Addendum to SI: The Monogamy Constraint in Two Coordinate Systems

Topological Robustness of V + χ = 7 and Implications for the Weinberg Angle

---

A.1 Motivation

The monogamy constraint on three features in Cl(5) was originally formulated as a linear budget constraint in normalised correlation space, producing a polytope with V=5, E=9, F=6, χ=2. The combination V+χ=7 enters both the fine structure constant (as √(7/30)) and the Weinberg angle (as 7/30).

The fundamental geometric constraint for three unit vectors in $\mathbb{R}^5$ to be mutually distinguishable is not the linear budget but the Gram determinant:

$$\det G_3 = 1 - \cos^2\theta_{AB} - \cos^2\theta_{BC} - \cos^2\theta_{CA} + 2\cos\theta_{AB}\cos\theta_{BC}\cos\theta_{CA} \geq 0$$

This defines a LARGER region than the linear polytope. Computational verification (see companion script `monogamy_topology_verification.py`, Table 8) confirms: the Gram region occupies 71.3% of the unit cube in correlation-magnitude space, compared to 25.0% for the linear polytope, and the linear polytope is a strict subset (zero samples in the linear region violate the Gram constraint, while 46.3% of the unit cube satisfies the Gram constraint but not the linear one).

The question addressed here: does the change in shape affect the topological invariants that enter the physical predictions?

The answer is no. V+χ=7 is preserved. But the analysis reveals something deeper about why, and what the difference between the two formulations means physically.

A.2 The Two Coordinate Systems

A.2.1 Dot-Product Space: (x, y, z) = (a·b, b·c, c·a) ∈ [-1, 1]³

The Gram determinant constraint defines:

$$\mathcal{R} = \{(x, y, z) \in [-1, 1]^3 : 1 - x^2 - y^2 - z^2 + 2xyz \geq 0\}$$

This is a spectrahedron — the intersection of the positive-semidefinite cone with the affine constraints that fix the diagonal entries of the 3×3 matrix to unity:

$$\mathcal{R} = \left\{ (x, y, z) : \begin{pmatrix} 1 & x & z \\ x & 1 & y \\ z & y & 1 \end{pmatrix} \succeq 0 \right\}$$

Since the PSD cone is convex and affine constraints preserve convexity, R is a convex body. Its boundary is the algebraic surface $f(x,y,z) = 1 - x^2 - y^2 - z^2 + 2xyz = 0$, consisting of configurations where the three features are coplanar (rank-deficient Gram matrix).

Proof of convexity. R is the intersection of (a) the PSD cone $\{M \succcurlyeq 0\}$, which is convex, and (b) the affine subspace $\{M_{ii} = 1, M \in \mathrm{Sym}_3\}$, which is an affine plane. The intersection of convex sets is convex. (The Hessian of f is not everywhere negative-semidefinite — it has a positive eigenvalue at the boundary vertices — but this is irrelevant since convexity of R follows from the spectrahedron structure, not from concavity of f.) Numerical verification confirms: over 38,000 random pairs of points in R, with 20 interpolation points per pair, no convexity violation is found (Table 5 in companion script).

Topology of ∂R:

Since R is a compact convex body in $\mathbb{R}^3$ with nonempty interior, it is homeomorphic to a closed 3-ball. Its boundary ∂R is therefore homeomorphic to S^2, giving $\chi(\partial R) = 2$.

The CW structure of ∂R can be determined explicitly:

Vertices (0-cells). The extreme points of R on the boundary of the ambient cube $[-1,1]^3$ are the rank-1 correlation matrices. A rank-1 matrix with unit diagonal has the form vv^T where $v \in \{\pm1\}^3$. Since v and −v produce the same matrix, there are $2^3/2 = 4$ distinct rank-1 points (Table 1 in companion script):

Vertex	Dot-product coords	Physical meaning	f value
(1, 1, 1)	All positive maximal	a ∥ b ∥ c (all aligned)	0
(1, −1, −1)	AB aligned, others anti	a ≈ b, c opposite	0
(−1, 1, −1)	BC aligned, others anti	b ≈ c, a opposite	0
(−1, −1, 1)	CA aligned, others anti	c ≈ a, b opposite	0

These four points form a **regular tetrahedron** inscribed in the cube, with all pairwise distances equal to $2\sqrt{2}$ (Table 2).

Edges (1-cells). On each face of the cube (e.g. $x = +1$), the constraint reduces to $f(1,y,z) = -(y-z)^2 \leqslant 0$, with equality only when $y = z$ (Table 4). Therefore R intersects each cube face in a line segment — a 1-dimensional edge. Each edge connects two vertices that share a cube face (i.e. differ in exactly two coordinates). Since every pair of the four tetrahedron vertices differs in exactly two coordinates, all $C(4,2) = 6$ pairs are connected. The six edges are in bijection with the six faces of the cube (Table 3):

Edge	**Cube face**	**Constraint**
(1,1,1) — (−1,−1,1)	z = +1	x = y
(1,1,1) — (−1,1,−1)	y = +1	x = z
(1,1,1) — (1,−1,−1)	x = +1	y = z
(−1,−1,1) — (−1,1,−1)	x = −1	y = −z
(−1,−1,1) — (1,−1,−1)	y = −1	x = −z
(−1,1,−1) — (1,−1,−1)	z = −1	x = −y

Faces (2-cells). The edges partition ∂R into connected patches. By Euler's formula, $F = 2 - V + E = 2 - 4 + 6 = 4$. Each face is a curved triangular patch of the algebraic surface $f = 0$, bounded by three edges.

$$V = 4, \quad E = 6, \quad F = 4, \quad \chi = V - E + F = 2, \quad V + \chi = 6$$

The boundary ∂R has the combinatorial type of a **tetrahedron** — the complete graph K_4 — but with curved faces rather than flat ones.

A.2.2 Correlation-Magnitude Space: (u, v, w) = (cos²θ_{AB}, cos²θ_{BC}, cos²θ_{CA}) ∈ [0, 1]³

Each feature in Cl(5) has a $\mathbb{Z}_2$ orientation ambiguity: the grade-1 vectors a_i and $-a_i$ represent the same distinguishable entity (they have the same bivector magnitudes with all other features, since $|a_i \wedge a_j| = |(-a_i) \wedge a_j|$). The signed dot product $a_i \cdot a_j$ changes sign under this $\mathbb{Z}_2$ action; the correlation magnitude $(a_i \cdot a_j)^2 = \cos^2\theta_{ij}$ is invariant. The independent $\mathbb{Z}_2$ choice for each of N features gives a $(\mathbb{Z}_2)^N$ gauge group, under which the squared dot products are the gauge-invariant observables.

The natural gauge-invariant map from dot-product space to correlation-magnitude space is therefore the squaring map:

$$(x, y, z) \mapsto (u, v, w) = (x^2, y^2, z^2)$$

Under this map, the Gram constraint becomes:

$$\mathcal{R}' = \{(u, v, w) \in [0, 1]^3 : 1 - u - v - w + 2\sqrt{uvw} \geq 0\}$$

where the sign ambiguity in xyz is resolved by taking the positive square root: the squaring map identifies (x,y,z) with all eight sign-flips (±x, ±y, ±z), and among the pre-images of any point (u,v,w), at least one has $xyz \geqslant 0$ (since flipping any two signs preserves the product's sign, and flipping all three reverses it — so among the 8 pre-images, exactly 4 have $xyz \geqslant 0$ and 4 have $xyz \leqslant 0$). The constraint $f \geqslant 0$ is satisfied whenever $1 - u - v - w + 2\sqrt{(uvw)} \geqslant 0$ (the case with the larger value).

Topology of ∂R':

R' is the image of the convex body R under a continuous surjection, so it is compact and connected. The boundary structure is determined by examining R' restricted to each face of $[0,1]^3$ (Table 7):

On each {var = 0} face: Setting $u = 0$ gives $g(0,v,w) = 1 - v - w$, so the constraint is $v + w \leqslant 1$. This is a **linear** triangle in the remaining two variables — exactly the linear monogamy constraint. The three coordinate planes each contribute a flat triangular boundary patch.

On each {var = 1} face: Setting $u = 1$ gives $g(1,v,w) = -v - w + 2\sqrt{(vw)} = -(\sqrt{v} - \sqrt{w})^2 \leqslant 0$, with equality only when $v = w$. The intersection is a line segment from the relevant axis vertex to (1,1,1).

Vertices (0-cells): The cube vertices of $[0,1]^3$ satisfying $g \geqslant 0$ (Table 6) are:

Vertex	g value	Status
(0,0,0)	1	Interior of R' (deep inside)
(1,0,0), (0,1,0), (0,0,1)	0	On boundary
(1,1,1)	0	On boundary
(1,1,0), (1,0,1), (0,1,1)	−1	Outside

The vertex (0,0,0) represents all features mutually orthogonal ($\cos^2\theta = 0$ for all pairs). This is a well-defined point in R (indeed, $f(0,0,0) = 1$, the global maximum), mapping to a corner of $[0,1]^3$. It becomes an extreme point of R' because the image space $[0,1]^3$ has a boundary at $u = v = w = 0$ that the interior of R maps onto.

Total: V = 5 vertices: (0,0,0), (1,0,0), (0,1,0), (0,0,1), (1,1,1).

Edges (1-cells): Nine edges, naturally grouped:

- Three coordinate edges: (0,0,0)→(1,0,0), (0,0,0)→(0,1,0), (0,0,0)→(0,0,1). These lie on pairwise intersections of {var = 0} faces.
- Three triangular edges: (1,0,0)→(0,1,0), (0,1,0)→(0,0,1), (0,0,1)→(1,0,0). These lie on individual {var = 0} faces, along the line u + v = 1 (etc.), which is the boundary of the linear triangle.
- Three {var = 1} edges: (1,0,0)→(1,1,1), (0,1,0)→(1,1,1), (0,0,1)→(1,1,1). These are line segments on the {var = 1} faces, parameterised by the constraint $\sqrt{v} = \sqrt{w}$ (etc.).

Faces (2-cells): By Euler's formula, F = 2 − V + E = 2 − 5 + 9 = 6. These are:

- Three flat triangular faces on the {var = 0} coordinate planes, each bounded by two coordinate edges and one triangular edge.
- Three curved faces connecting the triangular edges to (1,1,1) via the {var = 1} edges, each a patch of the surface g = 0.

$$V = 5, \quad E = 9, \quad F = 6, \quad \chi = V - E + F = 2, \quad V + \chi = 7$$

A.2.3 Where the Fifth Vertex Comes From

The four vertices of R in dot-product space all map to the same point under the squaring map: $(\pm 1)^2 = 1$, so all four go to (1,1,1). The tetrahedron's four vertices collapse to a single vertex of R'.

The fifth vertex (0,0,0) in correlation-magnitude space — representing all features mutually orthogonal — is the image of the interior maximum of R, not of any boundary vertex. The point (0,0,0) in dot-product space satisfies f(0,0,0) = 1, placing it at the centre of R (the configuration of maximum "room" within the Gram constraint). Under the squaring map, this interior point maps to a corner of $[0,1]^3$, where it becomes extreme because the image space has a boundary that the domain space does not.

So the combinatorial change V: 4 → 5 happens because:

- The $(\mathbb{Z}_2)^3$ gauge symmetry of R (independent sign-flip of each dot product) is quotiented by the squaring map
- Four gauge-inequivalent vertices of R merge into one gauge-invariant vertex of R'
- The interior maximum of R maps to a corner of the image space, becoming a new extreme point

The fifth vertex is **measurement-induced**: it exists because gauge-invariant measurement of correlation magnitudes cannot distinguish positive from negative correlation. The intrinsic constraint geometry has V=4 (dot-product space); the observable constraint geometry has V=5 (correlation-magnitude space).

A.3 Why $V + \chi = 7$ Is the Physically Correct Count

The α and θ_W derivations use correlation magnitudes, not signed dot products. The electromagnetic coupling α measures the STRENGTH of correlation between charged features — it doesn't distinguish whether the correlation is positive or negative. Similarly, the Weinberg angle measures the MAGNITUDE of mixing between U(1) and SU(2) sectors.

The gauge argument is as follows. Each feature $a_i \in Cl(5)$ is a grade-1 vector. The vectors a_i and $-a_i$ produce identical bivector magnitudes with every other feature: $|a_i \wedge a_j|^2 = 1 - (a_i \cdot a_j)^2 = 1 - ((-a_i) \cdot a_j)^2$. They therefore represent the same relational profile — the same distinguishable entity. This $\mathbb{Z}_2$ ambiguity for each of the N features gives a $(\mathbb{Z}_2)^N$ gauge group. The signed dot product $a_i \cdot a_j$ transforms nontrivially under this group (flipping the sign of either a_i or a_j negates the product); the squared dot product $(a_i \cdot a_j)^2 = \cos^2\theta_{ij}$ is gauge-invariant.

An observer measuring the relational field therefore has access to the gauge-invariant quantities — the correlation magnitudes — not to the signed dot products. The physical monogamy region is R' in correlation-magnitude space, and the physical vertex count is $V = 5$.

Therefore:

- The physical monogamy region is R' in correlation-magnitude space
- The physical vertex count is $V = 5$ (the extreme configurations accessible to gauge-invariant measurement)
- The physical topological weight is $V + \chi = 7$
- The linear polytope and the Gram region AGREE on this count despite having different shapes and volumes

The dot-product space analysis ($V=4$, $V+\chi=6$) describes the INTRINSIC geometry of the constraint. The correlation-magnitude space analysis ($V=5$, $V+\chi=7$) describes the OBSERVABLE geometry — what an observer measuring correlation strengths would reconstruct. The α formula uses the observable geometry because α is itself an observable quantity.

A.3.1 Connection to Trivector Gauge Equivalence

The $(\mathbb{Z}_2)^N$ gauge structure manifests concretely at the level of trivector Gram matrices. For four features in $\mathbb{R}^5$ with all pairwise |dot products| equal to $d = 1/4$, there are three distinct sign patterns on the dot products that produce exact eigenvalue spectrum {5, 1, 1, 1} (proportional to the I+J matrix), classified by graph-theoretic structure (Table 9 in companion script):

- **Matching-negative** (4 positive, 2 negative dot products): negative edges form a perfect matching of K_4 (orbit size 3 under S_4, stabiliser order 8)
- **Triangle-negative** (3 positive, 3 negative): negative edges form a triangular face of K_4 (orbit size 4, stabiliser order 6)
- **All-negative** (0 positive, 6 negative): all dot products have the same sign (orbit size 1, stabiliser order 24)

These are not distinct geometries. Each produces a trivector Gram matrix TG that factors as $\alpha \cdot D(I+J)D^T$ where $D = \mathrm{diag}(\pm1, \pm1, \pm1, \pm1)$ is a diagonal sign matrix (Table 9). The matrix D is precisely the $(\mathbb{Z}_2)^4$ gauge transformation acting on the four triangular faces of K_4. The gauge-invariant content — the eigenvalue spectrum — is identical for all three. There are 3 + 4 + 1 = 8 distinct sign patterns in total (Table 11), all related by $(\mathbb{Z}_2)^4$ gauge transformations.

This provides a concrete example of the general principle: signed quantities (dot products, trivector Gram entries) carry gauge-dependent information; their magnitudes and eigenvalue spectra carry the physically meaningful content.

A.4 The Linear Polytope as a Sub-Region

The linear monogamy polytope (the original formulation) is a strict subset of the Gram region R' in correlation-magnitude space (Table 8):

Property	Linear polytope	Gram region R'
Volume (fraction of $[0,1]^3$)	0.250	0.713
Boundary	Flat (planar faces)	Curved (algebraic surface)
"Democratic" vertex	(½, ½, ½) — extreme point	(½, ½, ½) — interior (g = 0.207 > 0)
Maximal vertex	Not present	(1, 1, 1) — extreme point
V, E, F, χ	5, 9, 6, 2	5, 9, 6, 2

Every configuration in the linear polytope is in R'. This follows from the fact that the linear constraints $u + v \leqslant 1$, $v + w \leqslant 1$, $u + w \leqslant 1$ (with $u, v, w \geqslant 0$) imply $1 - u - v - w + 2\sqrt{(uvw)} \geqslant 0$ (the Gram constraint). Proof: summing the three linear constraints gives $u + v + w \leqslant 3/2$, so $1 - u - v - w \geqslant -1/2$. Meanwhile $2\sqrt{(uvw)} \geqslant 0$. But the sharper argument is: on each {var = 0} face, the Gram constraint reduces to the linear constraint exactly, and the linear polytope's interior maps into R's interior. Monte Carlo verification over 5×10^6 samples confirms: zero points satisfy the linear constraints but violate the Gram constraint (Table 8).

The physical interpretation: the linear constraint $\lambda\_AB + \lambda\_CA \leqslant \Lambda$ is a sufficient condition for monogamy. It's stricter than necessary — it excludes configurations that the Gram determinant allows. But it preserves the correct topology and vertex structure. The extra volume in R' (0.713 − 0.250 = 0.463, or 46% of the unit cube) represents configurations where the pairwise correlations exceed the linear budget but the three features are still genuinely distinguishable (det $G_3 > 0$).

A.5 Comparison of the Two Formulations

Aspect	Linear polytope	Gram region R'	Implication
Constraint	$\lambda\_AB + \lambda\_CA \leqslant \Lambda$ (linear)	det $G_3 \geqslant 0$ (cubic in dot products)	Gram is fundamental; linear is sufficient
V	5	5	✓ Same
E	9	9	✓ Same
F	6	6	✓ Same
χ	2	2	✓ Same
V + χ	7	7	✓ α and θ_W derivations preserved
Volume	1/4	0.713	≠ Different
Fifth vertex	(½, ½, ½) democratic	(1, 1, 1) all-maximal	Different physical meaning
Boundary	Flat	Curved	Different differential geometry

The derivations of α and θ_W use V + χ and 5 × 3!. Neither depends on volume, boundary curvature, or the specific location of the fifth vertex. The topological invariants are robust to the choice of formulation.

A.6 Implications for the Weinberg Angle Discrepancy

A.6.1 The Current Status

The base formula $\sin^2\theta\_W = 7/30 = 0.2333$ agrees with experiment (0.23121 at M_Z) to 0.9%. The refined formula 49/212 = 0.23113 agrees to 0.03%. The residual discrepancy is approximately

0.00008, or 0.035%.

A.6.2 The Volume Ratio as a Potential Correction

The two formulations have different volumes:

- Linear polytope: 1/4 of the unit cube
- Gram region: 0.713 of the unit cube
- Ratio: 0.713 / 0.250 = 2.853

The physical monogamy constraint (Gram) allows 2.85× more configuration space than the linear approximation. If the Weinberg angle depends not only on the topological weight V+χ but also on the FRACTION of configuration space that is monogamy-constrained, the volume ratio enters.

The base formula uses (V+χ) / (5 × 3!) = 7/30, which can be interpreted as "topological weight of the constrained region divided by the embedding complexity." If we weight this by the volume fraction:

$$\sin^2 \theta_W = \frac{V + \chi}{5 \times 3!} \times f(V_{\text{Gram}}/V_{\text{cube}})$$

where f is some function of the volume ratio. The volume of the Gram region relative to the cube is 0.713, and the linear polytope's volume is 0.250. The ratio of volumes is:

$$\frac{V_{\text{linear}}}{V_{\text{Gram}}} = \frac{0.250}{0.713} = 0.3506$$

This represents the fraction of the geometrically valid configuration space that satisfies the stricter linear budget constraint.

A.6.3 A Speculative Interpretation of Running

The experimental $\sin^2\theta_W$ runs with energy:

- At M_Z (91 GeV): 0.23121
- At low energy: ~0.238
- At GUT scale: ~0.21 (projected)

In the framework, energy scale might correspond to the "tightness" of the monogamy constraint. At high energy (short distances, strong coupling), the constraint approaches the linear limit (volume → 1/4, tighter restriction). At low energy (long distances, weak coupling), the constraint relaxes toward the Gram limit (volume → 0.713, looser restriction).

If $\sin^2\theta_W$ interpolates between the two formulations:

- Tight (linear) limit: $\sin^2\theta_W \to 7/30 = 0.2333$
- Loose (Gram) limit: $\sin^2\theta_W \to 7/30 \times$ (volume correction)
- Physical value at M_Z: somewhere between

This would mean the running of θ_W with energy reflects the progressive tightening of the monogamy constraint at shorter distances — a geometric interpretation of the electroweak running that doesn't invoke virtual particles or loop corrections.

Note that the direction is physically natural: the linear polytope *excludes* the high-correlation region of configuration space (its democratic vertex sits at (½,½,½), far from the maximal-correlation boundary). "Tighter" at high energy means features are prevented from becoming highly correlated — which has the flavour of asymptotic freedom. This is the regime where the standard model also predicts $\sin^2\theta_W$ approaches 7/30 from below.

A.6.4 Why Different Measurement Procedures Give Different Values

The experimental $\sin^2\theta_W$ depends on how it's measured:

- From Z-pole asymmetries: 0.23153 ± 0.00016
- From W mass: 0.2229 ± 0.0004
- From neutrino scattering: 0.2397 ± 0.0013

These differ by more than their uncertainties. In the standard model, this is attributed to radiative corrections (different processes probe different loop structures).

In the framework, different measurement procedures correspond to different observation protocols — different ways of projecting the monogamy-constrained region. The correlation-magnitude space R' has curved boundaries, and different measurement protocols project onto different cross-sections. A protocol that accesses the flat (linear) part of the boundary sees a value closer to 7/30. A protocol that accesses the curved part sees a different effective geometric weighting.

This is the measurement-as-construction principle at the level of fundamental constants: $\sin^2\theta_W$ is not a single number but a family of protocol-dependent values, all of which are consistent with the same underlying topology ($V+\chi=7$) but which differ in the geometric weight they assign to different parts of the monogamy region.

The prediction: the spread of experimental $\sin^2\theta_W$ values across different measurement procedures should be bounded by the geometric range of the monogamy region. The maximum variation is set by the difference between the linear polytope value and the Gram region value.

A.6.5 Testable Consequences

If this interpretation is correct, the range of measured $\sin^2\theta\_W$ values should correlate with how each measurement protocol samples the monogamy region — specifically, which cross-section of R' each observable projects onto.

The framework predicts that the spread of $\sin^2\theta\_W$ values is geometrically bounded: all measurements must yield values consistent with $V+\chi = 7$, but the effective geometric weight can vary within a range set by the difference between the linear and Gram formulations. The full mapping between specific experimental protocols and specific cross-sections of R' remains to be worked out; until this is done, the framework predicts the *existence* of a bounded spread without committing to which direction each protocol shifts.

The intrinsic geometry (dot-product space) has $V+\chi = 6$, while the observable geometry (correlation-magnitude space) has $V+\chi = 7$. Whether any measurement protocol could access the intrinsic value — which would require sign-sensitivity beyond gauge-invariant correlation magnitudes — is an open question. If such a protocol exists, it would see effective $\sin^2\theta\_W \to 6/30 = 0.200$, a sharp and falsifiable prediction.

A.7 Implications for the Fine Structure Constant

A.7.1 Topological Robustness

The formula $\alpha = \sqrt{3}/(24\pi^2 + \sqrt{7/30})$ uses the same $V+\chi=7$ that survives the coordinate transformation. The formula is therefore topologically robust — it does not depend on whether the monogamy constraint is formulated linearly or via the Gram determinant.

A.7.2 Volume Dependence

If a future refinement of the α derivation introduces volume-dependent corrections (analogous to the $\chi/(V+\chi)$ correction in the θ_W formula), the volume to use would be ambiguous between 1/4 (linear) and 0.713 (Gram). The ratio $0.250/0.713 \approx 0.351$ might enter as a correction factor.

At present, the formula achieves 1 ppm agreement without volume corrections, suggesting that α depends purely on topological and symmetry factors ($V+\chi$, $3!$, $(2\pi)^2$, $\sqrt{3}$) and not on metric properties of the monogamy region. This is consistent with α being a dimensionless coupling constant — it should be determined by counting and topology, not by volumes and metrics.

A.8 Connection to the Trivector Gram Structure

For four unit vectors in $\mathbb{R}^5$, there are $C(4,3) = 4$ triangular faces, giving a 4×4 trivector Gram matrix TG where $TG[f_1,f_2] = T\_\{f_1\} \cdot T\_\{f_2\}$ is the inner product of the trivectors associated with two faces. The entry $TG[f_1,f_2]$ is equal to the determinant of the 3×3 cross-Gram matrix between the features of face f_1 and face f_2 (verified numerically to machine precision).

A.8.1 The One-Parameter Family

When all six pairwise |dot products| are equal to some value $d \in (0, 1/3)$, the trivector Gram eigenvalue spectrum has the form {λ_big, λ_small, λ_small, λ_small} — a threefold degeneracy — for *every* value of d, regardless of the sign pattern. The maximum spread $|\lambda_2 - \lambda_4|$ across 1000 values of d is less than 2×10^{-15} — machine epsilon, confirming exact degeneracy (Table 12 in companion script). The eigenvalue ratio R = λ_big/λ_small is a smooth monotonic function of d (Table 10):

d	Ratio R	Corresponding matrix structure
0	1	αI (isotropic, no edge-sharing)
1/5	**3**	**α(2I+J)** — the original prediction
1/4	**5**	**α(I+J)** — the numerically optimal value
1/3	∞	Rank 1 (equilateral simplex, maximally degenerate)

The "2I+J target" from the Three-Dimensional Geometry paper and the "I+J optimum" found by numerical optimisation are not competing predictions — they are two specific points on a continuous one-parameter family. The value d=1/4 (giving eigenvalue ratio 5:1:1:1, i.e. the I+J structure) is special in that it produces the feature Gram eigenvalues {5/4, 5/4, 5/4, 1/4} — a 3:1 degeneracy in the *feature* Gram matrix that echoes the threefold degeneracy in the trivector Gram.

A.8.2 What Selects d?

The algebra provides a family; a selection principle is needed to pin d to a specific value. Candidates include:

- The monogamy constraint (the viable region boundary): some d values may be preferred by the Gram determinant structure
- The efficiency potential Φ: gradient flow on the viable region may have attractors at specific d values
- The distinguishability threshold ε: the minimum pairwise distance $|a_i - a_j|^2 \geq \varepsilon^2$ constrains the maximum |dot product|, and the relationship d(ε) may select a specific ratio

This is a sharp, answerable question for future work.

A.9 Summary

1. The Gram determinant det $G_3 \geqslant 0$ is the fundamental monogamy constraint. The linear budget constraint λ_AB + λ_CA ⩽ Λ is a sufficient but not necessary condition. The linear constraint is recovered exactly on the {var = 0} faces of correlation-magnitude space, where it appears as the boundary of the Gram region restricted to one-vanishing-correlation planes.
2. In dot-product space $[-1,1]^3$, the Gram region is a spectrahedron (convex, compact, with nonempty interior). Its boundary has the combinatorial type of a tetrahedron: V=4, E=6, F=4, χ=2. The four vertices are rank-1 correlation matrices forming a regular tetrahedron inscribed in the cube.
3. In correlation-magnitude space $[0,1]^3$ (the gauge-invariant coordinates under the $(\mathbb{Z}_2)^N$ orientation group), the Gram region has V=5, E=9, F=6, χ=2, V+χ=7. The fifth vertex arises because the squaring map sends the interior maximum of R to a corner of the image space.
4. The linear polytope and the Gram region have IDENTICAL combinatorial type in correlation-magnitude space. V+χ=7 is a topological invariant that does not depend on whether the boundary is flat or curved, or whether the volume is 1/4 or 0.713.
5. The α derivation ($\alpha = \sqrt{3}/(24\pi^2 + \sqrt{(7/30)}) = 1/137.036$) is topologically robust and survives the correction from linear to Gram formulation.
6. The θ_W discrepancy (0.2333 vs 0.2312 experimental) may reflect the difference between the linear and Gram formulations, with the physical value interpolating between them. Different measurement procedures may probe different parts of the monogamy region, producing protocol-dependent values bounded by the geometric range.
7. The two coordinate systems correspond to two levels of description: dot-product space describes the intrinsic geometry of the relational field; correlation-magnitude space describes what an observer measuring gauge-invariant correlation strengths can reconstruct. The physical predictions use the observer-accessible coordinates, consistent with the framework's emphasis on observer-dependent construction.
8. The trivector Gram structure at N=4 exhibits a one-parameter family of exact threefold degeneracy indexed by d = |dot product|, with 2I+J and I+J as specific members. The sign patterns that distinguish apparently different configurations are gauge artefacts of the $(\mathbb{Z}_2)^4$ orientation group, with identical gauge-invariant eigenvalue spectra. The question of what selects a particular d value connects to the dynamics of the viable region.

Companion Script

All quantitative claims in this addendum are verified by `monogamy_topology_verification.py`, which requires only NumPy and SciPy. The script produces Tables 1–12 referenced throughout.