Witness Geometry Beyond Scalar Fee

Repair Entropy, Operational Sparsity,
and the Decision Theory of Disclosure

John Komkov
Independent Researcher

April 2026

Abstract

The coherence fee of a tool composition is a single integer measuring hidden convention obligation. Working in the linear disclosure model (repair = field exposure) under dimension-field disjointness, we show that the fee is the first of three increasingly operational projections of the witness Gram $K(G)$. The per-dimension witness matroid is uniform—$U(n{-}1, n)$ on each connected component—correcting Paper III Corollary 4.3. We define *repair entropy* $H_{\text{repair}} = \log\beta$, where $\beta = \prod_{d,j} |C_{d,j}|$ is the product of witness-component sizes, and prove sharp bounds: $\varphi{+}1 \leq \beta \leq 2^{\varphi}$, both attained. But repair entropy measures *structural* flexibility—the size of the exact repair space. Under cost perturbation, most of that space is operationally irrelevant. We prove that every basis is the unique optimum under some cost vector (full realizability), so operational sparsity is not intrinsic to the matroid but depends on the cost family. On a 240-composition MCP corpus (239 testable for basis enumeration at fee $\leq$ 14): (i) β(formula) = β(enumerated) on all 239; (ii) 81.7% of compositions sit at the rigid bound $\beta = \varphi{+}1$; (iii) under integer costs in $[1, 10]$, at $\beta = 378$ only 13.5% of bases are ever optimal. Fee counts the obligations. Repair entropy counts the structural repair space. Operational sparsity under a given cost family identifies the decision-relevant subset that survives deployment constraints. In this regime, the exact minimum repair cost is total active hidden cost minus the sum of per-component cost maxima (the *geometry dividend*). Decision activation—the condition under which witness geometry changes the optimal repair—requires both structural repair multiplicity and cost heterogeneity on the witness support; removing either ingredient deactivates the geometry.

1 Introduction

The coherence fee $\text{fee}(G) = \text{rank}(\delta_{\text{full}}) - \text{rank}(\delta_{\text{obs}})$ measures how many hidden convention obligations a tool composition carries. Papers I–III of this program [1, 2, 3] established the fee as a matroid corank (Paper II [2]), showed its rank lives in the witness Gram $K(G) = H_c^{\top}(I - P_O)H_c$ (Paper III [3]), and proved that no bounded-local audit can compute it (the obstruction paper [4]).

But the fee is a single integer. Two compositions with fee = 11 may have very different repair landscapes: one may have a unique minimum disclosure set, the other hundreds of equally valid alternatives; one may require disclosing expensive fields, the other only cheap ones. The scalar fee cannot distinguish these cases.

This paper asks: **what invariant, beyond fee, determines the structure of repair?**

The answer is the witness Gram $K(G)$ itself—and the central new invariant it yields is **repair entropy**.

We prove four results:

1. **K-sufficiency (Theorem 1).** In the linear disclosure model (repair = field exposure), $K(G)$ determines the witness matroid M/O and therefore the exact repair basis structure:

basis count, minimum-cost repair under any cost vector, and cost sensitivity. The fee is the rank of K; everything else lives in its finer structure.

2. **Uniform witness structure and repair entropy (Theorem 3).** Under DFD + CHP, the per-component witness matroid is $U(n{-}1,\, n)$—the uniform matroid, not the graphic matroid. The corrected basis count is $\beta(G) = \prod_{d,j} |C_{d,j}|$, and the repair entropy $H_{\mathrm{repair}} = \sum \log |C_{d,j}|$ is additive over witness components. Fixed fee does not determine repair entropy.

3. **Sharp bounds at fixed fee (Theorem 8).** For any composition with fee φ in the DFD + CHP regime: $\varphi{+}1 \le \beta(G) \le 2^{\varphi}$, with both bounds attained.

4. **Full realizability (Theorem 10).** Every basis of M/O is the unique minimum-cost repair under some cost vector. Operational sparsity is a property of the cost family, not the matroid.

5. **Exact repair functional (Theorem 12).** In the DFD + CHP regime, the minimum repair cost is total active hidden cost minus the sum of per-component cost maxima. The *geometry dividend* $A(G, c)$ is the operational value of witness geometry.

6. **Decision activation (Proposition 19).** Witness geometry changes the optimal repair if and only if the composition has structural repair multiplicity *and* the cost vector is heterogeneous on the witness support. Neither condition alone suffices. This follows algebraically from Theorem 12.

The paper also corrects an error in Paper III Corollary 4.3, which stated $|B(M/O)| = \prod \tau(\tilde{G}_d)$ (spanning-tree counts). The correct formula is $\beta = \prod |C_{d,j}|$ (component-size products). See §4 and the correction memo for the full impact analysis.

Notation. See the companion NOTATION.md for the canonical symbol table. We follow Papers II and III throughout.

2 K-Sufficiency: The Witness Gram Determines Repair Structure

2.1 The characterization theorem

Theorem 1 (K-sufficiency)**.** *Let G be a Bulla composition with linear disclosure model (field exposure as the repair action). The witness Gram $K(G) = H_c^\top (I - P_O) H_c$ determines:*

(a) The witness matroid M/O: a subset $S \subseteq H$ is independent in M/O if and only if the columns of K indexed by S are linearly independent. (Independence functions in this sense were axiomatized by Rado [6].)

(b) All bases $B(M/O)$: the minimum disclosure sets of G are exactly the column-bases of $K(G)$.

(c) The repair multiplicity $\beta(G) = |B(M/O)|$.

(d) For any cost vector $c : H \to \mathbb{Q}_{>0}$, the minimum-cost repair basis $B^*(G, c)$ and minimum repair cost $\Sigma^*(G, c)$.

(e) The leverage scores $\ell_j = (K^+K)_{jj}$ and effective resistances $R_{\mathrm{eff}}(i, j)$.

Proof. (a) By Paper III [3] Theorem 2.1, the column matroid of $K(G) = W^\top W$ coincides with the column matroid of $W = (I - P_O)H_c$, which is the contraction M/O. Hence $S \subseteq H$ is independent in M/O iff the S-columns of K are linearly independent.

(b) Immediate from (a): bases of the column matroid of K are bases of M/O.

(c) Follows from (b): $\beta(G) = |B(M/O)|$ is determined by K.

(d) Given $B(M/O)$ from (b) and a cost vector c, the minimum-cost basis is $B^*(G, c) = \arg\min_{B \in B(M/O)} \sum_{h \in B} c(h)$. This is a function of $B(M/O)$ and c, both determined by K and c. Computationally, this is the matroid greedy problem of Edmonds [5], with $B^*(G, c)$ recovered in $O(|H| \log |H|)$ time given the independence oracle on M/O.

(e) Standard: leverage and effective resistance are defined in terms of K and K^+. □

Remark 2 (What K-sufficiency does and does not say). $K(G)$ is sufficient for repair combinatorics—it determines what can be repaired, at what cost, and with what alternatives. It is *not* sufficient for all properties of G: it does not determine the composition graph, the tool schemas, or the observable structure (these are inputs to K, not outputs). The sufficiency is for one specific question: given that a composition has hidden obligations, what is the structure of resolving them?

2.2 K versus fee: a hierarchy of invariants

The witness Gram $K(G)$ sits at the top of an invariant hierarchy, each level a lossy compression of the one above:

$$\underbrace{K(G)}_{\text{full Gram}} \xrightarrow{\text{matroid}} \underbrace{M/O}_{\text{independence oracle}} \xrightarrow{\text{bases}} \underbrace{B(M/O)}_{\text{basis set}} \xrightarrow{\text{count}} \underbrace{\beta(G)}_{\text{multiplicity}} \xrightarrow{\text{rank}} \underbrace{\text{fee}(G)}_{\text{integer}}$$

Each arrow discards information. Fee discards all basis structure. β discards which bases exist. M/O discards the metric structure (effective resistance). K retains everything.

3 Fee Does Not Determine Repair Multiplicity

3.1 The separation theorem

Theorem 3 (Fee–multiplicity separation)**.** *For every $\varphi \geq 2$, there exist Bulla compositions G_0 and G_1 with:*

1. $\text{fee}(G_0) = \text{fee}(G_1) = \varphi$
2. $\beta(G_0) \neq \beta(G_1)$

Proof. We work in the DFD + full CHP regime (all convention fields hidden, each field in exactly one dimension), so that $K(G) = \bigoplus_d L_d$.

The correct basis-count formula. Under DFD + full CHP with each carrier graph G_d connected on n_d hidden vertices, the column matroid of the signed incidence matrix δ_d is the uniform matroid $U(n_d - 1, n_d)$: every $(n_d - 1)$-subset of columns is a basis. (This is because each column of δ_d is a signed characteristic vector of a vertex, and any $n_d - 1$ vertices of a connected graph are linearly independent in the signed incidence matrix.) Under the direct-sum decomposition, the total basis count is:

$$\beta(G) = \prod_d n_d$$

Each connected carrier graph G_d contributes $\text{fee}_d = n_d - 1$ and n_d bases to the product. The fee constrains the sum $\sum(n_d - 1) = \varphi$, but the product $\prod n_d = \prod(\text{fee}_d + 1)$ depends on the partition of φ across dimensions—not just the total.

Construction.

G_0: φ dimensions, each with carrier graph K_2 (2 hidden fields, one bilateral edge). Partition: $\mathrm{fee}_d = 1$ for all d.

- $\mathrm{fee}(G_0) = \varphi \cdot 1 = \varphi$
- $\beta(G_0) = 2^{\varphi}$

G_1: $(\varphi - 2)$ dimensions with carrier graph K_2, plus one dimension with carrier graph C_3 (3 hidden fields in a triangular routing). Partition: fee $= 1 \cdot (\varphi - 2) + 2 = \varphi$.

- $\mathrm{fee}(G_1) = (\varphi - 2) \cdot 1 + (3 - 1) = \varphi$
- $\beta(G_1) = 2^{\varphi-2} \cdot 3$

Verification. $\beta(G_0)/\beta(G_1) = 2^{\varphi}/(2^{\varphi-2} \cdot 3) = 4/3 \neq 1$, so $\beta(G_0) \neq \beta(G_1)$ for all $\varphi \geq 2$. Computationally verified for $\varphi \in \{2, 3, 5, 10\}$ by brute-force basis enumeration:

φ	$\beta(G_0)$	$\beta(G_1)$	Ratio
2	4	3	4/3
3	8	6	4/3
5	32	24	4/3
10	1024	768	4/3

□

Remark 4 (Paper III [3]. Corollary 4.3 correction] Paper III [3] states $|B(M/O)| = \prod_d \tau(\tilde{G}_d)$ (product of spanning-tree counts). This conflates the graphic matroid (ground set = edges, bases = spanning trees) with the column matroid of the signed incidence matrix (ground set = vertices, bases = $(n-1)$-subsets of vertices). The correct formula in the DFD + full CHP regime with connected carrier graphs is $\beta = \prod_d n_d$. The two coincide only when $n_d = \tau(G_d)$ for all d—for example, when every carrier graph is a cycle (where $\tau(C_n) = n$). For general connected graphs, $n \neq \tau$.

Corollary 5 (Unbounded multiplicity gap)**.** *For every $\varphi \geq 2$, the ratio $\beta_{\max}/\beta_{\min}$ among fee-φ compositions is unbounded as the number of available tools grows.*

Proof. The partition $[\varphi]$ (single dimension with $n = \varphi + 1$ hidden fields) gives $\beta = \varphi + 1$. The partition $[1, 1, \ldots, 1]$ (φ K_2-dimensions) gives $\beta = 2^{\varphi}$. Since $2^{\varphi}/(\varphi + 1) \to \infty$, the ratio is unbounded. □

Corollary 6 (Multiplicity is partition-determined)**.** *Under DFD + full CHP with connected carrier graphs, $\beta(G) = \prod_d(\mathrm{fee}_d + 1)$ is determined entirely by the fee partition* $(\mathrm{fee}_1, \ldots, \mathrm{fee}_D)$*— the distribution of the scalar fee across convention dimensions.*

Two compositions with the same fee and the same fee partition have the same β, but compositions with the same fee and different partitions generically have different β. The fee partition is a strictly finer invariant than the scalar fee.

3.2 Repair entropy

Definition 7. *The* repair entropy *of a composition G is $H_{\text{repair}}(G) = \log \beta(G) = \sum_{d,j} \log |C_{d,j}|$.*

Invariance class. H_{repair} is a matroid invariant of M/O: it depends only on the witness matroid, not on the choice of coboundary representation. Under DFD + CHP it decomposes additively over carrier graph components. It is invariant under tool renaming, field renaming, and edge reorientation.

Regime of validity. The product formula $\beta = \prod |C_{d,j}|$ and the additive decomposition $H_{\text{repair}} = \sum \log |C_{d,j}|$ require DFD (dimension-field disjointness) and CHP (convention-hidden partition). Without DFD, the witness Gram $K(G)$ is a single Kron-reduced Laplacian and the basis count does not factor by dimension. The definition $H_{\text{repair}} = \log |B(M/O)|$ remains valid in general, but the closed-form product formula does not.

What repair entropy does not measure. H_{repair} counts minimum-size disclosure sets (bases of M/O). It does not measure:

- the *cost* of repair (which requires a cost model $c : H \to \mathbb{Q}$),
- the *stability* of the optimal repair under cost perturbation (which requires the cost-sensitivity geometry),
- or the *difficulty* of implementing a repair (which is external to the matroid).

H_{repair} says how many interchangeable minimum disclosure configurations exist—not how hard, how expensive, or how robust any particular repair is.

The fee–entropy pair. Fee and repair entropy are two projections of the witness Gram:

- **fee** = obligation count (additive: $\sum(n_{d,j} - 1)$)
- H_{repair} = repair flexibility (additive: $\sum \log n_{d,j}$)

Fee says how many obligations exist. Repair entropy says how many interchangeable ways those obligations can be satisfied.

3.3 Sharp bounds at fixed fee

Theorem 8 (Sharp repair-entropy bounds)**.** *For any composition with fee $\varphi \geq 1$:*

$$\varphi + 1 \leq \beta(G) \leq 2^{\varphi}$$

equivalently:

$$\log(\varphi + 1) \leq H_{\text{repair}}(G) \leq \varphi \cdot \log 2$$

Both bounds are attained:

- ***Minimum*** *$\beta = \varphi + 1$: all fee in one connected witness component of size $\varphi + 1$ (rigid repair—few alternatives).*
- ***Maximum*** *$\beta = 2^{\varphi}$: fee distributed across φ independent bilateral pairs, each of size 2 (flexible repair—exponentially many alternatives).*

Proof. The repair multiplicity $\beta = \prod n_i$ where $n_i = |C_{d,j}| \geq 2$ and $\sum(n_i - 1) = \varphi$. Set $m_i = n_i - 1 \geq 1$; then $\sum m_i = \varphi$ and $\beta = \prod(m_i + 1)$.

Upper bound. If any $m_i \geq 2$, split it: replace $m_i = a$ with two parts $m_i = 1$ and $m_{\text{new}} = a - 1$. The product changes by factor $2a/(a+1) > 1$ for $a \geq 2$. Repeat until all $m_i = 1$, giving $\beta = 2^{\varphi}$.

Lower bound. If there are two parts $m_i = a$ and $m_j = b$, merge them into $m_i = a + b$. The product changes by factor $(a + b + 1)/((a + 1)(b + 1)) < 1$ since $ab > 0$. Repeat until one part remains, giving $\beta = \varphi + 1$.

Attainment. All-K_2 compositions (φ bilateral pairs, each with 2 hidden fields) give $\beta = 2^{\varphi}$. A single carrier component on $\varphi + 1$ hidden fields gives $\beta = \varphi + 1$. □

Computational verification: exhaustive enumeration of all integer partitions of φ confirms the bounds for $\varphi \in \{1, 2, 3, 4, 5, 8, 10\}$:

φ	$\beta_{\min}$	$\beta_{\max}$	Distinct β values
1	2	2	1
2	3	4	2
3	4	8	3
5	6	32	7
10	11	1024	39

At fee = 10, there are 39 distinct values of β—39 structurally distinguishable repair landscapes, all invisible to the scalar fee.

3.4 What Theorems 3 and 8 mean together

Two compositions with fee = 11 have:

- $\beta \in [12, 2048]$—a 170× range in repair flexibility
- The rigid extreme ($\beta = 12$): one large witness component, nearly unique repair path
- The flexible extreme ($\beta = 2048$): 11 independent bilateral obligations, each with a binary choice

The scalar fee says "11 obligations." The repair entropy says whether those obligations are rigid or flexible—and by how much. A system near $\beta_{\min}$ is brittle; a system near $\beta_{\max}$ has deep redundancy in its repair landscape.

This is not theoretical. In the MCP corpus (§4), fee-matched groups show up to 14× variation in β—variation invisible to the scalar fee but fully predicted by K. The corpus also reveals an empirical regularity: most real-world compositions sit near the rigid bound, making the rare flexible cases structurally diagnostic.

4 K-Sufficiency in Practice: The MCP Corpus

4.1 Corpus and method

We compute $\beta(G)$ and $H_{\text{repair}}(G)$ for all 240 nonzero-fee pairwise compositions in the MCP server corpus (24 servers, 703 compositions total, 240 with fee > 0). For each composition, we build the full coboundary matrix, compute $K(G) = H^\top(I - P_O)H$ over exact rationals, extract the connected components of K's nonzero off-diagonal graph, and take $\beta = \prod |C_i|$ over components.

All 240 computations succeed with zero errors. For the stability analysis (§5), basis enumeration is tractable only at fee ≤ 14, which covers 239 of the 240 compositions (the excluded composition has fee = 22). All counts in §4 use the full 240; all counts in §5 use the 239 amenable to brute-force basis enumeration.

4.2 The corpus is overwhelmingly rigid

The central empirical finding: **81.7% of corpus compositions sit at the theoretical minimum** $\beta = \varphi + 1$ (normalized flexibility $F = 0$). The corpus mean $F = 0.088$, median $F = 0.000$. No composition exceeds $F = 0.585$.

Flexibility distribution:

Range	Label	Count	Percentage
[0.0, 0.2)	rigid	196	81.7%
[0.2, 0.4)	low-flex	1	0.4%
[0.4, 0.6)	mid-flex	43	17.9%
[0.6, 1.0]	high-flex	0	0.0%

This means: for the vast majority of real-world compositions, repair structure is nearly determined by the matroid alone. The fee constrains not just how many obligations exist, but also—empirically—how rigid the repair path is.

4.3 Fee-grouped entropy profile

Fee	N	β range	F range	Distinct motifs
1	89	2	0.000	3
2	67	3	0.000	2
3	10	4–6	0.000–0.585	4
4	1	9	0.505	1
10	25	84	0.448	1
11	36	12–168	0.000–0.513	6
12	7	13–252	0.000–0.515	5
13	3	36–378	0.148–0.517	3
14	1	15	0.000	1
22	1	588	0.268	1

Two patterns emerge:

1. **At low fee (1–2), no separation exists.** All fee-1 compositions have $\beta = 2$ (single K_2 dimension); all fee-2 compositions have $\beta = 3$ (single K_3 dimension). The repair multiplicity is determined by the fee alone.
2. **At higher fee (≥ 3), separation appears.** At fee = 11, β ranges from 12 (single 12-vertex component, $F = 0.000$) to 168 (four-component motif, $F = 0.513$)—a 14× spread. At fee = 12, the range is 13–252. The fee says "11 obligations"; the repair entropy reveals whether those obligations are rigid or flexible.

4.4 Motif atlas

The 240 compositions produce only 19 distinct component-size motifs. The five most common account for 95% of the corpus:

Motif	Count	β	Fee	Class
(2)	89	2	1	binary-flex
(3)	67	3	2	rigid
(12)	30	12	11	single-block
(7, 3, 2, 2)	25	84	10	component-heavy
(3, 2)	6	6	3	mixed

The (7, 3, 2, 2) motif is notable: all 25 compositions with fee = 10 share exactly this component structure. This reflects a structural regularity in the corpus—certain servers contribute a fixed-size block of hidden conventions that recurs across pairings.

Classification. We partition the 240 compositions into mutually exclusive motif classes:

- **single-pair** (one component of size 2, fee = 1): 89 compositions (37.1%). Trivially rigid—the fee forces a unique partition.
- **single-block** (one component of size ≥ 3): 103 compositions (42.9%). $F = 0$—all fee in one carrier component, giving minimum $\beta = \varphi + 1$.
- **multi-component** (two or more nontrivial components): 48 compositions (20.0%). These are the cases where fee underdetermines repair: same fee, different partition, different β.

Among the 48 multi-component compositions, the (7, 3, 2, 2) motif alone accounts for 25 (52%)—reflecting a specific structural pattern in the corpus where `filesystem`-paired compositions produce a dominant 7-vertex block alongside three smaller convention dimensions.

5 Operational Sparsity: Most Bases Are Never Optimal

5.1 The problem with structural flexibility

Repair entropy counts minimum-cardinality disclosure sets—bases of the witness matroid M/O. A composition with $\beta = 168$ has 168 structurally valid repairs. But how many of those are ever the cheapest repair under any plausible cost model?

The answer, empirically, is: far fewer than β.

5.2 Reachable bases and operational entropy

Definition 9. *Let $\mathcal{F}$ be a family of cost vectors $c : H \to \mathbb{Q}_{>0}$. The* reachable basis set *under $\mathcal{F}$ is*

$$B_{\mathcal{F}}(G) := \bigcup_{c \in \mathcal{F}} \arg \min_{B \in B(M/O)} c(B)$$

and the operational entropy *is $H_{\text{op}}(G; \mathcal{F}) = \log |B_{\mathcal{F}}(G)|$.*

The *stability ratio* is $\rho(G; \mathcal{F}) = |B_{\mathcal{F}}(G)|/\beta(G) \in (0, 1]$: the fraction of structurally valid bases that are ever decision-relevant under the cost family $\mathcal{F}$.

5.3 The realizability theorem

Theorem 10 (Full realizability under unrestricted costs)**.** *For any Bulla composition G with* fee > 0 *and any basis $B \in B(M/O)$, there exists a cost vector $c : H \to \mathbb{Q}_{>0}$ such that B is the unique minimum-cost repair.*

Proof. Set $c(h) = 1$ for all $h \in B$, and $c(h) = 2$ for all $h \notin B$. Since every basis has cardinality $\text{fee}(G)$, $\text{cost}(B) = \text{fee}(G)$. Any other basis $B' \neq B$ has $|B' \cap B| = \text{fee}(G) - d$ and $|B' \setminus B| = d$ for some $d \geq 1$. So $\text{cost}(B') = (\text{fee}(G) - d) \cdot 1 + d \cdot 2 = \text{fee}(G) + d \geq \text{fee}(G) + 1 > \text{cost}(B)$. □

Corollary 11. *Under unrestricted linear costs, $\beta_{\mathcal{F}}(G) = \beta(G)$: every basis is reachable, and $\rho = 1$.*

This is the key structural fact. It says **operational sparsity is not a property of the matroid**. It is a property of the restricted cost family. The witness geometry admits all repairs; the deployment constraints collapse them.

5.4 Corpus stability profile

We compute stability under random integer costs $c(h) \in \{1, \ldots, 10\}$ (50 random draws + 4 structured cost models per composition). Basis enumeration limits this to fee ≤ 14 (239 of 240 compositions).

β	Reachable bases	ρ (stability ratio)
2	2	1.000
3	3	1.000
4–6	4–6	1.000
12	11	0.917
84	36	0.429
168	44	0.262
252	47–48	0.186–0.191
378	51	0.135

Two patterns:

1. **At low β, all bases are reachable.** For $\beta \leq 9$, $\rho = 1.000$—integer costs in $[1, 10]$ suffice to make every basis optimal. This is consistent with Theorem 10: the cost family is rich enough relative to the small basis set.
2. **At high β, ρ collapses.** For $\beta = 378$, only 51 of 378 bases (13.5%) are ever optimal. Most of the combinatorial repair space is operationally irrelevant.

5.5 Entropy does not fully predict stability

At fee = 11, the data shows a mild inversion:

Motif	β	Reachable	ρ
(7, 3, 2, 2, 2)	168	44	0.262
(7, 3, 3, 2)	126	39	0.310
(7, 4, 2, 2)	112	42	0.375
(8, 3, 2, 2)	96	38	0.396
(7, 3, 2, 2)	84	36	0.429
(12)	12	11	0.917

Notice: $\beta = 112$ produces 42 reachable bases, while $\beta = 126$ produces only 39. Lower entropy, more sensitivity. The component partition—not just the product—shapes the operational frontier.

This means repair entropy is a strong but not sufficient predictor of operational stability. The motif geometry carries information beyond the scalar β.

5.6 Live replay: same fee, different cheapest repair

The canonical pair is github+gtasks-mcp (motif (7, 4, 2, 2), $\beta = 112$) versus github+playwright (motif (7, 3, 3, 2), $\beta = 126$). Both have fee = 11. The structural difference is one component: the `state` dimension has 4 hidden fields in gtasks (including a cross-server match `github::state` $\leftrightarrow$ `gtasks::status`) but only 3 in playwright (all github-internal).

Under a semantic cost model that prices stateful fields at 5, path fields at 3, pagination at 1, and directional fields at 2:

Composition	Motif	β	Repair cost	Expensive omission
github+gtasks-mcp	(7, 4, 2, 2)	112	26	`gtasks::status` (cost 5)
github+playwright	(7, 3, 3, 2)	126	24	`github::state` (cost 5)

The gtasks composition pays 8.3% more for repair—not because it has more obligations (same fee), but because the cross-server `state`↔`status` match creates a 4-element component where every element is expensive. The playwright composition can spread the state obligation across 3 cheaper github-internal fields.

This is the operational meaning of motif geometry. The scalar fee sees 11 obligations in both cases. The witness Gram sees that those obligations are distributed differently—and that difference costs real money.

5.7 Interpretation

The correct interpretation of repair entropy is now:

- β measures *structural* flexibility—the size of the exact repair space intrinsic to the matroid.
- $\beta_{\mathcal{F}}$ measures *operational* flexibility—how much of that space survives contact with deployment costs.
- $\rho = \beta_{\mathcal{F}}/\beta$ measures the *compression ratio*—how aggressively cost constraints collapse the combinatorial repair space.

Witness geometry creates a large combinatorial repair space. Deployment costs collapse that space to a sparse operational frontier. The deeper the combinatorial space (higher β), the more aggressive the compression (lower ρ).

6 The Exact Repair Functional and Decision Activation

Sections 3–4 established that fee does not determine repair entropy, and that repair entropy does not determine operational stability. This section derives the *exact* minimum repair cost as a closed-form functional of the witness components and cost vector, then uses it to identify when geometry changes the recommendation.

6.1 The sum-of-maxima theorem

Theorem 12 (Exact minimum repair cost)**.** *Let G be a composition in the DFD + CHP regime with nontrivial witness components $C_1, \ldots, C_k$ (each $|C_j| \geq 2$). For any additive cost vector $c : H \to \mathbb{Q}_{>0}$, the minimum repair cost is:*

$$\Sigma^*(G, c) \;=\; \sum_{j=1}^{k}\Big(\sum_{h\in C_j} c(h) \;-\; \max_{h\in C_j} c(h)\Big) \;=\; \sum_{h\in H_{\text{active}}} c(h) \;-\; A(G, c)$$

where $H_{\text{active}} = \bigcup_j C_j$ is the union of nontrivial components and

$$A(G, c) \;:=\; \sum_{j=1}^{k} \max_{h\in C_j} c(h)$$

is the geometry dividend*—the total disclosure cost that witness geometry allows the optimal repair to avoid.*

Proof. Under DFD + CHP, each nontrivial component C_j carries the uniform matroid $U(|C_j|-1,\, |C_j|)$ (Theorem 3). A basis of the full witness matroid selects exactly $|C_j|-1$ elements from each C_j—equivalently, it omits exactly one element per component. To minimize total cost, the optimal basis omits the most expensive element from each component. Hence:

$$\Sigma^*(G,\, c) = \sum_j \Big(\sum_{h \in C_j} c(h) - \max_{h \in C_j} c(h) \Big). \qquad \square$$

Computationally verified by brute-force basis enumeration on all 51 nonzero-fee compositions in the corpus at fee ≤ 14. Zero failures.

Definition 13 (Geometry dividend)**.** *The* geometry dividend $A(G,\, c) = \sum_j \max_{h \in C_j} c(h)$ *is the total disclosure cost that the optimal repair avoids by omitting the most expensive element from each witness component. It is the economic value of witness geometry under cost vector c: the amount of cost that structure lets you not pay.*

Corollary 14 (Robustness as top-two gap)**.** *The optimal repair basis is unique if and only if each component has a unique most expensive element. The robustness margin—the cost perturbation required to change the optimal omission in any component—is:*

$$\text{margin}(G,\, c) = \min_j \big(m_j^{(1)} - m_j^{(2)}\big)$$

where $m_j^{(1)} \geq m_j^{(2)}$ *are the two largest costs in component* C_j*.*

Remark 15 (The four projections of repair)*.* The invariant hierarchy of §2.2 now has a sharper operational reading:

- **fee** $= \sum_j (|C_j| - 1)$ tells you how many obligations exist.
- $\beta = \prod_j |C_j|$ tells you how many structural repairs exist.
- $A(G, c) = \sum_j \max_{h \in C_j} c(h)$ tells you how much the optimal repair saves.
- **margin** $= \min_j (m_j^{(1)} - m_j^{(2)})$ tells you how stable that savings is.

Fee is a sum of component sizes minus one. β is a product of component sizes. Geometry dividend is a sum of componentwise maxima. All three are projections of the same component decomposition, but they compress different information.

6.2 Decision activation as a corollary

The sum-of-maxima theorem makes the activation principle exact, not merely empirical.

Definition 16 (Decision-relevant geometry)**.** *Let* G_0 *and* G_1 *be compositions with* $\text{fee}(G_0) = \text{fee}(G_1)$ *and different witness component motifs. The witness geometry is* decision-relevant under cost vector c *if* $\Sigma^*(G_0, c) \neq \Sigma^*(G_1, c)$*—the optimal repair cost changes as a function of motif at fixed fee.*

Definition 17 (Structural repair multiplicity)**.** *A composition G has* structural repair multiplicity *if its witness Gram* $K(G)$ *has two or more nontrivial connected components. In the DFD + CHP regime with connected carrier graphs, this is equivalent to* $\beta(G) > \text{fee}(G) + 1$*.*

Definition 18 (Cost heterogeneity on the witness support)**.** *A cost vector c is* heterogeneous on the witness support *of G if the per-component maxima* $\{\max_{h \in C_j} c(h)\}$ *are not all proportional to* $|C_j| - 1$*. Concretely: different components contribute different amounts to geometry dividend per unit of fee, so the motif partition affects* Σ^**.*

Proposition 19 (Decision activation). *In the DFD + CHP regime with additive costs, the witness geometry is decision-relevant at fixed fee if and only if:*

(i) the fee-matched compositions have different component partitions, and

(ii) the cost vector is heterogeneous on the witness support (Definition 18).

Under uniform costs ($c(h) = c_0$ for all h), $\Sigma^* = c_0 \cdot \text{fee}$ regardless of motif: geometry is inert.

Proof. By Theorem 12, at fixed fee with components $C_1, \ldots, C_k$:

$$\Sigma^*(G, c) = \sum_{h \in H_{\text{active}}} c(h) - \sum_j \max_{h \in C_j} c(h).$$

The first term $\sum c(h)$ depends on the field set but not on how it is partitioned into components. At fixed fee and fixed hidden field set, two motif partitions of the same fields yield different Σ^* exactly when their sums of componentwise maxima differ, which requires heterogeneous costs.

Under uniform costs $c(h) = c_0$, every component contributes $\max_{h \in C_j} c(h) = c_0$, so $A(G, c) = k \cdot c_0$ depends only on the number of components, not their sizes. But $\Sigma^* = c_0 \cdot \sum(|C_j| - 1) = c_0 \cdot \text{fee}$, independent of partition. □

6.3 The minimal activation diagram

At fee = 3 in the corpus, the two motif classes—(3, 2) and (4)—populate all four quadrants of the activation diagram under the environment-derived cost model:

	Motif	Mode	Composition	Bases	Costs	Outcome
Q1	(4)	rigid	notion+ns-mcp	4	all 9.0	nothing to decide
Q2	(3, 2)	flexible	notion+todoist	6	all 9.0	indifferent
Q3	(4)	rigid	xmind+playwright	4	17–21	no routing freedom
Q4	(3, 2)	flexible	xmind+notion	6	13–17	**routes to 13**

Q1 (rigid + uniform): Single component, all fields at sensitivity 2. Four bases, all cost 9.
Q2 (flexible + uniform): Two components, but all fields at sensitivity 2. Six bases, all cost 9. Geometry exists but the cost model assigns equal cost to every basis: $A = 3 + 3 = 6$ in either partition.
Q3 (rigid + heterogeneous): Single component with cheap (`directory`, cost 3) and expensive (`path/filePath`, cost 7) fields. Geometry dividend $A = 7$ (one max). No routing freedom.
Q4 (flexible + heterogeneous): Two components with different cost profiles. Geometry dividend $A = 7+3 = 10$, versus the single-component $A = 7$. The extra component contributes +3 of geometry dividend, reducing Σ^* from 17 to 13—a 23.5% reduction at identical fee.

Replication. Replacing notion with ns-mcp-server reproduces Q4 exactly: fee 3, $\beta = 6$, motif (3, 2), cost range [13, 17], $A = 10$.

6.4 Necessity and sufficiency

By Theorem 12 and Proposition 19:

- **Structural multiplicity is necessary but not sufficient.** Q2 has multiplicity ($\beta = 6$) but uniform costs make Σ^* partition-independent.

- **Cost heterogeneity is necessary but not sufficient.** Q3 has heterogeneous costs (cost spread 4 within the component) but only one component, so the partition cannot differ.
- **Together, they are sufficient.** Q4 has both, and $\Sigma^*(Q4) < \Sigma^*(Q3)$ at fixed fee.

Remark 20 (Separating exogenous costs from endogenous geometry)*.* The sum-of-maxima theorem (Theorem 12) holds for *any* additive cost vector—it is a structural fact about the uniform matroid, not a property of a particular cost model. The environment-derived cost model used in the corpus verification is one instance: it assigns costs from field-name sensitivity and cross-server exposure, quantities that are exogenous to the witness geometry. Leverage scores, which are endogenous to $K(G)$, enter the cost model as a secondary weight. To verify that activation is not an artifact of this coupling, we confirm that activation persists under the purely exogenous cost assignment $c(h) = \text{field\_sensitivity}(h)$ alone (without leverage): $\Sigma^* = 10$ in Q4 versus $\Sigma^* = 14$ in Q3, a 28.6% separation.

Remark 21 (Scope and boundary of the exact regime)*.* Proposition 19 and Theorem 12 are proved for the DFD + CHP regime, where each witness component carries a uniform matroid and every element is omittable. Outside this regime, two things can go wrong, and they push in opposite directions:

- **Coloops** (elements in every basis) make the formula *underestimate* Σ^*: the singleton coloop is ignored by the componentwise computation, but its cost is unavoidable.
- **Non-uniform essential structure** (components with rank deficit > 1) can make the formula *overestimate* Σ^*: the optimal basis omits more than one element per component, so the actual savings exceed the single per-component maximum.

These errors can partially cancel. The general exact decomposition is $\Sigma^* = B_{\text{forced}} + \Sigma^*_{\text{ess}}$, where B_{forced} is the sum of coloop costs and Σ^*_{ess} is the minimum-cost basis of the essential matroid (coloops contracted, loops deleted). The geometry dividend formula is exact when the essential matroid is a product of uniform matroids—precisely the DFD + CHP regime. Outside it, the formula is a surrogate induced by the uniform-product approximation, not a universal bound in either direction.

The entire current corpus is DFD (product formula matches brute-force enumeration on all 373 testable compositions), so the exact regime covers the observed operational landscape. See the companion boundary note for synthetic stress tests and the full three-level hierarchy.

7 Discussion

7.1 Four projections of one object

The witness Gram $K(G)$ is a single algebraic object. Its four increasingly operational projections are:

- **fee** $= \text{rank}(K) = \sum(|C_j| - 1)$—obligation count. How many hidden conventions?
- $\beta = \prod |C_j|$—structural flexibility. How many minimum repairs exist intrinsically?
- $A(G, c) = \sum \max_{C_j} c(h)$—geometry dividend. How much does the optimal repair save?
- **margin** $= \min_j(m_j^{(1)} - m_j^{(2)})$—robustness. How stable is that savings?

Fee is a sum of component sizes minus one. β is a product of component sizes. Geometry dividend is a sum of componentwise maxima. All four are projections of the same component decomposition; each compresses different information from K.

7.2 Why is the corpus so rigid?

The corpus's overwhelming rigidity (81.7% at $F = 0$) has a structural explanation. Most MCP server pairs share conventions through a single dominant semantic dimension. This produces a single large carrier graph component per dimension, giving $\beta = n_d$ and F near zero. Multi-component decompositions arise only when a server pair shares conventions across several independent dimensions—which requires richer semantic overlap than most pairs exhibit.

This is not a weakness of the theory. It means the real control-plane question is often not "which of many equivalent repairs should I choose?" but "do I have any real freedom at all?" The default operational regime is rigidity.

7.3 Relationship to Paper III [3]

This paper corrects Paper III [3] Corollary 4.3 (basis count via spanning trees) and replaces it with the uniform-matroid product formula $\beta = \prod |C_{d,j}|$. The correction is confined to basis counting; all other Paper III results (rank identity, Kron reduction, leverage, effective resistance) are unaffected.

The correction simplifies the theory: the uniform matroid $U(n{-}1, n)$ is simpler than the graphic matroid, the product formula $\beta = \prod n_d$ is simpler than $\beta = \prod \tau(\tilde{G}_d)$, and the sharp bounds (Theorem 8) follow from elementary integer-partition arithmetic rather than spanning-tree combinatorics.

7.4 Limits and frontier

- **Regime.** The separation theorem (Theorem 3), product formula, sharp bounds, and exact repair functional (Theorem 12) work in the DFD + CHP regime, where each witness component carries a uniform matroid. Outside this regime, the formula can fail in either direction: coloops cause underestimation, non-uniform essential structure causes overestimation (§6). The general exact decomposition $\Sigma^* = B_{\text{forced}} + \Sigma^*_{\text{ess}}$ holds universally. K-sufficiency (Theorem 1) and full realizability (Theorem 10) hold in the general linear disclosure model without regime restrictions. The entire current corpus is DFD.

- **Cost family.** The stability computation uses a specific cost family (integer costs in $[1, 10]$). Operational sparsity is relative to this family; a broader or narrower family would produce different stability ratios. A full characterization would require the normal fan of the matroid basis polytope—the partition of cost space into regions where each basis is optimal.

- **Corpus support.** The current corpus has 24 servers producing 240 nonzero-fee pairwise compositions with only 19 distinct motifs. The high-entropy regime ($\beta > 100$) is populated almost entirely by `github`-paired compositions. Deliberate construction of compositions that stress underrepresented regions of the repair-flexibility landscape is needed to test whether the patterns generalize.

- **Inversions.** The entropy-stability inversion (§5.5) shows motif geometry carries information beyond β. Characterizing $\beta_{\mathcal{F}}$ for structured cost families, and understanding which motif features predict operational stability, is the natural next frontier.

8 Conclusion

In the linear disclosure model under DFD + CHP, the witness Gram $K(G)$ is a single object with three increasingly operational projections. Fee counts hidden obligations. Repair entropy counts structurally valid repairs. The stability ratio under a cost family identifies how many of those repairs are decision-relevant.

The central empirical finding is that, under integer costs in $[1, 10]$, repair spaces are combinatorially large but operationally sparse. At $\beta = 378$, only 13.5% of bases are ever optimal. Theorem 10 proves this sparsity is not intrinsic to the matroid—under unrestricted costs, every basis is realizable—but depends on the cost family. The compression from structural to operational flexibility is the decision-relevant quantity.

Six results anchor this:

1. **K-sufficiency**—in the linear disclosure model, the witness Gram determines the exact repair basis structure.

2. **Fee–multiplicity separation**—fixed fee does not determine repair entropy (DFD + CHP regime).

3. **Sharp bounds**—$\varphi + 1 \leq \beta \leq 2^{\varphi}$, both attained (DFD + CHP regime).

4. **Full realizability**—every basis is the unique optimum under some cost vector (general).

5. **Exact repair functional**—minimum repair cost equals total active hidden cost minus the sum of per-component maxima (DFD + CHP regime). The *geometry dividend* $A(G, c) = \sum_j \max_{h \in C_j} c(h)$ is the operational value of witness geometry.

6. **Decision activation**—witness geometry changes the optimal repair if and only if structural multiplicity and cost heterogeneity on the witness support coincide. This follows algebraically from the exact repair functional, not merely from corpus observation.

The corpus validates the theory (β formula = β enumerated on all 239 testable compositions; sum-of-maxima formula = brute-force minimum on all 51 at fee ≤ 14) and reveals the dominant regime: 81.7% of compositions are structurally rigid (minimum β at their fee). When flexibility exists, the exact repair functional (§6) identifies precisely when that flexibility changes action: structural freedom alone is not sufficient (uniform costs make Σ^* partition-independent), and cost heterogeneity alone is not sufficient (a single-component motif offers no routing freedom). The recommendation changes only when both ingredients coincide. In the DFD regime, the geometry dividend is the exact operational value of witness geometry—recovered as a special case of the universal decomposition $\Sigma^* = B_{\text{forced}} + \Sigma^*_{\text{ess}}$, where the forced cost vanishes and the essential matroid is uniform-product. Outside DFD, the formula is a surrogate that can fail in either direction: coloops cause underestimation, non-uniform essential structure causes overestimation. The exact decomposition, not the formula, is the universal invariant. The witness Gram sees all of this; the scalar fee sees none of it.

References

[1] J. Komkov, "The Coherence Fee," 2026.

[2] J. Komkov, "The Column-Matroid Backbone," 2026.

[3] J. Komkov, "The Witness Gram as a Kron-Reduced Laplacian," 2026.

[4] J. Komkov, "Local Validity Does Not Compose," 2026.

[5] J. Edmonds, "Matroids and the greedy algorithm," *Math. Programming* **1**, 127–136, 1971.

[6] R. Rado, "Note on independence functions," *Proc. London Math. Soc.* **7**, 300–320, 1957.