TYK2 — V4 Whitespace & Trajectory Analysis

Stage E.7 of the V4 Target Opportunity Report playbook. Companion to `TYK2_IP_Portfolio_Analysis_V4_native.md` (Stage C, the KC read). This document carries the scope-coordinate work: the mechanism × direction grid, the cross-examination of its zeros, the filing trajectory, and the pan-JAK fence that bounds what any of it can mean.

Revision 1.1 (post-adversarial-review ×1). Every number in a table or attached to a label below is computed by `scripts/figures.py` and injected by `scripts/render_report.py`. That claim has a measured limit rather than being absolute: 397 numeric tokens across all four templates and their shared fragments are still hand-authored (mostly section numbers and thresholds), and the renderer's stale-number detector is **structurally blind to small integers** — with ~914 figures computed, every digit 0–9 coincides with some figure value, so a hand-typed small number matches by luck. Small load-bearing numbers are defended by explicit binding checks in `scripts/verify_tyk2_report.py`, not by the detector. Regenerate with `python3 scripts/figures.py && python3 scripts/render_report.py`, then verify with `python3 scripts/verify_tyk2_report.py`.

> **Public edition.** This is a trimmed version of an internal analyst deliverable. The uncalibrated claim-architecture (KC) scoring layer, the knowledge graph, the per-family scored data and the internal audit records are **withheld** — not because they are unfavourable but because an uncalibrated per-patent score reads as precision that does not exist, and because the machine-readable estate map is the method rather than the finding. What is published here is the **structural layer**: where claim coverage is present, absent or thin, and where filings are moving. Every guardrail from the internal edition is reproduced below verbatim.

0. Read this first — the guardrails, once

These bind every sentence that follows. They are not boilerplate; each one changes what a specific claim below is allowed to mean.

1. **This is candidate generation, NOT freedom-to-operate.** An OPEN cell is a place to look, never a place shown to be clear. Nothing here is a clearance opinion, and no cell verdict survives contact with counsel unassisted.
2. **An empty cell is a measurement, and the instrument is part of it.** A zero can mean *nobody filed here*, or *we did not search here*, or *the pairing is incoherent so nobody could file here*. Section 2.3 separates these, because conflating them is how whitespace analysis manufactures opportunity.
3. **Genus recall is a LOWER BOUND.** 5 of 23 coherent cells have any recall figure at all. A cell's "breadth" is the breadth we could compile, not the breadth that was claimed.
4. **The pan-JAK fence is measured but NOT ingested.** a **lower bound of 765 pan-JAK discovery rows**. Read that as rows, not families: 519 of the 765 carry a distinct title, and the ids are a **mixed namespace** — 721 numeric INPADOC family ids plus 44 PCT application dockets, which are not families at all. The true family count is therefore lower than 765 and is **UNMEASURED**, so no rows-to-families ratio against this 279-family corpus is like-for-like. Titles are not families either, so **no ratio here is like-for-like** and none should be quoted as the fence's size relative to the corpus. It is also a lower bound in the other direction: 5 of the 46 slices across both passes are recorded saturated at the 100-row cap, the first pass (16 slices) recorded no saturation field at all, and 5 legs in pass 2 plus 8 in pass 1 returned an upstream error and were skipped — so the fence

is larger than 765 by an unknown margin. Because those rows were screened at title level only, the three reportable zeros in §2.4 are **NOT CONTRADICTED**, which is weaker than *unfenced*.

5. **No claim-quality or scoring layer appears in this edition at all.** Nothing below ranks cells by the quality of anyone's claims; §4 describes *who* is filing, not how well.

The fence layer is NOT merged into the on-pathway set, and no fence row was ingested for this analysis. It is **not** true, as earlier revisions of this text said, that no fence family appears in any KC or genus number: **7 of the fence rows are already corpus families** (67479598, 78725685, 80247730, 82260225, 85602036, 90573424, 93794614) and **3 of those sit inside the KC basis** (67479598, 80247730, 90573424). They are there because TYK2-anchored discovery found them first — they are TYK2 families that also answer a pan-JAK query — not because the fence layer imported them.

---

1. The headline

The TYK2 landscape is **not open at the mechanism level, and the interesting question is not where the art is thin but why it is thin.**

- Of 60 grid cells, only **23 are coherent** — 38.3% . The other 37 pair a mechanism node with a direction that cannot be filed against it. Their emptiness measures the **ontology**, not the landscape, and reporting them as whitespace would be the single largest available error in this analysis.
- Of those 23 coherent cells, **7 carry any art at all** and 16 are empty. But "empty and coherent" is still not "TYK2 opportunity": only 2 of the 12 OPEN cells sit on the TYK2 protein rather than on a downstream STAT, an upstream receptor, or solid-state process art (§2.2).
- Art is extraordinarily concentrated. The top two cells — `JH1_catalytic|orthosteric_inhibitor` (n=92) and `JH2_pseudokinase|allosteric_inhibitor` (n=89) — hold 79.7% of all 227 coherently-binned families between them.
- After cross-examination, **3 of 14** probed zeros survive — but they are **not three opportunities**. Only **1** (`JH2_pseudokinase|activator_agonist`) is a coherent cell the corpus genuinely left empty. The other **2** — `JH2_JH1_interface|allosteric_inhibitor`, `FERM_SH2_receptor_box|allosteric_inhibitor` — are **ARTEFACT_ZERO**: their whole mechanism row is unpopulated corpus-wide, so the zero measures our binner, not the landscape. See §2.4.
- The trajectory is a **selectivity pivot, not a growth story**: median genus recall ran 0.288 → 0.643 → 0.413 → 0.279, peaking in **2013-2018** at 0.643 and then falling in each of the 2 eras since to 0.279 — 43.4% of peak. It is *not* a monotonic decline across all four eras (2 of 3 steps fall); it is a rise into the deucravacitinib era and a sustained narrowing after it.

One-line summary: **TYK2 allostery is heavily crowded and the remaining coherent gaps are mostly not on the TYK2 protein at all; the largest single uncertainty is a pan-JAK fence we sized but did not read.** "Crowded" is what patent density supports; it is not evidence that the biology is solved, and this document makes no claim about that.

---

2. The coherent grid

2.1 The coordinate system

Rows are **10 mechanism_node** values; columns are **6 direction** values; the product is 60 cells. The axes are structural, not commercial — they describe where on the protein an intervention acts and what it does when it gets there, so that a cell is a *mechanism claim-space* rather than a market segment.

A **coherence mask** declares which pairings are meaningful: 23 pairs across 10 nodes. The mask is authored, not derived — it encodes a pharmacological judgement, and it is the most contestable object in this analysis. §2.5 reports where it is already known to be wrong.

Two thresholds are fixed in advance: a cell is **thin** at n ≤ 2, and a genus is **broad** at recall ≥ 0.5.

2.2 What is populated

cell (mechanism node | direction)	n	verdict	live granted	legal unreadable	broad+live genera	median recall
`JH1_catalytic	orthosteric_inhibitor`	92	COVERED	26	30	10
`JH2_pseudokinase	allosteric_inhibitor`	89	COVERED	24	31	13
`solid_state_process	allosteric_inhibitor`	30	COVERED	3	10	1
`TYK2_biomarker_dx	diagnostic`	6	FENCE_UNMEASURED	0	5	UNMEASURED
`solid_state_process	orthosteric_inhibitor`	6	COVERED	2	3	1
`TYK2_protein_level	degrader`	3	PARTIAL	0	0	0
`TYK2_transcript	expression_knockdown`	1	THIN	0	0	UNMEASURED

Verdict vocabulary, derived from the grid rather than assumed: ARTEFACT_ZERO, COVERED, FENCE_UNMEASURED, OPEN, PARTIAL, THIN (6 kinds). Distribution across the 23 coherent cells: COVERED 4, OPEN 12, ARTEFACT_ZERO 4, PARTIAL 1, THIN 1, FENCE_UNMEASURED 1.

The OPEN count is the most over-readable number in this document, so decompose it before using it. Of the 12 OPEN cells, only **2 sit on the TYK2 protein itself**; 10 do not. By node: `downstream_STAT` 4, `solid_state_process` 3, `upstream_cytokine_receptor` 3, `JH1_catalytic` 1, `JH2_pseudokinase` 1. `downstream_STAT` and `upstream_cytokine_receptor` are STAT1/3/4 and IL-23R/IFNAR1 — **different proteins**, so an opening there is not TYK2 whitespace — and 3 are on `solid_state_process`, which the ontology itself describes as not a mechanism (a polymorph or process claim on an antisense construct is a coherent filing and an incoherent *opportunity*). A strict reverse audit on "is this actually TYK2 mechanism space" would take 12 OPEN down to 2, and cross-examination has already killed one of those two. That audit has not been run systematically — see §2.5.

The concentration is the finding. `JH1_catalytic|orthosteric_inhibitor` and `JH2_pseudokinase|allosteric_inhibitor` are the JAK-family orthosteric lane and the TYK2-selective allosteric lane respectively — the two places the entire field has been working — and they carry 26 and 24 live-granted families. Third place, `solid_state_process|allosteric_inhibitor` (n=30), is solid-

state and process chemistry: not a mechanism at all, but formulation and salt/polymorph art accreting around molecules that already exist. Its presence at rank 3 is a maturity signal.

Below that the counts collapse to single digits — `TYK2_biomarker_dx|diagnostic` (n=6), `solid_state_process|orthosteric_inhibitor` (n=6), `TYK2_protein_level|degrader` (n=3) — which is why no ranking below rank 3 should be read as a trend.

2.3 Why most zeros are not whitespace

17 coherent cells sit at or below the thin threshold of n ≤ 2. That is the **candidate** set, and it is the number a careless version of this analysis would have published as the opportunity.

Two bookkeeping corrections, both of which an earlier revision got wrong:

- `data/whitespace_thin_cells.json` holds 18 rows, not 17. The extra row is `TYK2_biomarker_dx|diagnostic` (n=6, FENCE_UNMEASURED) — above the thin threshold, and in that file only because of its verdict. Publishing the file's row count as "cells at or below the thin threshold" was wrong; the count is now derived from the grid.
- **Not every candidate was cross-examined.** 14 of the 17 were probed; **3 were never examined at all** — `FERM_SH2_receptor_box|activator_agonist`, `JH2_JH1_interface|activator_agonist`, `TYK2_transcript|expression_knockdown`. Those are recorded as *unexamined*, which is not the same thing as empty, and they are excluded from every rate below rather than counted as survivors.

The 14 probed cells were cross-examined against four questions:

1. Is the cell occupied by art we failed to retrieve? → **externally occupied** (note: "occupied" here means the probe returned hits. Two of the 4 rest on a single hit each, and one of those — a liquiritin-nanoparticle psoriasis filing — is adjudicated *semantic noise, not mechanism art* in `panjak_fence_findings.json`. The direction of that error is conservative: it removes candidate whitespace rather than creating it.)
2. Did we ever actually search for it? → **not searched**
3. Does it sit downstream of a cell that is itself thin, so its emptiness is inherited rather than independent? → **downstream of thin**
4. Does it survive all three? → **reportable**

Of 14 probed zeros: **4 returned external hits**, **4 were never searched**, **3 are downstream of a thin cell**, leaving **3**. That is a 78.6% attrition rate on the probed set — the only population every member of which was actually adjudicated. Read against the full 17-cell candidate set, one coherent zero survives as a genuine empty cell and 3 cells remain unexamined.

Separately, 4 cells are classed **ARTEFACT_ZERO**: their emptiness is a property of the retrieval instrument and is reported as a measurement gap, never as opportunity.

2.4 The three cells that survive

`JH2_pseudokinase|activator_agonist`, `JH2_JH1_interface|allosteric_inhibitor`, `FERM_SH2_receptor_box|allosteric_inhibitor`. Those are what survived cross-examination of 14 probed zero cells: 4 proved externally occupied, 4 were never searched, and 3 sit downstream of a thin cell — so 3 of 14 is the reportable remainder, not a count of empty space.

Their grid verdicts are not the same, and the difference is decisive:
`JH2_pseudokinase|activator_agonist` (OPEN); `JH2_JH1_interface|allosteric_inhibitor` (ARTEFACT_ZERO); `FERM_SH2_receptor_box|allosteric_inhibitor` (ARTEFACT_ZERO).

`ontology.json` defines ARTEFACT_ZERO as a zero that "measures the instrument, not the landscape", and §2.3 of this document says such cells are reported as a measurement gap **never as opportunity**. Both ARTEFACT_ZERO cells here have `row_unpopulated_corpus_wide: true` — across all 279 corpus families, on- and off-pathway, the binner assigned **nothing** to either mechanism node. So the correct reading is:

- **1 genuine coherent zero** — `JH2_pseudokinase|activator_agonist`.
- **2 instrument artefacts** — `JH2_JH1_interface|allosteric_inhibitor`, `FERM_SH2_receptor_box|allosteric_inhibitor` — which tell you the binner has no vocabulary for those rows, and tell you **nothing** about whether the art exists. Treating them as opportunity would be this document's worst available error, and an earlier revision did exactly that by publishing all 3 undifferentiated.

What each one actually is:

- **`JH2_pseudokinase|activator_agonist` — OPEN**, the only one of the three that is a genuine coherent zero in a row the corpus does populate. A TYK2 *activator* binding the regulatory pseudokinase domain. The entire field is inhibitors; agonism is the mirror image. Of 519 pan-JAK fence titles, **0 name TYK2** and the single vocabulary match was on the word *"activator"* — zero titles contain *"agonist"*. Thin evidence of absence, and thin is the operative word.
- **`JH2_JH1_interface|allosteric_inhibitor` — ARTEFACT_ZERO.** Targeting the interface *between* the regulatory and catalytic domains rather than either domain alone. Mechanistically credible and adjacent to the most crowded cell on the board — but the entire `JH2_JH1_interface` row is empty corpus-wide, so this cell reports that our binner never classified anything to that node. It is a vocabulary gap first and a hypothesis second.
- **`FERM_SH2_receptor_box|allosteric_inhibitor` — ARTEFACT_ZERO.** Intervening where TYK2 docks to its cytokine receptor, upstream of the kinase domains entirely. The most differentiated of the three and the least validated: same caveat, the whole row is unpopulated corpus-wide.

All three carry a fence caveat on top of the above. Of the 5 allosteric probe legs, 4 recorded an upstream error on a sub-leg while still returning a hit, so the probe is substantially — not entirely — unmeasured. It is **not** "4 of 5 legs returned nothing", which an earlier revision implied. These cells are **NOT CONTRADICTED**: not shown to be occupied, and equally not shown to be open. For the two ARTEFACT_ZERO cells the fence is not even the binding constraint; the binner is.

2.5 Where the mask is already wrong

Absence-first requires reporting this rather than absorbing it. **1 family sits in a cell the mask calls incoherent**: `TYK2_protein_level|expression_knockdown` (n=1). This is why the coherent cells sum to 227 while §3 of the KC report counts 228 on-pathway families — the difference is exactly 1.

The family is `80628904` (UNIV OXFORD INNOVATION LTD) — a gene-modified cell therapy. Knocking out TYK2 by editing does act at the protein level, but it gets there through the transcript, so the mask's refusal to pair `TYK2_protein_level` with `expression_knockdown` is too narrow rather than the binning being wrong.

The consequence for this document: the mask is an authored artefact with at least one known defect, so **every cell verdict inherits an unquantified mask risk**. A mask error in the other direction — wrongly calling a pairing *coherent* — would **manufacture** whitespace, and no systematic audit for that class has been run.

We can, however, bound it partially, and the bound is unflattering: §2.2 shows 10 of 12 OPEN cells are not on the TYK2 protein at all. Those are exactly the wrongly-coherent class — pairings that are *filable* but

are not TYK2 mechanism space. So the 16 empty coherent cells and the OPEN 12 are both **upper bounds on anything a reader should treat as TYK2 opportunity**, and the honest floor is 2 before cross-examination and 1 after it.

2.6 Breadth, where it could be measured

5 of 23 coherent cells carry a genus recall figure. The highest cell-level median is 0.421. Across the grid, 25 distinct live families hold a genus at or above the 0.5 breadth threshold.

Those 25 families are the practical fence: broad, granted, alive. They are enumerated per-cell in `data/whitespace_grid_node_dir.json` under `broad_live_family_ids`. Recall is a lower bound, so this understates them.

---

3. Trajectory — the selectivity pivot

3.1 Volume plateaued; breadth collapsed

era	families	share	median genus recall	genera compiled
pre-2013	16	7.0%	0.288	12
2013-2018	27	11.8%	0.643	20
2019-2022	95	41.7%	0.413	56
2023-2026	90	39.5%	0.279	24
total	**228**			

Filing volume rose steeply (16 → 27 → 95), peaking in 2019-2022 at 95 families, and then **plateaued or fell** (90) — and the last era is under-counted by pendency, so it will grow. Volume alone says "mature".

Median genus recall is the more informative series: **0.288 → 0.643 → 0.413 → 0.279**. It peaks in 2013-2018 and falls in every step after the peak, ending at 43.4% of it. The precise claim is **narrowing since 2013-2018**, over 2 eras — not a decade-long monotonic decline, since breadth *rose* into the peak.

Read together: a maturing field where the broad genus positions were taken early, and everything since has been differentiation inside them.

A confound that may explain the whole slope. Breadth is measured only on families whose genus compiled, and that coverage rate collapses across exactly the same eras:

era	families	genera compiled	coverage
pre-2013	16	12	75.0%
2013-2018	27	20	74.1%
2019-2022	95	56	58.9%
2023-2026	90	24	26.7%

Coverage falls from 75.0% to 26.7%. Younger families are systematically less likely to have a compiled genus, so the newest era's breadth rests on a 26.7% sample of it. That is a **direction-carrying selection effect and a live alternative explanation for the narrowing**, and its magnitude is **UNMEASURED**. An earlier revision called the slope "robust" and the denominators "small and shrinking over time"; they are neither uniformly small nor monotonically shrinking (12 → 20 → 56 → 24). Both statements are withdrawn. The slope is **consistent with** narrowing and **not separable** from declining genus-compilation coverage on this data.

3.2 The orthosteric → allosteric inversion

Within TYK2-labelled art, the direction mix inverts:

era	n	orthosteric inhibitor	allosteric inhibitor	degrader	expression knockdown	diagnostic	median genus recall
pre-2013	16	14	1	0	0	1	0.288
2013-2018	27	8	18	0	0	1	0.643
2019-2022	95	45	43	2	2	3	0.413
2023-2026	90	31	57	1	0	1	0.279

Early art is orthosteric — ATP-competitive inhibitors hitting the catalytic domain, i.e. pan-JAK chemistry that happens to touch TYK2. By the most recent era allosteric filings **outnumber orthosteric 57 to 31**. That is the deucravacitinib effect: JH2-binding selectivity became the paradigm, and the field followed.

This inversion is scoped to TYK2-labelled art only, and that scope is a retraction. An earlier draft presented it as a statement about JAK-family chemistry generally. It is not: the pan-JAK fence excluded 765 discovery rows that are overwhelmingly orthosteric. By volume those excluded rows are dominated by ATP-competitive chemistry, but their era distribution is **UNMEASURED**, so the inversion cannot be tested against them at all — no per-era ordering over undated rows is available, even hedged. The correct claim is narrow — *among patents that name TYK2, the centre of gravity moved to allostery* — and the fence's own era profile is **UNMEASURED**, because discovery rows carry no filing date.

3.3 What the trajectory does not tell you

Degrader and knockdown directions appear at n=1 and below in the newest era. Whether that is a field declining to pursue them or a field that has not yet published on them is **not resolvable from filing counts**; pendency alone could account for it. Recorded as indeterminate.

---

4. Who is filing, and where

This section describes participants. It does not rank cells, and it carries no claim-quality scoring of any kind.

The two crowded cells are crowded by different populations. The orthosteric cell is legacy pharma chemistry with long-granted estates; the allosteric cell is the post-deucravacitinib cohort. Which of them holds the structurally broadest position is a claim-architecture question, and that layer is withheld from this edition.

Legal state, across the whole on-pathway set: 149 of 228 families have readable legal status and **55 are live-granted**. The other 79 are **unreadable, not "not live"**.

Loss events need their arithmetic shown, because the raw figure is unusable: **18 raw -> 12 on-pathway -> 10 de-duplicated -> 8 after excluding country-prefixed EP cease codes -> 4 distinct losses**. Specifically, 6 of the 18 raw events belong to 1 off-pathway family; 2 are exact duplicate rows; and 2 are country-prefixed EP cease codes ("GB: EUROPEAN PATENT CEASED…") — territorial lapses the exclusion filter missed because it matched only the literal string `LAPSED IN A CONTRACTING STATE`. An earlier

revision published 18 events over 5 families *in the same sentence that claimed territorial over-counting had been fixed*. It had not.

The defensible figures are **8 events over 4 families**, and because one lapse is reported under several codes on the same date (LAPS + STCH + FP) those correspond to only **4 distinct losses**. Separately, 642 territorial EP lapses were excluded by the original filter.

---

5. Structural adjacency

Scope first, because the denominators differ. The structural layer holds 230 rows — the families that were on-pathway when it was harvested — of which 2 were later reversed off-pathway and are excluded from every figure in this section. So the figures below are scoped to the current 228 on-pathway families, not to 230.

Of those, 119 have a compiled `family_best` genus and 112 a recall figure (median 0.391). The layer holds **17016 genus-overlap edges** and 345 sequence edges.

Genus edges are the mechanism by which a family in one cell *could* reach chemistry in another. Two honest qualifications, both of which an earlier revision elided:

- The mesh is **concentrated, not uniform**: 17016 edges come from 20 of the 119 compiled genera, and 99 families have **zero** edges with `neighbors_state: measured` — genuine measured zeros, not gaps.
- **No cell-to-cell adjacency was ever computed.** Edge hosts are arbitrary corpus publications, not grid cells. So the intuition that an empty cell beside a crowded one is more suspect than an isolated one is *reasoning*, not measurement — and it cannot be applied to `JH2_JH1_interface` at all, which has n=0 and therefore no edges.

Sequence adjacency is thin here (345 edges) and unsurprising — TYK2 is a small-molecule target and the corpus is small-molecule art.

---

6. Knowledge graph

809 nodes and 2483 edges linking families, assignees, mechanism nodes, directions and modalities. It is a navigation aid built from the same substrate as the grid, not an independent source of evidence — anything it shows is already in `families_scoped.json`.

---

7. Method, and what a reader must not do with this

How the grid was built. Families from the Stage-D scoped set were binned to a `(mechanism_node, direction)` pair by claim and title reading; the coherence mask was authored from the pharmacology; cell statistics were then computed mechanically. Thin cells were enumerated by threshold, then each probed zero was cross-examined against external search before any of it was called whitespace.

Do not read an OPEN cell as clear. 3 cells survived cross-examination; all three remain NOT CONTRADICTED rather than unfenced, because 4 of 5 probe legs did not execute.

Do not treat 16 empty coherent cells as 16 opportunities. 78.6% of apparent whitespace did not survive examination.

Do not port the coherence mask to another target. It is TYK2 pharmacology, it is authored, and §2.5 shows it already has one known defect.

Do not use the era series as a forecast. The newest era is incomplete by pendency and its breadth figure rests on 24 compiled genera.

7.1 Open measurement gaps

13 hard ingest failures (5 `gp_404_family_resolution`, 8 `bigquery_snapshot_lag_no_rows`), of which **8 are individually named in the ledger** and 5 are counted but not itemised there — `US-12655108-B2`, `US-RE47929-B1`, `WO-2026096291-A1`, `WO-2026147992-A1`, `WO-2026156128-A1`

Specific to this document:

- The pan-JAK fence — 765 rows — is screened at **title level only**. Ingesting it so the genera compile is the single highest-value next action, and until then §2.4's cells cannot be upgraded past NOT CONTRADICTED.
- **1 stored verdict is absent from the published rule** in `ontology.json`: FENCE_UNMEASURED. The rule as applied lives in the analysis script, so `ontology.json` documents 6 verdicts' worth of grid using 5 published definitions.
- **The reverse mask audit has not been run.** §2.2 bounds it partially (10 of 12 OPEN cells are off the TYK2 protein) but no systematic pass exists.
- The fence's **era profile is UNMEASURED**; discovery rows carry no filing date.
- The coherence mask has **no reverse audit**: we found one wrongly-incoherent pairing and did not look for wrongly-*coherent* ones, which is the error direction that manufactures whitespace.
- 5 coherent cells carry a fence-unmeasured flag covering 79 family-slots.
- 14 coherent cells are flagged pan-JAK undercounted.

---

8. Files

file	what it holds
`data/families_scoped.json`	the binned family set, one row per family
`data/whitespace_grid_node_dir.json`	the 60-cell grid, per-cell statistics and `broad_live_family_ids`
`data/ontology.json`	axes, the 23-pair coherence mask, cell-verdict rule
`data/whitespace_thin_cells.json`	the 17 candidate thin/empty cells
`data/zero_cross_examination.json`	the 14 probed zeros and their adjudication
`data/dynamic_signals.json`	era binning, direction/node/modality mix per era
`data/structural_layer.json`	genus + sequence adjacency
`data/legal_layer.json`	legal status, challenges, loss events
`data/panjak_fence_{raw,hits,findings}.json`	the 765-row fence and its findings
`data/tyk2_knowledge_graph.json`	809 nodes / 2483 edges

Rendered from `templates/TYK2_V4_Whitespace_and_Trajectory.md.tmpl` by `scripts/render_report.py`; all figures from `scripts/figures.py`. Verified by `scripts/verify_tyk2_report.py`.