to position $n + 1 - i$, so it acts on positions; their complement sends the value v to $n + 1 - v$, so it acts on values; and neither induces the other. Ours is neither: complementation acts on the vertex labels and, through the ordering, moves the positions at the same time. The delimitation just made covers this third site unchanged, and it is the reason Lemma 0 carries a bit and their arguments do not.

In the vocabulary of that literature, the corollary above says that the indicator of discordance is $1/2$ -mesic on the forced orbits, in the sense of Definition 1 of Propp and Roby (2015). We use the term and do not present it as ours.

3.3 Forcing by a uniform witness

Lemma 2. If some g in G has $\text{epsilon}(g, p) = 1$ for every p in an orbit, that orbit is forced.

Proof. By Lemma 0 the state alternates along every cycle of g inside the orbit. A cycle of odd length would force a state to equal its own negation, so every cycle has even length and contains as many pairs of state one as of state zero. Summing over cycles, the orbit contributes half its cardinality, and it does so for either value of the free bit. QED

Such a g is called a **witness**. A witness is a finite certificate that can be checked on its own: run through the orbit and verify that epsilon is one.

Prior owner of the mechanism. The mechanism of this lemma is stated in one line in Roby's (n.d.) survey of homomesy, in his Example 4 on inversions under the ninety degree rotation of permutation matrices, whose average is again $n(n-1)/4$: "the proof of homomesy is easy: Q takes inversions to non-inversions, and vice-versa." That is the witness argument, said there for another object.

The same survey also contains the direction from a subgroup to the group: if a triple exhibits homomesy for a subgroup, it does for the whole group, because merging orbits with equal averages gives a larger orbit with that average. That is the direction we use whenever a larger group is compared with a smaller one.

3.4 Forcing by a matching

Lemma 3. Let S be a set of pairs of positions. If S admits a perfect matching into couples $\{p, q\}$ such that for each couple there is some g in G with $g p = q$ and $\text{epsilon}(g, p) = 1$, then the contribution of S is exactly half its cardinality.

Proof. By Lemma 0, $\text{epsilon}(g, p) = 1$ gives $\text{state}(q) = \text{state}(p) \text{ XOR } 1$, so exactly one of the two pairs of a couple is discordant, without knowing which. Since the couples partition S , the contribution is the number of couples. QED

Lemma 2 is the special case in which one and the same g serves for every couple: its cycles inside the orbit have even length and split into consecutive couples. What Lemma 3 adds is that the witness may change from couple to couple, and that the hypothesis is required only on the chosen half rather than on the whole orbit.

The lemma is not free. If a set does not contribute exactly half, its two parity classes have different cardinalities and no perfect matching can exist, because every edge of the relation graph joins opposite parities. The check has been carried out where it matters, in both directions. Where the lemma applies, it settles a class of 96 pairs that Lemma 2 cannot reach, with a contribution of exactly 48. Where it does not apply, it does not pretend to: of the 19 free orbits of the third historical ordering, 19 admit no perfect matching, which is what has to happen if the lemma is not to force what is not forced.

3.5 What this accounting is for

The four lemmas turn a question about one ordering into a question about the orbits of a group action on pairs. Each orbit is either forced, and then its contribution is fixed by the structure, or free, and then it contributes one of two values. The total therefore lies in an interval, and the set of totals compatible with the structure is the minimum plus the subset sums of the gaps.

This is where the present work parts company with the literature just cited. In homomesy, a case in which the average is not constant on orbits is a case that falls outside the phenomenon being collected: the phenomenon is the constancy. Here it is that case that is measured. The free orbits are not failures of a phenomenon; they are the quantity of freedom that the construction leaves, and the interval they define is the object of the sections that follow.

The other large phenomenon of this area, cyclic sieving (Reiner et al., 2004), is a different question again: it counts the fixed points of a cyclic action through a generating function at roots of unity, not the average of a statistic over orbits. It is named here to place the work, and is not used.

4. The parity obstruction

The accounting of section 3 leaves each orbit either forced or free, and the free ones make the total an interval rather than a number. This section proves a constraint that cuts across that freedom: under two hypotheses on the group, the parity of the total is fixed by the structure, whatever the free bits do. When that fixed parity differs from the parity of the tie, the tie becomes impossible.

4.1 The lemma that makes the whole section independent of the ordering

Lemma. The action on pairs of positions is conjugate, by the ordering, to the action on pairs of vertices. Hence the cardinalities of the orbits do not depend on the ordering at all: they depend only on the group.

Proof. The ordering is a bijection from positions to vertices, and it carries one action to the other by conjugation. Conjugate actions have the same orbit cardinalities. QED

The exact scope of this, because it is easy to overstate. What is independent of the ordering is the list of orbit cardinalities, and therefore also, as section 5 shows, the parity of the differences between compatible totals. What is **not** independent of the ordering are the values of c , and with them the absolute parity of the count and which of the three outcomes the ordering falls into. The group fixes the shape of the accounting; the ordering fills it.

4.2 The theorem, with two hypotheses

Theorem. Let G be a subgroup of B_n acting on the vertices, and suppose G contains the translations by a subspace V with $\dim V \geq 2$. Then every orbit of G on unordered pairs has even cardinality.

Proof. Write T for the group of translations by V , of order 2^k with $k = \dim V$. First, T is normal in G : conjugating the translation by v gives the translation by $P(v)$, where P is the linear part of the conjugating element, so the set of translation vectors of G is closed under the linear parts. If V itself is not invariant, replace it by W , the span of all images of V under the linear parts of G ; the translations by W lie in G , since conjugates of translations are translations and products of translations are translations, W is invariant by construction, and $\dim W \geq \dim V$. So we may assume T normal.

Second, T acts freely on the vertices, since a nonzero translation fixes nothing. Now take a pair $p = \{x, y\}$. A nonzero translation fixing the set $\{x, y\}$ must swap its two elements, so the stabiliser of p inside T contains the identity and at most the translation by $x + y$, and its order is one or two. The T orbit of p therefore has cardinality 2^k or 2^{k-1} , and both are even because $k \geq 2$. Finally, since T is normal, G permutes the T orbits, so all T orbits inside a single G orbit have the same cardinality. A G orbit is thus a disjoint union of T orbits of equal even size, and its cardinality is even. QED

Corollary. Under the same hypothesis, the two options of each orbit have the same parity, since they differ by the cardinality minus twice c , which is even. Hence the total is congruent modulo two to the sum of the c over the orbits, **whatever the free bits are**. The parity of the count is fixed by the structure.

4.3 The second hypothesis is necessary, and here is the witness

The condition $\dim V \geq 2$ cannot be dropped. Take G to be the two element group consisting of the identity and the translation by a nonzero vector v . It satisfies everything else: its elements are translations by a subspace contained in G , they act freely, and T is normal in G because G is abelian. And the conclusion fails at once: for any vertex x , the pair $\{x, x + v\}$ is fixed by the translation, so its orbit has cardinality one, which is odd.

The witness is exhibited for three dimensions, and it is the same witness in all three: the pair whose two vertices differ exactly by v . It produces 2 orbits of odd cardinality in the first dimension checked, 4 in the second and 8 in the third, so the failure is not an artefact of one small case.

4.4 The normality hypothesis is redundant, and that is a theorem

An earlier formulation of the theorem asked for three hypotheses: that T be the translations by a subspace contained in G , that T be normal in G , and that the dimension be at least two. The middle one is not needed, and the proof above already contains the reason: given the translations by V inside G , the group also contains the translations by W , the span of the images of V under the linear parts of G . That W is invariant by construction, so its translations are normal, and its dimension is at least that of V . The theorem then applies with W in place of V .

This is checked on an instance where T is genuinely not normal: a group of order 16 built from a V of dimension 2 and a coordinate permutation that moves it. The subgroup T fails to be normal, the closure argument produces an invariant subspace W of dimension 3, and of the 5 orbits of pairs, 0 have odd cardinality.

A declared family of 50 further cases, over two dimensions and every subspace of dimension at least two together with every coordinate permutation, produces no orbit of odd cardinality either.

4.5 When the obstruction bites

The tie is $C(N, 2)/2$. If the structural parity of the count differs from the parity of that number, then no choice of the free bits can reach it, and the tie is not merely unattained but **impossible**. This is the mechanism that distinguishes one of the three historical orderings of section 6 from the other two.

It is worth saying what this mechanism does **not** explain, since we tried it and it failed. The parity argument is not what excludes the tie in every case where the tie is excluded: there are orderings whose compatible totals share the parity of the tie, whose interval contains the tie, and which still do not reach it, because the set of compatible totals is coarse. Parity is one obstruction among others, and we claim only what it proves.

5. The characterisation

Sections 3 and 4 build an apparatus: orbits, parities, forced and free. This section removes it again. Two theorems turn the question of whether an ordering is forced into a plain count, and a third tells exactly how much of the outcome belongs to the group and how much to the ordering.

5.1 Forced means contributing half

Theorem 1. An orbit is forced if and only if the number of discordant pairs it contains is exactly half its cardinality.

Proof. The contribution of an orbit is c or its cardinality minus c , according to the single free bit. If it equals half the cardinality then c or the cardinality minus c equals half, and either

equation gives c equal to half, which is the definition of forced. Conversely, if c is half then the two options coincide at half. QED

This removes the apparatus from the question of **whether**. To decide if an ordering is forced one need not propagate parities or compute any epsilon: one counts. The accounting is still needed to say **what interval** remains when the ordering is not forced, but not to answer whether it is.

5.2 Forced means every difference class splits in two

Theorem 2. Let G contain all the translations. Then the orbits of G on pairs of vertices are the difference classes, that is, the sets of pairs whose difference runs over an orbit of the linear parts. Consequently an ordering is forced **if and only if, in every difference class, exactly half of its pairs are discordant**.

Proof. The translations act transitively on the pairs with a fixed difference, since the pair at x goes to the pair at x plus v . The linear parts send a difference d to its image. Hence the orbit of a pair is the set of pairs whose difference lies in the orbit of d under the linear parts. Applying Theorem 1 orbit by orbit gives the statement. QED

The argument uses neither a particular dimension nor a particular partition, so it holds for every group containing the full translation group, which covers all the $B(n, k)$ block systems of section 7.

It does not cover the other two constructions of section 6. The group that respects the palaces is the translations by a three dimensional subspace, and the group that respects the pairs of the third ordering is the centraliser of the half turn; neither contains all the translations of the cube, so Theorem 2 does not apply to them and their orbits are not the difference classes. Those two cases are settled by the orbit accounting of section 3, which needs no such hypothesis.

5.3 Forcing by a matching, and what is proved versus counted

Lemma 3 of section 3 gives a sufficient condition for an orbit to be forced: a perfect matching of its pairs into couples, each couple carrying a witness. The converse question, whether every forced orbit admits such a matching, has two directions with different status, and we keep them apart.

One direction is a theorem. If an orbit is not forced, no perfect matching can exist. Every edge of the relation graph joins pairs of opposite parity, so the graph is bipartite between the two parity classes; if the orbit is not forced those classes have different cardinalities, and no perfect matching exists. This direction needs no search at all.

The other direction is enumerative, and stays enumerative. If an orbit is forced, the two classes have equal cardinality, but that alone does not give a matching: Hall's condition is needed and could fail. It does not fail in any case we examined, and the search was declared in advance, bounded, and reported with its outcome branch fixed beforehand. Its size is the whole content of the word *enumerative*, so it is printed here rather than left to the results file: 1272 forced orderings traversed across three block systems and two dimensions, of which 200 were found by search in the larger one, giving 6960 forced orbits in which the matching was verified, and 0 failures of Hall's condition. That is a count, not a proof. **We do not promote it to a theorem.** A failure could occur in a higher dimension, in another partition, or in a case the search did not reach.

5.4 The forced class is closed under reversing the ordering

Theorem. If an ordering is forced, so is the ordering read backwards.

Proof. Reversing sends position i to $N - 1 - i$. This conjugates the induced permutation of every group element, so orbits of pairs go to orbits of pairs of the same cardinality. And within each pair, the vertex that was earlier is now later, so its discordance is negated. The contribution of an orbit becomes its cardinality minus the contribution, and being half is preserved. By Theorem 1, forced goes to forced. QED

The corresponding statement for relabelling the vertices by a group element is **false**, and measured to be false. That is not an anomaly: the group acts on the vertices while the value function stays fixed, so relabelling produces a genuinely different ordering, and there is no reason for the forced class to be invariant under it.

5.5 What the group decides and what the ordering decides

Theorem. The orbit cardinalities, and therefore the parity of the differences between compatible totals, are functions of the group alone. The values of c , the absolute parity of the count, and which of the three outcomes occurs, are not.

Proof of the first half. The cardinalities are independent of the ordering by the lemma of section 4.1. For the differences: the set of compatible totals is the minimum plus all subset sums of the gaps of the free orbits, and the gap of an orbit is the absolute value of its cardinality minus twice its c . Modulo two that gap is congruent to the cardinality, since the term in c is even. So the parity of every gap, and hence of every difference between two compatible totals, is fixed by the cardinality, which does not depend on the ordering. QED

Proof of the second half, by witness. Take the smallest block system in the smallest dimension we enumerate, where the space of orderings is small enough to be traversed in full: all 40320 of them, under one and the same group of order 16. Of those, 472 are forced and 39848 are not. One group, two outcomes; therefore the outcome is not a function of the group. QED

Those two counts partition the space exactly, which the results file checks. They are the witness of the theorem above and are not read for anything else.

5.6 What the characterisation is, and what it is not

There **is** a characterisation, and it is Theorem 2: forced if and only if every difference class splits exactly in two. It is proved, it is general, and it turns the question into a count with no apparatus in the way.

What there is **not** is a simpler invariant behind it. A closed list of candidate invariants was declared in advance and tested against the full partition of the enumerated space. Three of them fail to separate the forced orderings from the rest. One separates, but only because it encodes the definition. The one that separates informatively is the matching of section 5.3, whose two directions have the mixed status described there. That is the state of the question, and the negative half of it is a result too.

6. The three historical orderings

The three orderings are taken from a published replication package, read at a fixed tag, and only the sequences themselves are taken from it: no figure of that package enters any computation here. Two of the three are constructions with documented rules and are rebuilt from those rules in our own code, which halts if the rebuild fails to reproduce the extracted sequence. The third is a received datum and cannot be derived from anything.

That third sequence has since been checked against an independent artefact, the full binary table printed in the appendix of another paper, under that paper's own bit convention. All 64 positions agree. What is corroborated is the transcription, that is, that the list carries no error peculiar to a single source; nothing historical is corroborated by it.

Every count below is a number of discordant pairs out of 2016, the number of unordered pairs of the 64 positions, and the tie is 1008.

6.1 Mawangdui: forced, and the mechanism that was claimed for it is false

The count is 1008 in all four conventions, which is the tie exactly.

The block system is the fibres of the upper trigram, that is, the cosets of the subspace spanned by the three lower lines. The group of elements of B_6 respecting it is exactly the set of maps whose linear part does not mix the two trigrams, with any mask: 36 such permutations with all 64 masks, of order 2304. Its orbits of pairs are exactly the classes given by the pair of difference weights, one for each trigram, and there are 15 of them.

Every one of those classes is forced, and the demonstration is in two parts. 9 of them, covering 1568 pairs and contributing 784, are forced by a witness that is argued rather than exhibited: complement the lower trigram, and the value order of the pair reverses while the block, and therefore the position order, does not. The remaining classes are forced by witnesses that are exhibited and verified, and one of them has no uniform witness at all in the whole group and needed the matching lemma, with a certificate that is deposited pair by pair and can be checked line by line. All 15 orbits are forced and 0 are free, so by Lemma 1 the forced contribution is the whole count, 1008, and the width of the interval is zero.

The mechanism that was previously claimed for this count does not hold, and saying so is part of the result. It was claimed in the preregistration of this repository, in the section that listed a prior result to be verified and reported as a retrodiction, and that is where the reader can find it stated before any measurement was taken. The claim was that closure under complementation forces the count. Closure does hold, and it is concrete: the complement sends each octet onto another octet. But the group generated by complementation alone forces only 484 of the 2016 pairs, leaves 540 orbits free, and admits 1049 different compatible totals: an interval, not a value. The control settles it exhaustively rather than by sample. The family order of the octets can be permuted in every one of the 40320 possible ways, all such rearrangements keep the closure intact, and the tie comes out in 3836 of them, a rate of 0.09514. Closure is true of the construction and is not what forces the count. What forces it is the full group of order 2304 together with the received family order.

6.2 Jing Fang: forced, and the demonstration is complete

The count is again 1008 in all four conventions.

Each palace is a translate of one and the same set of masks, read off the construction rules, and the palace heads are exactly the diagonal subspace. None of the 28 differences of that mask set lies in the diagonal, which is why the palaces partition the vertices.

The group respecting the palaces is exactly the translations by the diagonal, of order 8, and this is proved rather than enumerated: any respecting map must preserve the multiset of differences of the mask set, whose multiplicities line by line are 12 15 16 12 15 0, so one line never appears in any difference and the remaining multiplicities are distinct enough to pin the permutation down to the identity, after which the mask is forced into the diagonal.

The 280 orbits are classified by hand into three families, and every one of them is forced: those inside a single palace by an argued witness, the complement, which here is the translation by the all ones vector of the diagonal; and the rest by exhibited witnesses, deposited orbit by orbit. There are 0 free orbits and exactly 1 compatible total, so by Lemma 1 that total is the tie. **For this ordering the demonstration is complete: there is no enumerative residue.**

6.3 King Wen: the tie is impossible

The count is 1013 under the two conventions that read yang as one, and 1003 under the two that read yang as zero. The two are symmetric about the tie, since exchanging the roles of yang and yin turns every discordant pair into a concordant one, and 1013 plus 1003 is twice 1008. Neither is the tie. Everything below explains why neither could have been.

The **half turn** is the element of B_6 whose coordinate permutation is $(1\ 6)(2\ 5)(3\ 4)$, reversing the order of the six lines, and whose complementation mask is zero. It is an involution, it fixes the vertex of all yin, and a hexagram it fixes is a palindrome.

The construction pairs the sequence into 32 adjacent blocks of two: the 28 orbits of size two of the half turn, together with the 8 hexagrams that the half turn leaves fixed, paired among themselves by complementation.

Prior owner of this characterisation. It is the complete equivariance of Radisic (2026), Theorem 3.3 of arXiv:2601.07175v3, which states that every King Wen pair is either the complement or the reversal of its partner and splits the 32 pairs into palindromes paired by complement, anti-symmetric ones where reversal and complement coincide, and generic ones paired by reversal. That statement has prior date and a formal verification in Lean 4, which we do not have. We reached it independently and before opening our review of the literature, which the commit history dates, and independent does not mean first.

From that block system, the group respecting it is exactly the centraliser of the half turn inside B_6 , of order 384, and this is proved rather than counted: the block differences are precisely the nonzero vectors of the subspace of size 8 fixed by the half turn, a respecting map must preserve that subspace, its only 3 elements of weight two are the ones that pair the mirror lines, so the linear part commutes with the half turn and there are 48 such permutations, and the mask is then forced to be fixed by the half turn, which leaves 8 masks. The predicted order 48 times 8 is 384, and it agrees with the enumeration.

Every orbit of that group on pairs has even cardinality, by the theorem of section 4 applied to the translations by the fixed subspace. Hence the parity of the count is fixed by the structure. That parity is odd, and the tie is even. **The tie is therefore impossible for this ordering:** not merely unattained, but outside the set of totals compatible with the structure.

Here is the whole account in figures. Of the 36 orbits of that group on pairs, 17 are forced and contribute 374, and 19 are free. The structure therefore leaves the interval $[957, 1059]$, and inside it exactly 52 totals are compatible, of the 103 integers the interval contains. The observed 1013 is one of the 52. The tie, 1008, lies well inside the interval and is **not** one of them: the nearest compatible total to it is 1007, one below. The tie is not missed by a small amount, it is absent from the list.

6.4 One structure narrows it, and it is the construction's own rule

A closed list of five candidate structures was declared before measuring anything: the nuclear hexagram operation with its exact definition, the two part division of the received sequence, whose literature has a located owner (Hacker & Moore, 2003), and three maps defined through the positions. Nothing outside that list was tried.

Two of the five contribute no bijection at all, so they cannot narrow anything, and that was known before running. Two of the remaining three narrow a great deal but generate groups beyond a bound that was declared in advance, and a narrowing obtained with a large unstructured group was declared in advance not to count. It does not count.

The one that counts is the pairing involution itself, the map that sends each hexagram to its partner in the construction. **It is not affine**, which is why it was not in the group at all, even though that group is precisely its centraliser. Adding it takes the group from 384 to 768, cuts the free orbits from 19 to 17, and narrows the interval from $[957, 1059]$ to $[961, 1055]$, that is, from a width of 102 to one of 94. It does not force, the tie is still not reachable, which the table

records as 0, and the observed count does not move: changing the group changes neither the sequence nor its discordance.

6.5 The anatomy of the residue, as a table

The count differs from the tie by 5, that being 1013 minus 1008, and the natural question is where those 5 sit. Each free orbit contributes its half plus a deviation, and the deviations sum to exactly 5. Table 1 gives the whole decomposition over the 19 free orbits, one row each, with cardinality, half, contribution and deviation. Two things can be read off Table 1 directly, and nothing else is read off it here:

Table 1. *The residue of the third ordering, decomposed over its free orbits.*

orbit	cardinality	half	contribution	deviation
o0	192	96	90	-6
o1	96	48	50	2
o2	96	48	52	4
o3	96	48	46	-2
o4	96	48	46	-2
o5	96	48	46	-2
o6	96	48	46	-2
o7	96	48	46	-2
o8	96	48	54	6
o9	96	48	52	4
o10	48	24	28	4
o11	48	24	22	-2
o12	48	24	22	-2
o13	24	12	16	4
o14	12	6	5	-1
o15	12	6	8	2
o16	12	6	4	-2
o17	4	2	3	1
o18	4	2	3	1

The rows are labelled o0 to o18 in the order the accounting produces them; the label is a row name and not a measurement.

first, **every free orbit deviates**; there is no orbit carrying the difference and none carrying none of it, and the deviations very nearly cancel;

second, the deviations of odd size are exactly those of the orbits that contain the construction's own pairs, and the even ones are all the rest.

The crossing of the deviating orbits against the two structural features of the declared list is **not** printed here. It stays in results/residuo5.tsv, under the keys that name each orbit against each feature, and the reason for leaving it there rather than promoting it is the guard below: it is a set of counts with their provenance and nothing is read from it. There is no declared null, no family of comparisons fixed in advance and no discipline of multiplicity behind it, and in this work that means it is not read as evidence of anything.

6.6 Where this stops

None of the five declared structures explains the residue. After the one that narrows, the residue is unchanged, the tie is still impossible, and the remaining free orbits still spread the deviation across almost all of them.

The residue is therefore declared informative relative to this list. Informative in the structural sense, as the part of the count not determined by the declared structures; no inferential claim is attached to the word. The decomposition exists, it is complete, and it has no culprit.

This is where the declared structures run out. What remains of the difference, once the symmetry that the construction respects has been exhausted and the four additional declared structures have been tried, is not a residue that any of them absorbs. It is the part of the count that no declared structure determines, and that is a statement about the list of structures that were tried and about nothing else.

That sentence is about the limits of this method, not about the people who ordered the sequence. Nothing here is claimed about design, about intention or about ancient knowledge, and the declaration is relative to the declared list: another structure, not tried here, could absorb what this one leaves.

7. The landscape $B(n,k)$

The three orderings of section 6 live in one dimension and carry historical weight. This section removes both features: it varies the dimension and replaces the received orderings by a family with a parameter, to see what the apparatus does when nothing is inherited.

7.1 One block system with a parameter, and what it does not cover

For a dimension n and a level k between 1 and n minus 1, let $B(n, k)$ be the partition of the vertices into the cosets of the subspace spanned by the k low coordinates. Two vertices share a block when they differ only in low coordinates. Every level is reported; the level was not chosen after seeing results.

This single family covers two things at once. At the middle level it is the analogue of the first historical ordering, whose blocks are the fibres of the upper trigram. Running over all levels it is the tower of recursion levels of the reflected Gray code. For odd n there is no middle level, and rather than force the analogy we say that the analogue is simply not defined there.

What the family does not cover, checked and not assumed. The palaces of the second historical ordering are the cosets of a set that is **not a subspace**, and the pairs of the third are the orbits of the half turn, whose differences are not a single vector but several. Neither is a $B(n, k)$, and putting them in the table would be forcing the analogy. They keep their own treatment in section 6 and have no row here.

7.2 The table

Table 2 gives dimensions three to six, every level, with the same columns throughout. The order of the group respecting $B(n, k)$ is k factorial times n minus k factorial times 2 to the n , which was stated in the general preregistration before any of this was run and is verified in all 14 rows. The order of the group depends only on the partition and not on the ordering, which the same table confirms across thousands of orderings.

Table 2. *The landscape $B(n, k)$ in dimensions three to six.*

n	k	group order	Gray	canonical	forced	bounded	barred	cases
3	1	16	bounded	barred	0	24	24	48, enumerated
3	2	16	bounded	barred	0	24	24	48, enumerated
4	1	96	bounded	barred	0	827	1173	2000, sampled
4	2	64	bounded	bounded	36	228	312	576, enumerated
4	3	96	bounded	barred	0	828	1172	2000, sampled
5	1	768	bounded	barred	0	781	1219	2000, sampled
5	2	384	bounded	barred	10	794	1196	2000, sampled
5	3	384	bounded	barred	9	753	1238	2000, sampled
5	4	768	bounded	barred	0	771	1229	2000, sampled
6	1	7680	bounded	barred	.	.	.	not sampled
6	2	3072	bounded	barred	.	.	.	not sampled
6	3	2304	bounded	bounded	.	.	.	not sampled

6	4	3072	bounded	barred	.	.	.	not sampled
6	5	7680	bounded	barred	.	.	.	not sampled

The last three columns of Table 2 are the distribution of the three outcomes over the parametrised family, by block order and internal order; a dot means the distribution was not measured. The **forced** column is empty in every row at an extreme level and non empty only at the intermediate ones, which is the shape that section 7.3 takes up and cuts down to size.

Each row reports the outcome for two fixed orderings, the reflected Gray code and the canonical member of the parametrised family.

The **parametrised family** of $B(n, k)$ is indexed by a pair of permutations: one of the 2^{n-k} blocks, which fixes the order in which the blocks are laid out, and one of the 2^k positions inside a block, which is the same for every block. Given the pair, the ordering places at position i times 2^k plus j the vertex whose upper n minus k coordinates encode the image of i and whose lower k coordinates encode the image of j . The **canonical member** is the one obtained when both permutations are the identity, so that position i times 2^k plus j carries the vertex whose integer code is exactly i times 2^k plus j : the map from positions to vertices is the identity. That is not the binary reference order, because the reference order reads the bottom line as the most significant bit while the vertex code reads it as the least significant one, which is why the canonical member has a discordance to report at all.

Where the space of those pairs has at most 5000 members it is traversed in full, and the row says *enumerated*. Where it is larger, 2000 pairs are drawn and the row says *sampled*. Each draw takes one permutation of the blocks and one of the inner positions, each uniform over its symmetric group and independent of the other, from a Mersenne Twister seeded once per sampled row at 20260809, a seed frozen in the source before any of this was run. The draws are independent and are not deduplicated, so the sample is with replacement and a pair can in principle come up twice.

One historical row belongs here and is kept apart. The first historical ordering does fit the family, at the middle level of dimension six, since its octets are exactly those cosets. Its row is Table 3, kept out of Table 2 so that the parametrised family and the received sequence are not read as one population:

Table 3. *The one historical ordering that fits the family, reported apart.*

ordering	n	k	group order	free orbits	compatible totals	count	outcome
Mawangdui, received	6	3	2304	0	1	1008	forced

Table 3 reports 0 free orbits, exactly 1 compatible total, and that total is 1008: the forced outcome of section 6 read inside the general frame rather than beside it, and the same figure that the demonstration there produces.

7.3 A shape that we tried to raise to a theorem, and could not

Across the grid, the forced outcome appeared only at intermediate levels and never at the extremes. That was a real observation about what had been measured, and the temptation was to state it as a property of the extremes.

It is false, and the refutation is by exhaustive enumeration. In the smallest dimension the space of all orderings can be traversed in full, all 40320 of them, and at both extreme levels it contains forced orderings, with a witness exhibited: 472 of them at the lower extreme and 600 at the upper. One also turns up at an extreme level in the next dimension. Note that neither count is zero while the corresponding cell of the table above is: the forced orderings exist, and they lie outside the family the table samples.

What had been measured was not false, it was narrower than the sentence suggested: within the parametrised family there is indeed no forced ordering at the extremes, and that remains true. The property belongs to the family, not to the extremes.

The lesson about sampling is worth stating, because it is general. The forced class is small, and at the extreme levels it lies entirely outside the family that was being sampled. No amount of sampling inside that family would ever have produced a counterexample: the sample was not too small, it was drawn from the wrong set. Enumerating the whole space in the one dimension where that is possible is what settled it.

7.4 The reflected Gray code as a reference, and an error of ours

The reflected Gray code enters as a **reference ordering** and not as a relative of anything. It is built from its recursive definition and checked against the closed form. A separation is needed on first appearance: in the Gray code literature *balanced* refers to the transition counts per coordinate, and not to the tie measured here.

Its results in this table are corrected ones, and the correction is ours to report. An earlier run used a construction that added the wrong coordinate at each recursive step, so it was not the ordering that had been declared. The discrepancy surfaced only when dimension six was brought into the same table as the smaller ones and the same object produced two different outcomes. The declared construction reproduces the closed form and the earlier one does not. With the declared construction the reference ordering falls in the same outcome in every row of the table, and the apparent anomaly that an earlier draft tried to explain **did not exist**. Both reports carry a visible amendment saying so.

That episode produced a working rule that this paper follows: a reimplementing of an object already built inherits the verifications of the original, or states in writing why it does not.

7.5 What the landscape shows and what it does not

The table shows that the outcome is not a function of the group: within a row the group is fixed and the outcomes differ. It shows that the forced outcome is possible in this family only at intermediate levels, and that outside the family it is possible at the extremes too. And it shows the reference ordering behaving uniformly across four dimensions.

It does not show a classification. Which combinations of group and ordering fall into which outcome was declared, before any of this was measured, as a question on which no prediction would be made, and it remains unanswered. The 14 rows in four dimensions are a shape, not a theorem.

8. Open problems

Four things are open. Each is stated with what was measured, and with what would close it, so that the section is a list of specified tasks rather than of wishes.

8.1 Does Hall's condition ever fail on a forced orbit?

What is proved. One direction needs no search: if an orbit is not forced, no perfect matching of the relation graph can exist, because every edge joins pairs of opposite parity and the two parity classes then have different cardinalities.

What is measured. The converse, that a forced orbit always admits such a matching, requires Hall's condition, which could fail. A bounded search was declared in advance, with its space, its seed and both outcome branches fixed before running: the 472 and 600 forced orderings of two block systems in the smallest dimension, taken whole, and 200 forced orderings of a third system in the next dimension, found using Theorem 2 of section 5. That is 6960 forced orbits in which the matching was verified, across three block systems and two dimensions, and 0 failures.

What was deliberately not done. The status was not promoted. Absence of a failure in a finite search is not a theorem, and the branch that said so was written before the search ran.

What would close it. Either a proof that the relation graph of a forced orbit always satisfies Hall’s condition, or a single counterexample: a forced orbit whose two parity classes admit no perfect matching. A counterexample would also refute the equivalence between the matching criterion and being forced. That criterion is the one of section 5.3, and it is not the characterisation: one of its two directions is proved and the other is enumerative. The characterisation is Theorem 2, as section 5.6 says.

8.2 How many orderings are forced?

What is measured. In the smallest dimension the space of all 40320 orderings is enumerable, and the number of forced ones is known exactly for both block systems, under a single group in each case: 472 for one and 600 for the other. Theorem 2 of section 5 says exactly which orderings those are: the ones in which every difference class splits in two.

What is missing. A formula, or even an asymptotic. The characterisation turns membership into a countable condition but does not count the members. The condition is a system of simultaneous exact-halving constraints, one per difference class, over the symmetric group on the vertices, and nothing here says how many permutations satisfy such a system.

What would close it. A count, in closed form or asymptotically, of the permutations satisfying the difference-class conditions of Theorem 2, in terms of n and of the block system.

8.3 The residue of the third historical ordering

What is measured. The count, 1013, differs from the tie by 5, that difference is decomposed completely over the 19 free orbits, and the decomposition has no culprit: every free orbit deviates and the deviations very nearly cancel. A closed list of five candidate structures was declared before measuring; one of them narrows the interval and none explains the difference. The residue is declared **informative relative to that list**.

What would close it, structurally. A structure, outside the declared list, that the construction respects and that absorbs the difference. The declaration is relative to the list precisely so that this remains possible.

The one principled route that remains, named as future work

There is one route we can name rather than gesture at. Radisic (2026) proves that the pairing of the third ordering is the unique cost-minimising equivariant matching under his criteria. Conditioning on that matching leaves a family of orderings: all those that realise it, differing only in the order of the pairs and in the orientation within each pair. Inside that family, the received ordering is one point, and its count could be located among the counts of the rest.

That is an inferential question, and it is not asked in this paper. Locating a value inside a family and reading anything from where it falls requires a null declared before looking, a family of comparisons fixed before looking, and discipline of multiplicity over everything that could have been looked at and was not. This work builds none of the three, and therefore cannot support such a reading with these numbers or with larger ones.

The method that would govern it has a name and a citable version, given here as a pointer and not as a read artefact: the stopping-criterion work of the same author on nested reference sets, *Uninformative rungs: an order-theoretic stopping criterion for nested reference sets* (García Hurtado, 2026), deposited at doi 10.5281/zenodo.21750029. We have not read it in the course of this work and make no claim about its contents; we name it as the framework under which the question above would have to be posed.

8.4 The shape of the classification

What was asked. For which subgroups, and under which conditions on the ordering, is the count forced to the tie, bounded in an interval, or barred from the tie by parity. The general preregistration fixed that question and declared, before any measurement, that no prediction would be made about its answer.

What is measured. The 14 rows in four dimensions of section 7, the three outcomes populated, and a proved split of what belongs to the group and what to the ordering: the orbit cardinalities and the parity of the differences between compatible totals are functions of the group; the values of c , the absolute parity and the outcome are not.

What would close it. A criterion that, given the group and the ordering, returns the outcome, and that reduces on the historical cases to what sections 4 to 6 prove. Section 5 gives such a criterion for the forced outcome alone, through the difference classes; nothing here separates the other two outcomes in the same way.

8.5 One thing that is not open

For the record: the apparent anomaly of the reference ordering, which an earlier draft listed as a question, is not open. It was an error in our own construction of that ordering, it was corrected, and it is documented as an error in the amendments cited in section 7 rather than removed from the record.

9. Methods of verification

The repository is not an appendix to this paper. Its history is part of the result, because most of what is claimed here is a claim about **when** something was known: what was fixed before a measurement, what was measured afterwards, and what changed when the two disagreed. A reader who wants to check that ordering has the commit history, and the ordering is what a finished text cannot show.

9.1 The two preregistrations, and the defect in the first

The root commit of the repository contains a preregistration and nothing else: no analysis code, no data, no measured figure. It separates three things by name: a prior result that was to be verified and reported as retrodiction, a prediction made before running anything with its refutation criterion attached, and an open discrepancy on which no bet was placed. It also fixes, in advance, the bit conventions to be tried and the denominator, so that neither could be chosen after seeing results.

The second preregistration, for the general phase, was signed after a phase whose only permitted activity was proving things: two hypotheses that had been in doubt were settled by theorem before signing, which is why its prediction section is almost empty and says so. The temptation it records having avoided is worth naming: a prediction could have been manufactured by narrowing the space until the answer became unknown again, and that is choosing the question after seeing the answer.

The refutation criterion of the first preregistration was defective, and this is the place to say so. It asked, for the prediction, that a hexagram be exhibited whose image under complementation lay outside the construction. Since every ordering contains all the vertices, that image is always the same set: no sequence could ever have failed the criterion. A preregistered criterion that cannot be failed is not a criterion. It was reported as a defect in the very report that would have benefited from claiming the prediction as confirmed, the preregistration was left unamended, and the prediction was not cashed.

9.2 The effort log

Every working session opens and closes an entry in an append only log: 36 closed sessions in 176 records at the point this text was frozen. Each record carries the previous record's hash, so editing an old line breaks the chain and a verifier reports it, and the chain currently verifies with 0 problems; the tool that writes the log has no operation that rewrites or deletes. Exactly 1 entry is marked retroactive, the first, and it says that a reconstructed record is not equivalent to one taken live, which is the only honest thing to do with a log that begins one commit late.

The log also classifies every file as apparatus or analysis, separating what was written here from what was extracted, so that the proportion between building instruments and producing results is visible rather than anecdotal. At the same point, 98 files were classified, over 20519 lines: 7440 of apparatus against 13079 of analysis, of which 313 lines are extracted from elsewhere and are not counted as written here.

Dead ends are recorded as their own kind of entry, with their cost. There are 8 of them: a command written with the wrong working tree, a process left running after its replacement had been launched, an analysis that hung because a group closure exploded exactly as a declaration had warned it might, a tool of our own that resolved a key by prefix and so pointed a figure at the wrong line, the same prefix mistake made again by hand a session later while editing this very section, a criterion for binary files that let an uncompressed PDF pass as text and counted its line breaks as written work, a patch script that truncated a source file because the call that opens a file for writing empties it before it validates its own arguments, and a shell here-document that ate an escape sequence and left a control byte inside a checker, so that the checker failed on the very fix it was meant to confirm. The fourth, the fifth and the eighth were caught by the checkers written to catch exactly that; the sixth by the packaging run; the seventh by the file being under version control, which is the cheapest safety net in the list.

9.3 Rules that were born during the work, with their origin

A reimplementation inherits the verifications of the original, or states in writing why it does not. This rule exists because of a specific failure. The reference ordering of section 7 was first built with two checks tied to its definition, and when it was reimplemented for a variable dimension the reimplementation dropped both, built a different ordering, and cost a full session explaining an anomaly that did not exist. The check it dropped would have caught the error at the moment the function was written.

No figure enters a commit message before a checker has printed it. This rule exists because two commit messages in one day carried counts that had been written from memory rather than counted, in a repository whose whole point is that figures come with provenance. The counts were small and harmless; the habit was not.

Two further rules predate the work but were exercised by it: contact with the source of the data is read only and at a fixed tag, and the name of anything that leaves this work passes a term check before it is frozen. The repository name and the title of this paper each have their own record of that check, with the queries declared before running and the limits of the check declared after.

9.4 What a third party can verify today

These counts are a snapshot, and the apparatus makes going stale a failure rather than a silence. They are counted up to the last closed session, so opening a session or recording a decision does not move them; closing one does. When it does, the exporter is rerun and the figures above no longer match, and the checker that guards every figure in this paper refuses the assembly until they are brought up to date. A number that would quietly rot is instead a number that stops the build.

Clone the repository and run the programs. Every figure in this paper comes from a results file in it, every results file is produced by a program in it, and every program is deterministic: the only source of randomness is a seed that is frozen in the source and declared in the reports. The declarations that were made before each measurement are separate commits, earlier in the history than the measurements they govern, and that ordering is checkable without trusting anyone.

What a third party cannot verify by running anything is the literature review, which is a matter of what was read. That is why it carries its own doctrine, verdicts only from artefacts actually read, second hand marked as second hand with its chain of citation, and every figure with a pointer or returned; why the intersections found are recorded next to the results they touch rather than in a closing note; and why the single novelty declaration is a signed file with its scope stated as relative to a review whose limits are written down.

References

Each entry records the identity as the literature review fixed it, and whether the artefact was read here or enters as second hand. That mark is part of the reference.

1. Björner, A., & Brenti, F. (2005). *Combinatorics of Coxeter groups* (Graduate Texts in Mathematics, Vol. 231). Springer. ISBN 3-540-44238-3. Equation (1.25) and Proposition 1.5.2, p. 20; Sections 8.1 and 8.2, pp. 245, 252.
Status: read, and the cited pages verified against the PDF; the year, series, volume, publisher and ISBN were read from the front matter of the copy held here, its title page and its copyright page. That title page prints the first author as Bjorner, without the diaeresis; the accepted spelling is used above and the discrepancy is recorded here.
2. Cook, R. S. (2006). *Classical Chinese combinatorics: Derivation of the Book of Changes hexagram sequence* (STEDT Monograph Series, Vol. 5). University of California, Berkeley. ISBN 0-944613-44-6.
Status: read through its review in full, plus full text sweeps of the converted text; the book itself was not read cover to cover.
3. Drasny, J. *The solution of the King Wen sequence?* [Review of the book *Classical Chinese combinatorics*, by R. S. Cook]. Yijing Dao. <http://www.biroco.com/yijing/cook.htm>
Status: read in full; the review carries no date on the page read, so none is printed.
4. Elder, J., Lafrenière, N., McNicholas, E., Striker, J., & Welch, A. (2024). Homomesies on permutations: An analysis of maps and statistics in the FindStat database. *Mathematics of Computation*, 93(346), 921–976.
Status: read in the cited parts, verified against the rendered pages of the PDF: Definition 3.8 on p. 6, the treatment of a non homomesic statistic on p. 16, and Definitions 5.1 and 5.2 and Theorem 5.7 on pp. 21 and 22. The other fifty pages were not read.
5. García Hurtado, A. (2026). *Uninformative rungs: An order-theoretic stopping criterion for nested reference sets*. Zenodo. <https://doi.org/10.5281/zenodo.21750029>
Status: cited as a pointer to the framework under which the inferential question of Section 8.3 would have to be posed. ITS CONTENT IS NOT LOAD BEARING HERE: nothing in this paper depends on it, and it was not read in the course of this work.
6. Gritter, G. (2015). *The hidden pattern in the classical sequence of the I Ching*. Groningen: Self-published.
Status: read in full; the artefact names a place and a year and no publisher, so the entry says self-published, which is what APA asks for when a work carries none. If a copy turns up that names one, the artefact wins.
7. Hacker, E., & Moore, S. (2003). A brief note on the two-part division of the received order of the hexagrams in the Zhouyi. *Journal of Chinese Philosophy*, 30(2), 219–221.

- Status:* SECOND HAND: bibliographic identity taken from the bibliography of Moore (2005); not read. The pointer names the author and year and not a number, so that renumbering the list cannot turn it into a reference to itself, which is what it had just become.
8. Moore, S. (2005). *Structural elements in the King Wen sequence of hexagrams* (Oracle Paper No. 1).
Status: read; the quoted passage taken verbatim from the rendered PDF page, p. 6, and not from the OCR conversion.
 9. Mütze, T. (2023). Combinatorial Gray codes: An updated survey. *The Electronic Journal of Combinatorics*, 30(3), Dynamic Survey DS26.
Status: read in the cited part, Section 3.2 on p. 11, verified against the PDF.
 10. Propp, J., & Roby, T. (2015). *Homomesy in products of two chains* (Version 6) [Preprint]. arXiv. <https://arxiv.org/abs/1310.5201v6>
Status: read, and Section 2.1 on p. 4 verified against the PDF; its journal identity is second hand, since the artefact read is the arXiv version.
 11. Radisic, A. (2026). *Optimal equivariant matchings on the 6-cube: With an application to the King Wen sequence* (Version 3) [Preprint]. arXiv. <https://arxiv.org/abs/2601.07175v3>
Status: read in full, eleven pages; its Appendix A was transcribed and collated against our data, all sixty four positions agreeing. The year comes from the header of version 3, dated 25 May 2026, recorded when the artefact was read.
 12. Reiner, V., Stanton, D., & White, D. (2004). The cyclic sieving phenomenon. *Journal of Combinatorial Theory, Series A*, 108(1), 17–50. <https://doi.org/10.1016/j.jcta.2004.04.009>
Status: read in identity and definition; cited for context, since it counts fixed points and not orbit averages.
 13. Roby, T. *Dynamical algebraic combinatorics and the homomesy phenomenon*. Example 1, p. 3; Section 2.1 and Example 4, p. 4.
Status: read, and the cited pages verified against the PDF; its first page was checked again for this edition and carries title, author, abstract, key words and affiliation but neither year nor volume nor series, so neither is printed and the gap stays declared.
 14. Schöter, A. (1998). Boolean algebra and the Yi Jing. *The Oracle: The Journal of Yijing Studies*, 2(7), 19–34. ISSN 1463-6220.
Status: read; Definition 6, Sequence Parameters.