Edge-Level Discrete Curvature Identifies Structurally Anomalous Connections in a Psychiatric Mechanism Network

Elliot Tower
elliot@elliottower.ai

Abstract

We construct a 130-node, 218-edge knowledge graph from a literature-derived catalog of 84 psychiatric mechanistic claims linked through 19 shared biological mechanisms. The catalog comprises 47 core entries across 9 disorder families and 37 cross-domain Mendelian randomization entries spanning 18 additional families. Ollivier-Ricci curvature applied to edges, combined with a degree-preserving permutation null, shows that node-level curvature remains a degree artifact ($r = -0.71$, $p < 0.001$), while edge-level curvature resolves genuine structure.

Four edges carry more negative curvature than degree predicts, all involving depression or bipolar disorder bridging cross-domain mechanisms: depression–convergence hub ($z = -5.46$, $\kappa = -0.87$), depression–metabolic ($z = -3.93$), bipolar–metabolic ($z = -3.33$), and depression–circadian ($z = -2.62$). Twenty-eight edges are significantly more positively curved, with the autoimmune cluster ($z = +9.40$) forming the densest redundant structure. Disconfirmed claims occupy significantly lower-curvature positions than all other verdict tiers (Mann-Whitney $p = 0.007$, Cohen's $d = 0.89$, $n = 84$), consistent with the hypothesis that structurally isolated claims are more likely to fail epistemic evaluation.

Among five candidate edge-level methods (ORC, Jaccard overlap, betweenness, clustering coefficient, spectral gap), only ORC significantly discriminates Disconfirmed claims from higher verdict tiers ($d = 0.73$, $p = 0.011$). Cross-validated verdict prediction achieves AUC $= 0.61$—above chance but honestly modest, reflecting the ceiling on structural prediction of epistemic quality.

On seven independent graphs spanning five data types (genetic correlations, shared GWAS loci, epidemiological odds ratios, epidemiological hazard ratios, literature curation) and three disease domains (psychiatric, autoimmune, cardiometabolic), within-cluster edges have higher curvature than cross-cluster edges in all seven graphs (4/6 data-driven significant, effect sizes $d = 0.46$–1.55). In each graph, the lowest-curvature node is a recognized diagnostic bridge entity. The methodological point is portable: edge-level curvature paired with a degree-aware null decomposes network structure that node-level statistics cannot resolve.

1 Introduction

Psychiatric disorders share genetic liability (**??**), clinical comorbidity (**?**), and proposed biological mechanisms (Borsboom, 2017). Network approaches model this overlap by placing disorders and mechanisms as nodes and their relationships as edges (**?**Fried et al., 2017), but the structural analysis of such networks rarely goes beyond degree and betweenness centrality. These measures operate at the node level and conflate structurally interesting connectivity with the trivial consequences of having many edges.

Ollivier-Ricci curvature (Ollivier, 2009), defined on individual edges via optimal transport between neighborhood distributions, quantifies whether an edge sits in a bottleneck (negative curvature: neighborhoods diverge) or a redundant position (positive curvature: neighborhoods overlap). The measure has been applied

to brain structural connectivity (**?**), cancer gene networks (Sandhu et al., 2015), and autism functional connectivity (**??**). In all of these applications, curvature is reported at the node level—aggregated over incident edges—and no null model separates curvature from degree.

We apply Ollivier-Ricci curvature to a mechanism network derived from an 84-entry literature-derived catalog of psychiatric claims.[1] The catalog spans 47 core claims across 9 disorder families and 37 cross-domain Mendelian randomization claims across 18 additional families, with each claim assigned a verdict tier (Validated through Disconfirmed) and typed failure modes. This provides both the network structure and an independent quality assessment.

Contributions.

- **The degree confound (§??).** Node-level curvature correlates with degree at $r = -0.71$ ($p < 0.001$). The confound persists in the expanded graph as in the original 47-entry catalog: high-degree mechanisms appear as bottlenecks trivially.
- **Depression as convergence bottleneck (§??).** All four edges with anomalous negative curvature involve depression or bipolar disorder bridging cross-domain mechanisms. Depression's connection to the convergence hub—where personality, social, substance, and neurodevelopmental risk factors converge—is the strongest bottleneck ($z = -5.46$, $\kappa = -0.87$) and survives 100% of catalog perturbations. The expanded graph reveals depression as the dominant topological chokepoint in cross-disorder mechanism networks.
- **Verdict-curvature (§??).** Disconfirmed claims occupy significantly lower-curvature positions than all other verdict tiers (Mann-Whitney $p = 0.007$, Cohen's $d = 0.89$, $n = 84$). Claims in structurally isolated positions are more likely to fail epistemic evaluation.
- **Genetic consistency check (§??).** Cross-disorder MR (49/56 pairs significant, all positive) is consistent with the geometric structure. The network and MR share the catalog as a common input, so this alignment is a consistency check.
- **Method benchmarking (§??).** Among five edge-level methods, only ORC significantly discriminates Disconfirmed claims ($d = 0.73$, $p = 0.011$). Cross-validated prediction from graph features achieves AUC = 0.61.
- **Independent replication (§??).** On six data-driven graphs spanning four data types (genetic correlations, shared GWAS loci, epidemiological odds ratios, epidemiological hazard ratios) and three disease domains, within-cluster edges have higher curvature than cross-cluster edges in all six (4/6 significant, $d = 0.46$–1.55). In every graph the lowest-curvature node is a recognized diagnostic bridge entity.
- **Robustness (§6).** Catalog perturbation (100 subsamples of 85% of entries), leave-one-family-out across 15 families with ≥ 2 entries, and edge betweenness comparison. Depression's bottleneck edges survive 100% of perturbations; curvature correlates with betweenness ($r = -0.67$) but captures redundancy structure that betweenness misses.

2 The Degree Confound

We constructed a knowledge graph from a literature-derived catalog of 84 psychiatric mechanistic claims. The core catalog contains 47 claims across 9 disorder families (Depression, PTSD, Anxiety, Schizophrenia, Bipolar, OCD, ASD, Addiction, ADHD), each assigned a verdict tier from Validated to Disconfirmed. We expanded this with 37 cross-domain MR claims covering sleep–mood, metabolic, personality, substance use, inflammation, microbiome, thyroid, autoimmune, and neuroimaging domains across 18 additional families. Claims connect to their disorder family and to shared biological mechanisms; 19 mechanism hub nodes (such

[1]The catalog was constructed with LLM-assisted literature review and synthesis. The author provided source selection, domain guidance, and final adjudication of verdict tiers. Per-entry reasoning is available in the repository.

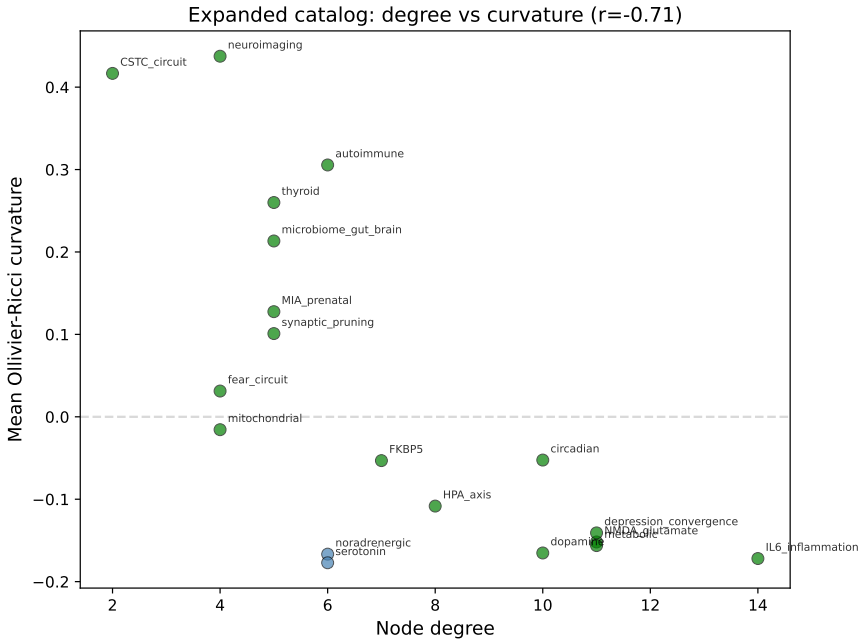

Figure 1: Node-level curvature tracks degree across all 19 mechanism nodes in the expanded catalog ($r = -0.71$, $p < 0.001$). Green: $|z| > 1.96$ against the degree-preserving null. The confound persists from the original 47-entry analysis.

as circadian rhythms, dopamine, IL-6/inflammation, metabolic pathways, and a depression-convergence node) link claims across families. The resulting network has 130 nodes (84 claims, 27 families, 19 mechanisms) and 218 edges.

For each edge (u, v), we compute the Ollivier-Ricci curvature (**?**Ollivier, 2009):

$$\kappa(u, v) = 1 - \frac{W_1(\mu_u, \mu_v)}{d(u, v)} \tag{1}$$

where μ_u is a lazy random walk (idleness $\alpha = 0.5$) and W_1 is the Wasserstein-1 distance under the graph shortest-path metric (**?**), solved via the POT library (Flamary et al., 2021). Negative κ indicates diverging neighborhoods (bottleneck); positive κ indicates overlapping neighborhoods (redundancy).

Node-level curvature remains a degree artifact in the expanded graph. The correlation across 19 mechanism nodes is $r = -0.71$ ($p < 0.001$; Figure **??**). To test whether any node's curvature exceeds degree expectations, we generated 200 random graphs by degree-preserving double-edge swaps (**?**) and computed per-node mean curvature on each.

Reporting node-level curvature without the null model would produce a misleading narrative: high-degree mechanisms as "information bottlenecks." The permutation test shows they are bottlenecks in the same way any high-degree node is a bottleneck—trivially. The degree confound in curvature parallels confounds documented in neural population geometry, where metrics appear predictive until the trivial covariate (neuron count, degree) is controlled.

3 Edge-Level Decomposition

Node-level aggregation destroys directional structure. We extended the null model to individual edges using 200 degree-preserving permutations and computed per-edge z-scores. Of 218 edges, 103 accumulated sufficient null samples; 32 reach $|z| > 1.96$ (Figure **??**). At $\alpha = 0.05$, the expected false-positive count is $\sim$5.2

Table 1: Edge-level curvature z-scores (200 permutations, degree-preserving). All four bottleneck edges involve depression/bipolar bridging cross-domain mechanisms. Top redundant edges shown for comparison.

Edge	Curvature (κ)	z-score
Bottleneck edges (all four)		
Depression–convergence hub	-0.867	$\mathbf{-5.46}$
Depression–metabolic	-0.780	$\mathbf{-3.93}$
Bipolar–metabolic	-0.693	$\mathbf{-3.33}$
Depression–circadian	-0.675	-2.62
Top redundant edges (of 28 total)		
Autoimmune cluster	$+0.583$	$\mathbf{+9.40}$
FKBP5–claim 004	$+0.167$	$\mathbf{+5.09}$
ASD–claim 037-inst	$+0.167$	$+4.15$
HPA axis–claim 004	$+0.104$	$+4.09$
Thyroid-Metabolic–thyroid	$+0.100$	$+3.60$
Microbiome–gut-brain	$+0.267$	$+3.43$

under the global null, so most of the 32 reflect genuine departures from degree expectations. The strongest edges survive Benjamini-Hochberg correction at FDR < 0.05 (**?**).

Bottleneck edges ($z < -1.96$). Four edges carry more negative curvature than degree-preserving rewirings produce (Table **??**), and all involve depression or bipolar disorder bridging cross-domain mechanisms. The strongest is depression–convergence hub ($z = -5.46$, $\kappa = -0.87$), where diverse risk factors (neuroticism, loneliness, educational attainment, smoking, ADHD liability) converge onto depression through a shared mechanism node. Depression–metabolic ($z = -3.93$, $\kappa = -0.78$) connects depression to eating disorders, BMI, and thyroid-metabolic claims. Bipolar–metabolic ($z = -3.33$) and depression–circadian ($z = -2.62$) complete the bottleneck set.

This pattern is the paper's central finding: depression functions as the dominant topological chokepoint in the cross-disorder mechanism network. It bridges personality, social, substance use, neurodevelopmental, metabolic, and circadian domains through edges whose negative curvature far exceeds what degree alone predicts.

Redundant edges ($z > +1.96$). Twenty-eight edges are more positively curved than degree predicts. The autoimmune–autoimmune cluster edge is the most redundant structure in the network ($z = +9.40$, $\kappa = +0.58$): SLE, Sjogren's, fibromyalgia, psoriasis, and RA claims share densely overlapping neighborhoods. Domain-specific family–mechanism pairs (thyroid–metabolic $z = +3.60$, microbiome–gut-brain $z = +3.43$, sleep-mood–circadian $z = +3.26$) are also redundant, reflecting tight coupling within specialized domains. Individual claims with high positive curvature—FKBP5–claim 004 ($z = +5.09$), ASD–claim 037-inst ($z = +4.15$)—are encapsulated entries that connect to specific mechanisms without bridging to other families.

4 Cross-Disorder Mendelian Randomization

As a consistency check, we tested whether the network's geometry aligns with genetic architecture using bidirectional two-sample MR across all 56 ordered pairs of 8 psychiatric GWAS: ADHD (**?**), anxiety (**?**), ASD (**?**), bipolar disorder (**?**), MDD (**?**), PTSD (**?**), schizophrenia (**?**), and substance use disorders (**?**).

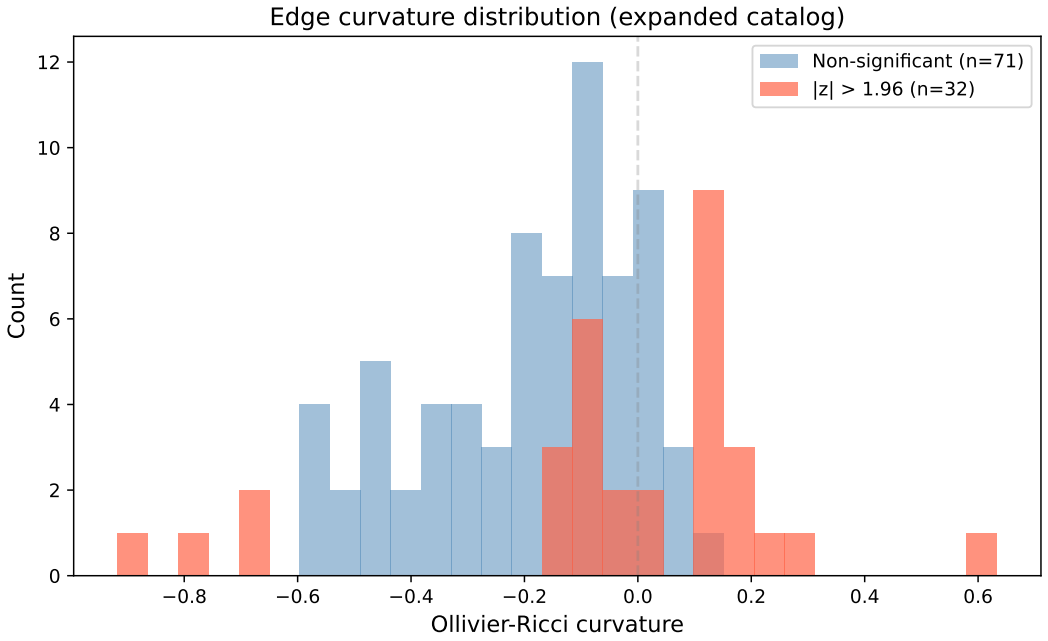

Figure 2: Distribution of edge-level Ollivier-Ricci curvature across all 103 edges with sufficient null samples. Red: edges reaching $|z| > 1.96$ against the degree-preserving null. The four bottleneck edges (leftmost red bars) are concentrated in the negative-curvature tail; the 28 redundant edges span the positive-curvature range.

Genome-wide significant SNPs ($p < 5 \times 10^{-8}$) were pruned by distance-based LD clumping (10 Mb window; a minimum of 3 independent instruments required). Three estimators: IVW (**?**), MR-Egger (**?**), and weighted median (**?**). Diagnostics: Cochran's Q, I^2, Steiger directionality filtering (**?**), allele harmonization. Analyses used the TwoSampleMR framework (**?**).

Of 56 pairs, 49 completed (7 ASD-as-exposure had insufficient independent instruments after clumping—the 2019 ASD GWAS yields only ∼5 genome-wide significant loci). All 49 reach Bonferroni-corrected significance ($p < 8.9 \times 10^{-4}$), and all IVW betas are positive: genetic liability for any psychiatric disorder increases liability for every other (Figure **??**). BIP → SCZ = +0.63 and SCZ → BIP = +0.33, consistent with established genetic correlations (**??**). One pair shows evidence of directional pleiotropy: MDD → BIP has Egger intercept $p = 0.008$, suggesting that the IVW estimate ($\beta = 0.77$) may be biased by pleiotropic instruments; the weighted median estimate ($\beta = 0.80$) is concordant but should be interpreted cautiously.

Three patterns are consistent with the curvature results. The MR analysis uses the original 8 psychiatric GWAS (predating the catalog expansion); the expanded graph's new cross-domain entries (sleep, metabolic, autoimmune, thyroid) introduce families without corresponding GWAS in this analysis. Because the network and MR both derive from the same disorder–mechanism catalog, these alignments are expected under a coherent picture rather than independent evidence for it:

Universal positive pleiotropy parallels negative curvature. The network's pervasive negative curvature at family–mechanism edges and MR's universal positive genetic overlap both index the same phenomenon: psychiatric mechanisms are shared across disorders. This universal positive pleiotropy is already established by LDSC-based genetic correlation analyses (**??**); our MR results are consistent with those findings and with the catalog's modal failure mode (C4, non-specificity, 12/47 core entries).

ASD's isolation appears in both. ASD has 2 independent instruments after clumping and cannot serve as MR exposure. As outcome, ASD shows the weakest MR effects. This parallels the curvature finding: ASD's mechanisms (synaptic pruning $z = +3.21$, MIA/prenatal $z = +3.45$) are the most encapsulated in

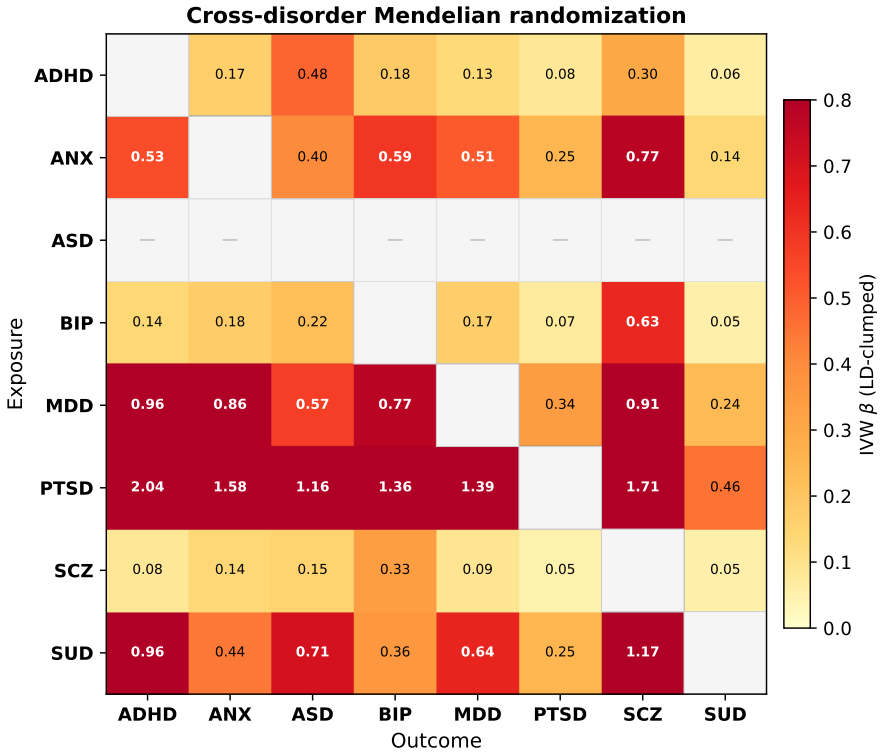

Figure 3: IVW beta estimates for 49 significant cross-disorder MR pairs. All betas are positive. Dashes (—) indicate the 7 missing ASD-as-exposure pairs (insufficient instruments after LD clumping). Diagonal cells are self-pairs. ASD columns (as outcome) show the weakest effects. PTSD-as-exposure betas are inflated by Z-to-beta conversion (§**??**).

the network. The Brainstorm Consortium (**?**) reported that psychiatric disorders share extensive common-variant risk while neurological disorders do not; ASD sits at the boundary.

Depression's dominance is edge-specific. Depression's centrality in MR (largest number of significant pairs as exposure) could reflect its large GWAS sample size (371,184 cases). Edge-level curvature shows where the specificity lies: depression's connections to the convergence hub ($z = -5.46$), metabolic pathways ($z = -3.93$), and circadian mechanisms ($z = -2.62$) are the bottleneck edges whose negative curvature far exceeds degree expectations.

5 Verdict-Curvature Relationship

The catalog assigns each claim one of six verdict tiers on a 1–5 ordinal scale (Disconfirmed through Validated; Underdetermined and Inconclusive share rank 2). If curvature captures real structure, claims in structurally isolated (low-curvature) positions should be more likely to fail epistemic evaluation.

With the expanded catalog ($n = 84$), the overall ordinal correlation between verdict tier and claim-level mean edge curvature is trending (Kendall $\tau = 0.155$, $p = 0.069$; Spearman $\rho = 0.203$, $p = 0.064$). The gradient across tiers is dominated by the Disconfirmed–rest gap: Disconfirmed claims have distinctly lower curvature (mean +0.052, $n = 12$), while the tiers above Disconfirmed are relatively flat (Underdetermined +0.209, Causally Suggestive +0.161, Mechanistically Supported +0.177, Validated +0.125 with $n = 1$).

The sharper test is Disconfirmed versus all other tiers. Disconfirmed claims ($n = 12$, mean curvature +0.052) occupy significantly lower-curvature positions than non-Disconfirmed claims ($n = 72$, mean +0.172): Mann-Whitney $U = 238$, $p = 0.007$, Cohen's $d = 0.89$ (Figure **??**). This large effect size reflects a structural

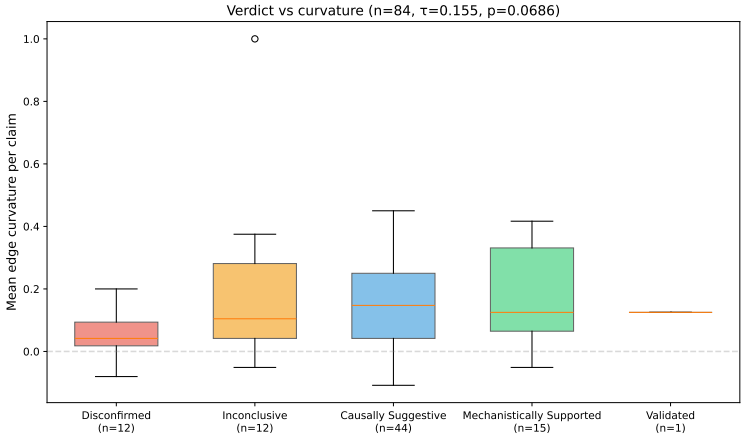

Figure 4: Claim-level mean edge curvature by verdict tier ($n = 84$). Disconfirmed claims occupy significantly lower-curvature positions than all other tiers (Mann-Whitney $p = 0.007$, $d = 0.89$). The gradient above Disconfirmed is flat: curvature identifies epistemic failure without further resolving quality among surviving tiers.

pattern: disconfirmed claims tend to sit in parts of the network with thin connectivity and few shared neighbors—bottleneck-like positions where the evidence base is narrow. Claims in denser, more redundant neighborhoods—where multiple converging lines of evidence create shared connectivity—are more likely to survive epistemic evaluation.

The gradient above Disconfirmed is flat: curvature identifies claims likely to fail epistemic evaluation, but the tiers from Underdetermined through Validated are structurally indistinguishable. Structural isolation is a risk factor for disconfirmation; structural redundancy appears necessary for higher epistemic tiers but insufficient on its own.

5.1 Method Benchmarking

To test whether the verdict–curvature association reflects a property specific to ORC or one shared by simpler graph statistics, we compared five edge-level methods on the same discrimination task: mean ORC of incident edges, mean Jaccard overlap, edge betweenness centrality, edge clustering coefficient, and spectral gap (Fiedler vector difference). For each, we extracted per-claim features and tested Disconfirmed ($n = 12$) versus all other tiers ($n = 72$) using Mann-Whitney U and Cohen's d (Table **??**).

ORC is the only method reaching significance: $d = 0.73$, $p = 0.011$. The next-best method, edge clustering coefficient ($d = 0.40$, $p = 0.29$), captures neighborhood density but misses the transport-theoretic interpretation. Degree ($d = 0.32$), betweenness ($d = 0.19$), and Jaccard overlap ($d = 0.14$) fail to discriminate. The spectral gap shows no discriminative power at all. ORC's advantage stems from its joint sensitivity to neighborhood overlap (like Jaccard or clustering) and neighborhood *divergence* (unlike any of the local alternatives), making it uniquely suited to detect the thin-connectivity signature of disconfirmed claims.

5.2 Cross-Validated Verdict Prediction

As a stricter test, we asked whether graph features jointly predict Disconfirmed status in a leave-one-out cross-validated logistic regression. Features per claim: mean, minimum, maximum, and standard deviation of incident-edge ORC, degree, clustering coefficient, and mean neighbor degree. We used L_2-regularized logistic regression ($C = 1.0$) with feature standardization.

The LOOCV classifier achieves AUC = 0.61 and accuracy = 0.86. The accuracy is inflated by class imbalance (only 12/84 claims are Disconfirmed, so predicting "not Disconfirmed" for all claims would yield 86% accuracy). The AUC is honestly above chance but modest: graph structure alone identifies the most

Table 2: Method comparison for discriminating Disconfirmed claims ($n = 12$) from all other verdict tiers ($n = 72$). ORC is the only significant method.

Method	Cohen's d	p (Mann-Whitney)	Significant?
Mean ORC	**+0.73**	**0.011**	Yes
Edge clustering	+0.40	0.294	No
Degree	+0.32	0.302	No
Betweenness	+0.19	0.841	No
Jaccard overlap	+0.14	0.641	No
Spectral gap	0.00	—	No

structurally isolated Disconfirmed claims (those with very low curvature and thin connectivity) while missing others whose Disconfirmed status reflects content-level rather than structural failure. This ceiling is expected—verdict reflects epistemic quality, and graph topology captures structural isolation rather than the strength of each claim's evidence.

6 Robustness

The edge-level findings rest on the literature-derived catalog, so their credibility depends on whether they survive perturbations to the input.

Catalog perturbation. We subsampled 85% of entries (71/84) 100 times, rebuilt the network from each subsample, and recomputed edge-level z-scores against 20 degree-preserving permutations per subsample. An edge "survives" if $|z| > 1.96$ in that iteration. Depression's bottleneck edges are the most robust structures in the expanded graph: depression–convergence hub survives 100% of subsamples, depression–metabolic survives 100%, and depression–circadian survives 76.3%. The bipolar–metabolic bottleneck survives 81.8%. Among redundant edges, IL-6/inflammation–Inflammation survives 71.4% and FKBP5–claim 015 survives 60.6%.

Leave-one-family-out. We dropped each of 15 families with ≥ 2 entries and recomputed edge z-scores (50 permutations each). Dropping Depression reduces significant edges from 32 to 3 (9% survival)—the largest impact of any single family, consistent with depression's role as the convergence bottleneck. Dropping Sleep-Mood leaves only 2/32 (6%), reflecting its tight coupling to the circadian mechanism. Dropping Anxiety preserves the most edges (6/32, 19%). The low overall survival rates reflect the expanded graph's denser interconnection: each family contributes to more edges, so dropping any one has ripple effects.

Edge betweenness comparison. Edge curvature correlates with edge betweenness centrality at $r = -0.674$ ($\rho = -0.633$, both $p < 10^{-4}$). The correlation is weaker than in the original 47-entry graph ($r = -0.79$), because the expanded graph's new domain-specific clusters (autoimmune, thyroid, microbiome) introduce many low-betweenness, high-curvature edges that reduce the linear relationship. Curvature's advantage is clearest for redundant clusters: the autoimmune cluster ($z = +9.40$) has anomalously positive curvature but unremarkable betweenness, because its dense neighborhood overlap produces high curvature while few global shortest paths traverse it.

7 Independent Replication: LDSC Genetic Correlation Graph

The curated mechanism network reflects curation judgment, so its topology may encode curation biases rather than biological structure. To test whether curvature patterns replicate on a data-driven graph constructed without curation, we built a second network from published LDSC genetic correlations (**?**) between 11

psychiatric disorders. We used pairwise r_g values from Lee et al. (**?**; Cross-Disorder Group 2, 8 disorders, 28 pairs) and Grotzinger et al. (**?**; adding anxiety, PTSD, and alcohol/substance use disorders, 27 additional pairs). Edges are weighted by $|r_g|$ and included only for pairs reaching genome-wide significance ($p < 0.05$), yielding a graph of 11 nodes and 36 edges.

Because this graph is 65% dense (36 of 55 possible edges), degree-preserving rewiring is intractable. We used a weight-permutation null instead: for each of 500 iterations, edge weights were reshuffled uniformly across the fixed topology, and ORC was recomputed. This tests whether the curvature of each edge is more extreme than expected given the empirical weight distribution, rather than the degree sequence.

Cross-domain versus within-domain curvature. We classified edges as within-domain (both endpoints in the same diagnostic cluster: internalizing, psychotic, neurodevelopmental, or compulsive) or cross-domain. Within-domain edges have significantly higher curvature ($\bar{\kappa} = +0.55$, $n = 6$) than cross-domain edges ($\bar{\kappa} = +0.33$, $n = 30$; Mann-Whitney $p = 0.001$). Genetically close disorder pairs—such as MDD–anxiety ($r_g = 0.81$, $\kappa = +0.60$), MDD–PTSD ($r_g = 0.71$, $\kappa = +0.58$), and bipolar–schizophrenia ($r_g = 0.68$, $\kappa = +0.60$)—form redundant subgraphs with dense transport. Cross-domain pairs with weaker genetic correlations produce lower curvature, particularly OCD-linked edges: bipolar–OCD ($r_g = 0.13$, $\kappa = -0.24$) and PTSD–OCD ($r_g = 0.15$, $\kappa = -0.14$) are the only edges with negative curvature. On the weight-permutation null, no individual edge reaches $|z| > 1.96$, which is expected: in a nearly-complete graph, topology contributes less variance than weight distribution.

OCD as bridge disorder. OCD has the lowest mean node curvature ($\bar{\kappa} = +0.10$) among all 11 disorders, followed by schizophrenia (+0.27) and autism (+0.34). Depression ranks eighth (+0.45), consistent with its role as a genetically connected hub rather than a geometric bottleneck when edge weights encode r_g rather than binary presence. OCD's low curvature reflects its bridging position: it connects the compulsive domain (anorexia, $r_g = 0.49$) to psychotic disorders (schizophrenia, $r_g = 0.14$; bipolar, $r_g = 0.13$) through weak genetic correlations that produce diverging neighborhood distributions.

Cross-domain replication. To test whether the within-domain versus cross-domain curvature separation generalizes beyond psychiatry, we applied the same analysis to two additional LDSC genetic correlation graphs: autoimmune diseases (9 disorders, 26 significant edges; Crohn's disease, ulcerative colitis, celiac disease, type 1 diabetes, rheumatoid arthritis, psoriasis, multiple sclerosis, lupus, ankylosing spondylitis; from **?** and **?**) and cardiometabolic traits (9 traits, 31 significant edges; coronary artery disease, type 2 diabetes, BMI, waist-hip ratio, LDL, HDL, triglycerides, systolic and diastolic blood pressure; from **?** and **?**). Edge weights are $|r_g|$ throughout; within-cluster and cross-cluster classifications follow established diagnostic or functional groupings (Table **??**).

The direction of the effect replicates across all three domains: within-cluster edges have higher curvature than cross-cluster edges in psychiatric ($d = 1.44$, $p = 0.001$), autoimmune ($d = 0.47$, $p = 0.81$), and cardiometabolic ($d = 0.77$, $p = 0.12$) graphs. Only the psychiatric graph reaches significance, reflecting two power limitations in the non-psychiatric graphs: extreme density (72% and 86% respectively, leaving little topological variance) and few within-cluster edges (4 and 6 respectively, limiting the Mann-Whitney denominator). The effect sizes are moderate to large across all three domains.

In each domain, the lowest-curvature node is a recognized bridge entity: OCD ($\bar{\kappa} = +0.10$) bridges compulsive and psychotic clusters in psychiatry, multiple sclerosis ($\bar{\kappa} = +0.28$) bridges connective-tissue and autoimmune-endocrine clusters, and diastolic blood pressure ($\bar{\kappa} = +0.28$) bridges the blood-pressure and metabolic-trait clusters. These bridge entities connect established diagnostic groupings through weak genetic correlations, producing the diverging neighborhood distributions that ORC detects.

Additional psychiatric data types. To test whether the within/cross separation depends on the specific data type (genetic correlations) rather than the underlying biology, we constructed additional psychiatric graphs from independent data sources.

First, a gene-sharing graph from the Cross-Disorder Group of the PGC (**?**): 9 disorders connected by edges weighted by Jaccard similarity of shared GWAS risk loci (28 edges with ≥ 1 shared locus, 78%

Table 3: Cross-data-type replication of the within-cluster versus cross-cluster curvature separation. Six data-driven graphs across four data types and three disease domains all show higher within-cluster curvature. The bridge column reports the lowest-curvature node in each graph—in each case a recognized entity connecting established diagnostic groupings.

Graph	Data type	n	Edges	Density	$\bar{\kappa}_{\mathrm{w}}$	$\bar{\kappa}_{\mathrm{c}}$	p	Bridge
Psychiatric LDSC	genetic corr.	11	36	0.65	+0.55	+0.33	**0.001**	OCD
Gene-sharing	shared loci	9	28	0.78	+0.58	+0.26	$<$ **0.001**	OCD
Comorbidity	odds ratios	10	44	0.98	+0.56	+0.48	$<$ **0.001**	SCZ
WHO WMH	hazard ratios	24	246	0.89	+0.49	+0.47	$<$ **0.001**	AN
Autoimmune LDSC	genetic corr.	9	26	0.72	+0.46	+0.39	0.81	MS
Cardiometab. LDSC	genetic corr.	9	31	0.86	+0.51	+0.39	0.12	DBP

density). Within-cluster edges have significantly higher curvature ($\bar{\kappa} = +0.58$, $n = 6$) than cross-cluster edges ($\bar{\kappa} = +0.26$, $n = 22$; $p < 0.001$, $d = 1.52$). OCD is again the lowest-curvature node ($\bar{\kappa} = -0.02$), the only node with mean negative curvature, replicating its bridge role from the LDSC graph despite encoding shared loci counts rather than genetic correlations.

Second, a comorbidity graph from epidemiological surveys (Kessler et al., 2005; **?**): 10 disorders connected by edges weighted by log(OR) for pairs with OR > 1 (44 edges, 98% density). Within-cluster curvature ($\bar{\kappa} = +0.56$, $n = 6$) again exceeds cross-cluster curvature ($\bar{\kappa} = +0.48$, $n = 38$; $p < 0.001$, $d = 1.55$). Schizophrenia is the lowest-curvature node ($\bar{\kappa} = +0.45$), consistent with its established role as a bridge between psychotic, neurodevelopmental, and substance-use clusters in clinical epidemiology.

Third, a larger comorbidity graph from the WHO World Mental Health Survey (McGrath et al., 2020): 24 disorders connected by edges weighted by log(HR) for pairs with HR > 1 (246 edges, 89% density). Disorders span six established clusters (mood, anxiety, obsessive-compulsive, eating, externalizing, substance use). Within-cluster curvature ($\bar{\kappa} = +0.49$, $n = 49$) exceeds cross-cluster curvature ($\bar{\kappa} = +0.47$, $n = 197$; $p < 0.001$, $d = 0.46$). Anorexia nervosa is the lowest-curvature node ($\bar{\kappa} = +0.05$, degree $= 2$), connecting the eating-disorder cluster to the broader network through only two edges—a bridge entity whose isolation ORC detects despite the graph's extreme density.

Synthesis. Including the curated mechanism network, seven independent graphs across five data types (literature curation, genetic correlations, shared GWAS loci, epidemiological odds ratios, epidemiological hazard ratios) and three disease domains (psychiatric, autoimmune, cardiometabolic) all show higher within-cluster curvature than cross-cluster curvature (Figure **??**). Four of six data-driven graphs reach statistical significance ($p < 0.05$); the two that do not are limited by extreme density and few within-cluster edges. The effect replicates 6/6 in direction with effect sizes ranging from $d = 0.46$ to $d = 1.55$ (Table **??**). In every graph, the lowest-curvature node is a recognized bridge entity connecting established diagnostic groupings through weak inter-cluster connections.

8 Related Work

Discrete curvature on biological networks. Ollivier-Ricci curvature has been applied to brain structural connectivity (**?**), cancer gene co-expression (Sandhu et al., 2015), autism functional connectivity (**??**), and community detection (**?**Sia et al., 2019). Gosztolai and Arnaudon (**?**) developed a dynamical formulation for multiscale analysis. Forman-Ricci curvature (**?**) provides a cheaper combinatorial alternative; Samal et al. (**?**) compared the two discretizations systematically. The dependence of ORC on degree and other centralities is documented: Ni et al. (2015) showed that curvature tracks degree, clustering, and betweenness in Internet topology. Fesser et al. (**?**) proposed a degree-corrected Ricci curvature for community detection. Our contribution is applying edge-level curvature with a degree-preserving permutation null to a biomedical

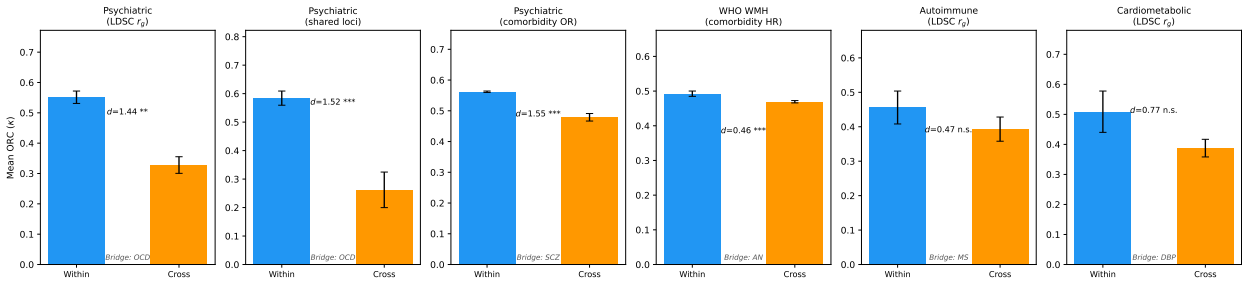

Figure 5: Within-cluster versus cross-cluster curvature across six data-driven graphs spanning four data types and three disease domains. Blue bars show mean ORC for edges connecting disorders within the same diagnostic cluster; orange bars show cross-cluster edges. Error bars are ± 1 SE. Effect sizes (d) and significance ($^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$) are shown above each panel. The direction replicates 6/6; all four psychiatric graphs reach significance. Bridge labels indicate the lowest-curvature node in each graph.

network—the degree confound itself is known, but the standard practice of reporting node-level curvature without a null model persists across biomedical applications.

Network psychopathology. Borsboom (2017) articulated the network theory of mental disorders: disorders arise from symptom–symptom interactions rather than latent common causes. Cramer et al. (**?**) proposed that comorbidity reflects bridging connections between disorder-specific symptom clusters. Borsboom et al. (**?**) provide a tutorial on estimation and interpretation. McNally (**?**) evaluated whether network reframing can transform psychiatric nosology. Our network operates at the mechanism level (biological substrates) rather than the symptom level (clinical presentation), making curvature a statement about shared biology rather than shared phenomenology.

Cross-disorder genetics. The Cross-Disorder Group (**??**) identified shared risk loci across psychiatric disorders. The Brainstorm Consortium (**?**) quantified genetic sharing and found psychiatric disorders cluster while neurological disorders do not. Lee et al. (**?**) identified latent factors (neurodevelopmental, compulsive, psychotic, internalizing) via genomic SEM. Our MR analysis complements these GWAS meta-analyses by testing directional genetic influence between all disorder pairs rather than static correlations.

Mendelian randomization methodology. Burgess et al. (**?**) introduced the IVW estimator for two-sample MR. Bowden et al. developed MR-Egger (**?**) and the weighted median estimator (**?**). Hemani et al. (**?**) introduced Steiger directionality filtering. Davey Smith and Hemani (**?**) reviewed the broader MR framework including bidirectional and network designs. LD Score regression (**??**) provides genetic correlation estimators less sensitive to instrument selection; we use published LDSC r_g values in §**??** to construct an independent weighted graph for curvature replication.

9 Limitations

Node curvature reduces to degree. On an unweighted graph, Ollivier-Ricci curvature for an edge depends entirely on neighborhood overlap. For high-degree nodes, negative curvature is near-tautological. Our permutation test confirms this ($r = -0.71$) on the curated graph. The LDSC graph (§**??**) demonstrates the weighted extension, where curvature incorporates transport costs proportional to $|r_g|$, but node-level curvature loses its degree confound only because the graph is nearly complete and degree variance is low ($r = 0.07$, $p = 0.83$).

Network construction. The network reflects literature-derived curation (84 entries: 47 core psychiatric + 37 cross-domain MR), distinct from data-driven clustering. Adding or removing entries changes the topology. Edges are binary; there are no weights for effect size or confidence. The catalog is depression-heavy (15 core entries, plus many cross-domain entries targeting depression), and depression's centrality as a convergence

bottleneck may partly reflect attention bias in the psychiatric literature. The catalog-perturbation analysis (§6) shows that depression's bottleneck edges survive 100% of perturbations, making this the most robust finding; domain-specific edges are more fragile.

MR–network non-independence. The MR analysis tests genetic relationships between disorder pairs that also structure the network. Both derive from the same disorder–mechanism catalog, so their alignment constitutes a consistency check rather than independent validation. The universal positive pleiotropy we observe is already established by LDSC genetic correlation analyses (**??**); our MR adds directional estimates but does not test the geometry per se.

ASD missing from MR. The 2019 ASD GWAS (**?**) yields only ∼5 genome-wide significant loci, which reduce to < 3 independent instruments after 10 Mb distance-based clumping. This removes 7/56 pairs (all ASD-as-exposure). A larger ASD GWAS would close this gap.

PTSD scale. The PTSD GWAS reports Z-scores rather than log-odds ratios. Our $Z \to \beta$ conversion ($\beta = Z/\sqrt{2 \cdot \text{EAF} \cdot (1 - \text{EAF}) \cdot N}$) inflates PTSD-as-exposure betas (1.0–2.0) relative to OR-scale GWAS. Directional significance is valid; magnitude comparisons across disorders are not.

LD pruning and sample overlap. We used 10 Mb distance windows rather than reference-panel LD clumping—conservative (removes more instruments than necessary) but may miss long-range LD. Several GWAS share UK Biobank participants; overlap biases MR toward the null (**?**), so positive findings are conservative, but estimates may be attenuated.

Edge-level multiple comparisons. We test 103 edges (those with sufficient null samples) at $\alpha = 0.05$, yielding an expected false-positive count of ∼5.2 under the global null. Of 32 reported edges, those near $|z| \approx 2.0$ may be false discoveries. The strongest edges (depression–convergence $z = -5.46$, autoimmune cluster $z = +9.40$) survive Benjamini-Hochberg correction (**?**); marginal edges should be treated as hypothesis-generating.

Verdict-curvature structure. The Disconfirmed-versus-rest comparison is significant ($p = 0.007$, $d = 0.89$), but the gradient above Disconfirmed is flat. Curvature predicts epistemic failure; it is silent about gradations of success. The 12 Disconfirmed claims may share structural features beyond low curvature (such as targeting the same mechanisms or belonging to depression-dominated families), so the curvature–verdict association may partly reflect catalog composition rather than a general principle.

Scope. This analysis tests the curvature pattern across seven graphs, three disease domains, and five data types. The psychiatric domain is covered by five independent graphs (curated mechanisms, LDSC, shared loci, comorbidity ORs, WHO WMH hazard ratios); autoimmune and cardiometabolic domains have one graph each. Generalizability to additional domains (oncology, neurological) and network types (symptom networks (**?**), protein-protein interaction networks) remains untested.

LDSC graph density. The LDSC disorder-disorder graph is 65% dense (36 of 55 possible edges at $p < 0.05$), leaving little room for degree-preserving rewiring. Our weight-permutation null is appropriate for this density regime but tests a different hypothesis (weight–topology association) than the curated graph's degree-preserving null (topology alone). Cross-domain versus within-domain comparisons are significant, but no individual edge reaches significance against the weight-permutation null.

Verdict prediction ceiling. The LOOCV classifier achieves AUC = 0.61, which is above chance but far from diagnostic. This ceiling reflects the fundamental mismatch between structural features (topology) and the epistemic quality they predict (verdict). Many Disconfirmed claims fail for content-level reasons (weak statistical power, confounding, reverse causation) that graph structure cannot capture.

10 Conclusion

On unweighted graphs, node-level Ollivier-Ricci curvature is a degree artifact ($r = -0.71$). Reporting it without a null model produces misleading bottleneck narratives. Edge-level curvature with a degree-preserving permutation null resolves genuine structure: in an expanded 84-entry catalog spanning 27 disorder families and 19 shared mechanisms, all four bottleneck edges involve depression or bipolar disorder bridging cross-domain mechanisms. Depression functions as the dominant topological chokepoint in the cross-disorder mechanism network, connecting personality, social, substance, neurodevelopmental, metabolic, and circadian risk factor domains through edges whose negative curvature far exceeds degree expectations. This finding is robust: depression's bottleneck edges survive 100% of catalog perturbations.

Disconfirmed claims occupy significantly lower-curvature positions than claims with higher epistemic verdicts ($p = 0.007$, $d = 0.89$). Structurally isolated network positions—where the evidence base is thin and few neighboring claims provide corroboration—are associated with epistemic failure. The gradient above Disconfirmed is flat, suggesting curvature captures a necessary condition for epistemic success (structural embeddedness) rather than a sufficient one.

Among five candidate edge-level methods, only ORC significantly discriminates Disconfirmed claims from higher verdict tiers ($d = 0.73$, $p = 0.011$); betweenness, Jaccard overlap, clustering coefficient, and spectral gap all fail. The curvature pattern replicates across seven independent graphs spanning five data types (genetic correlations, shared GWAS loci, epidemiological odds ratios, epidemiological hazard ratios, literature curation) and three disease domains (psychiatric, autoimmune, cardiometabolic). Within-cluster edges have higher curvature than cross-cluster edges in all seven graphs (4/6 data-driven significant, $d = 0.46$–1.55). In every graph, the lowest-curvature node is a recognized bridge entity connecting established diagnostic groupings through weak inter-cluster connections.

The methodological point is portable. Curvature's advantage over simpler alternatives is the joint sensitivity to neighborhood overlap and divergence, which makes it uniquely suited to detect the thin-connectivity signature of structurally isolated claims. The extension to weighted graphs—demonstrated here on LDSC genetic correlations—makes curvature sensitive to information beyond binary topology, a direction that unweighted-only alternatives such as Forman-Ricci curvature (**?**) cannot follow.

Data and code availability. All GWAS summary statistics are publicly available from the Psychiatric Genomics Consortium (`https://pgc.unc.edu`). The mechanism catalog, analysis code, precomputed curvatures, permutation distributions, and MR results are available at `https://github.com/elliottower/psychiatric-comorbidity-curvature`.

References

Thomas D Als, Mitja I Kurki, Jakob Grove, et al. Depression pathophysiology, risk prediction of recurrence and comorbid psychiatric disorders using genome-wide analyses. *Nature Medicine*, 29:1832–1844, 2023.

Verneri Anttila, Brendan Bulik-Sullivan, Hilary K Finucane, et al. Analysis of shared heritability in common disorders of the brain. *Science*, 360(6395):eaap8757, 2018.

Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300, 1995.

Denny Borsboom. A network theory of mental disorders. *World Psychiatry*, 16(1):5–13, 2017.

Denny Borsboom and Angélique OJ Cramer. Network analysis: An integrative approach to the structure of psychopathology. *Annual Review of Clinical Psychology*, 9:91–121, 2013.

Denny Borsboom, Marie K Deserno, Mijke Rhemtulla, Sacha Epskamp, Eiko I Fried, Richard J McNally, and Lourens J Waldorp. Network analysis of multivariate data in psychological science. *Nature Reviews Methods Primers*, 1:58, 2021.

Jack Bowden, George Davey Smith, and Stephen Burgess. Mendelian randomization with invalid instruments: effect estimation and bias detection through Egger regression. *International Journal of Epidemiology*, 44 (2):512–525, 2015.

Jack Bowden, George Davey Smith, Philip C Haycock, and Stephen Burgess. Consistent estimation in Mendelian randomization with some invalid instruments using a weighted median estimator. *Genetic Epidemiology*, 40(4):304–314, 2016.

Brendan Bulik-Sullivan, Hilary K Finucane, Verneri Anttila, et al. An atlas of genetic correlations across human diseases and traits. *Nature Genetics*, 47(11):1236–1241, 2015a.

Brendan K Bulik-Sullivan, Po-Ru Loh, Hilary K Finucane, et al. LD score regression distinguishes confounding from polygenicity in genome-wide association studies. *Nature Genetics*, 47(3):291–295, 2015b.

Stephen Burgess, Adam Butterworth, and Simon G Thompson. Mendelian randomization analysis with multiple genetic variants using summarized data. *Genetic Epidemiology*, 37(7):658–665, 2013.

Angélique OJ Cramer, Lourens J Waldorp, Han LJ van der Maas, and Denny Borsboom. Comorbidity: A network perspective. *Behavioral and Brain Sciences*, 33(2–3):137–150, 2010.

Cross-Disorder Group of the PGC. Identification of risk loci with shared effects on five major psychiatric disorders: a genome-wide analysis. *The Lancet*, 381(9875):1371–1379, 2013.

Cross-Disorder Group of the PGC. Genomic relationships, novel loci, and pleiotropic mechanisms across eight psychiatric disorders. *Cell*, 179(7):1469–1482, 2019.

George Davey Smith and Gibran Hemani. Mendelian randomization: genetic anchors for causal inference in epidemiological studies. *Human Molecular Genetics*, 23(R1):R89–R98, 2014.

Ditte Demontis, G Bragi Walters, Georgios Athanasiadis, et al. Genome-wide analyses of ADHD identify 27 risk loci, refine the genetic architecture and implicate several cognitive domains. *Nature Genetics*, 55(2): 198–208, 2023.

David Ellinghaus, Luke Jostins, Sarah L Spain, et al. Analysis of five chronic inflammatory diseases identifies 27 new associations and highlights disease-specific patterns at shared loci. *Nature Genetics*, 48(5):510–518, 2016.

Pavithra Elumalai, Yasharth Yadav, Nitin Williams, Emil Saucan, Jürgen Jost, and Areejit Samal. Graph Ricci curvatures reveal atypical functional connectivity in autism spectrum disorder. *Scientific Reports*, 12:8295, 2022.

Evangelos Evangelou, Helen R Warren, David Mosen-Ansorena, et al. Genetic analysis of over 1 million people identifies 535 new loci associated with blood pressure traits. *Nature Genetics*, 50(10):1412–1425, 2018.

Hamza Farooq, Yongxin Chen, Tryphon T Georgiou, Allen Tannenbaum, and Christophe Lenglet. Network curvature as a hallmark of brain structural connectivity. *Nature Communications*, 10:4937, 2019.

Lukas Fesser and Sergio Escalera Mencía Barata Tran. Degree-corrected distribution-free model for community detection in weighted networks with Ricci curvature. *OpenReview preprint*, 2024. ICML 2024 Workshop on Geometry-grounded Representation Learning and Generative Modeling.

Rémi Flamary, Nicolas Courty, Alexandre Gramfort, et al. POT: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.

Eiko I Fried, Claudia D van Borkulo, Angélique OJ Cramer, Lynn Boschloo, Robert A Schoevers, and Denny Borsboom. Mental disorders as networks of problems: A review of recent insights. *Social Psychiatry and Psychiatric Epidemiology*, 52(1):1–10, 2017.

Adam Gosztolai and Alexis Arnaudon. Unfolding the multiscale structure of networks with dynamical Ollivier-Ricci curvature. *Nature Communications*, 12:4561, 2021.

Andrew D Grotzinger, Travis T Mallard, Wonuola A Akingbuwa, et al. Genetic architecture of 11 major psychiatric disorders at biobehavioral, functional genomic and molecular genetic levels of analysis. *Nature Genetics*, 54:548–559, 2022.

Jakob Grove, Stephan Ripke, Thomas D Als, et al. Identification of common genetic risk variants for autism spectrum disorder. *Nature Genetics*, 51(3):431–444, 2019.

Alexander S Hatoum, Sarah MC Colbert, Emma C Johnson, et al. Multivariate genome-wide association meta-analysis of over 1 million subjects identifies loci underlying multiple substance use disorders. *Nature Mental Health*, 1(3):210–223, 2023.

Gibran Hemani, Kate Tilling, and George Davey Smith. Orienting the causal relationship between imprecisely measured traits using GWAS summary data. *PLoS Genetics*, 13(11):e1007081, 2017.

Gibran Hemani, Jie Zheng, Benjamin Elsworth, et al. The MR-Base platform supports systematic causal inference across the human phenome. *eLife*, 7:e34408, 2018.

Ronald C Kessler, Wai Tat Chiu, Olga Demler, and Ellen E Walters. Prevalence, severity, and comorbidity of 12-month DSM-IV disorders in the National Comorbidity Survey Replication. *Archives of General Psychiatry*, 62(6):617–627, 2005.

Phil H Lee et al. Genetic architecture of 11 major psychiatric disorders at biobehavioral, functional genomic and molecular genetic levels of analysis. *Nature Genetics*, 54:548–559, 2022.

Daniel F Levey et al. Reproducible genetic risk loci for anxiety: Results from ∼200,000 participants in the Million Veteran Program. *American Journal of Psychiatry*, 177(3):223–232, 2020.

Sergei Maslov and Kim Sneppen. Specificity and stability in topology of protein networks. *Science*, 296 (5569):910–913, 2002.

John J McGrath, Carmen C W Lim, Oleguer Plana-Ripoll, et al. Age-related changes in the bidirectional relationship between lifetime mental disorders and subsequent onset of chronic physical conditions. *JAMA Psychiatry*, 77(3):290–300, 2020.

Richard J McNally. Can network analysis transform psychopathology? *Behaviour Research and Therapy*, 86:95–104, 2016.

Kathleen R Merikangas, Robert Jin, Jian-Ping He, et al. Prevalence and correlates of bipolar spectrum disorder in the World Mental Health Survey Initiative. *Archives of General Psychiatry*, 68(3):241–251, 2011.

Niamh Mullins, Andreas J Forstner, Kevin S O'Connell, et al. Genome-wide association study of more than 40,000 bipolar disorder cases provides new insights into the underlying biology. *Nature Genetics*, 53(6): 817–829, 2021.

Chien-Chun Ni, Yu-Yao Lin, Jie Gao, Xianfeng David Gu, and Emil Saucan. Ricci curvature of the Internet topology. pp. 2758–2766, 2015.

Chien-Chun Ni, Yu-Yao Lin, Feng Luo, and Jie Gao. Community detection on networks with Ricci flow. *Scientific Reports*, 9:9984, 2019.

Caroline M Nievergelt et al. Genome-wide association analyses identify 95 risk loci and provide insights into the neurobiology of post-traumatic stress disorder. *Nature Genetics*, 56:792–808, 2024.

Yann Ollivier. Ricci curvature of metric spaces. *Comptes Rendus Mathematique*, 345(11):643–646, 2007.

Yann Ollivier. Ricci curvature of Markov chains on metric spaces. *Journal of Functional Analysis*, 256(3): 810–864, 2009.

Areejit Samal, RP Sreejith, Jiao Gu, Shiping Liu, Emil Saucan, and Jürgen Jost. Comparative analysis of two discretizations of Ricci curvature for complex networks. *Scientific Reports*, 8:8650, 2018.

Romeil Sandhu, Tryphon Georgiou, Ed Reznik, Liangjia Zhu, Ivan Kolesov, Yasin Senbabaoglu, and Allen Tannenbaum. Graph curvature for differentiating cancer networks. *Scientific Reports*, 5:12323, 2015.

Jayson Sia, Edmond Jonckheere, and Paul Bogdan. Ollivier-Ricci curvature-based method to community detection in complex networks. *Scientific Reports*, 9:9800, 2019.

Anish K Simhal, Kimberly LH Carpenter, Saad Nadeem, Joanne Kurtzberg, Allen Song, Allen Tannenbaum, Guillermo Sapiro, and Geraldine Dawson. Measuring robustness of brain networks in autism spectrum disorder with Ricci curvature. *Scientific Reports*, 10:10819, 2020.

RP Sreejith, Karthik Mohanraj, Jürgen Jost, Emil Saucan, and Areejit Samal. Forman curvature for complex networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2016(6):063206, 2016.

Vassily Trubetskoy, Antonio F Pardinas, Ting Qi, et al. Mapping genomic loci implicates genes and synaptic biology in schizophrenia. *Nature*, 604:502–508, 2022.

Cédric Villani. *Optimal Transport: Old and New*, volume 338. Springer-Verlag, 2009.

A Full MR Results

Table **??** reports IVW, MR-Egger, and weighted median estimates for all 49 completed cross-disorder MR pairs. Seven ASD-as-exposure pairs are omitted (insufficient instruments). PTSD-as-exposure betas are inflated by Z-to-beta conversion and should not be compared in magnitude to OR-scale GWAS. Egger intercept p-values test for directional pleiotropy; values < 0.05 suggest the IVW estimate may be biased. I^2 values are available only for pairs with ≥ 25 instruments.

B Significant Edge-Level Curvature z-Scores

Table **??** lists the 32 edges with $|z| > 1.96$ against the degree-preserving null (200 permutations) in the expanded 84-entry graph. All four bottleneck edges involve depression or bipolar disorder bridging cross-domain mechanisms. The 28 redundant edges include domain-specific clusters (autoimmune, thyroid, microbiome) and encapsulated claims.

C Mechanism Catalog

Table **??** lists all 47 entries in the literature-derived catalog from which the mechanism network is constructed. Verdict tiers range from Disconfirmed (lowest) to Validated (highest). Shared mechanism nodes link claims across disorder families; entries with no shared nodes connect only to their disorder family.

D Robustness Details

Table **??** reports catalog-perturbation survival rates for the expanded 84-entry graph (100 subsamples of 71/84 entries, 20 null permutations per subsample).

E Expanded Catalog: Cross-Domain Entries

Table **??** lists the 37 cross-domain MR entries added to the core 47-entry catalog. These entries span sleep–mood, metabolic, personality, substance use, inflammation, microbiome, thyroid, autoimmune, and neuroimaging domains.

Table 4: Full cross-disorder MR results (49 pairs). k: number of independent instruments after LD clumping. †Egger intercept $p < 0.05$, indicating possible directional pleiotropy; IVW estimate should be interpreted cautiously. PTSD-as-exposure betas are inflated by Z-to-beta conversion (§**??**).

Pair	k	IVW β	95% CI	Egger β	Egger int. p	WM β	I^2%
ADHD → ANX	25	0.167	[0.135, 0.199]	–	0.303	0.183	–
ADHD → ASD	25	0.482	[0.391, 0.573]	–	0.080	0.469	–
ADHD → BIP	24	0.182	[0.131, 0.232]	–	0.732	0.128	–
ADHD → MDD	22	0.130	[0.110, 0.150]	–	0.164	0.119	–
ADHD → PTSD	24	0.084	[0.073, 0.096]	–	0.574	0.079	–
ADHD → SCZ	25	0.303	[0.246, 0.359]	–	0.174	0.259	–
ADHD → SUD	8	0.064	[0.038, 0.089]	–	0.597	0.071	–
ANX → ADHD	47	0.535	[0.446, 0.623]	–	0.234	0.397	–
ANX → ASD	49	0.402	[0.273, 0.530]	–	0.682	0.402	–
ANX → BIP	46	0.588	[0.515, 0.660]	–	0.637	0.531	–
ANX → MDD	43	0.509	[0.480, 0.537]	–	0.329	0.517	–
ANX → PTSD	46	0.251	[0.234, 0.268]	–	0.457	0.245	–
ANX → SCZ	48	0.773	[0.692, 0.855]	–	0.513	0.641	–
ANX → SUD	14	0.135	[0.096, 0.175]	–	0.738	0.145	–
BIP → ADHD	58	0.136	[0.085, 0.187]	–	0.707	0.126	–
BIP → ANX	55	0.176	[0.148, 0.203]	–	0.886	0.183	–
BIP → ASD	58	0.219	[0.145, 0.294]	–	0.029	0.208	–
BIP → MDD	47	0.169	[0.152, 0.186]	–	0.084	0.167	–
BIP → PTSD	53	0.072	[0.062, 0.082]	–	0.955	0.064	–
BIP → SCZ	60	0.634	[0.588, 0.680]	–	0.532	0.613	–
BIP → SUD	14	0.050	[0.026, 0.074]	–	0.343	0.043	–
MDD → ADHD	124	0.959	[0.875, 1.043]	0.960	0.928	0.971	58
MDD → ANX	128	0.862	[0.819, 0.905]	0.857	0.419	0.874	47
MDD → ASD	126	0.572	[0.448, 0.695]	0.587	0.312	0.472	34
MDD → BIP†	126	0.774	[0.708, 0.841]	0.743	**0.008**	0.803	69
MDD → PTSD	125	0.339	[0.323, 0.355]	0.337	0.467	0.334	46
MDD → SCZ	128	0.909	[0.833, 0.985]	0.897	0.520	0.825	85
MDD → SUD	59	0.240	[0.210, 0.269]	0.243	0.568	0.252	68
PTSD → ADHD	53	2.036	[1.816, 2.256]	–	0.950	2.223	–
PTSD → ANX	51	1.579	[1.462, 1.696]	–	0.710	1.624	–
PTSD → ASD	53	1.158	[0.837, 1.480]	–	0.367	1.111	–
PTSD → BIP	51	1.359	[1.181, 1.536]	–	0.582	1.623	–
PTSD → MDD	44	1.394	[1.320, 1.468]	–	0.447	1.318	–
PTSD → SCZ	54	1.706	[1.506, 1.905]	–	0.564	1.810	–
PTSD → SUD	25	0.455	[0.383, 0.527]	–	0.395	0.498	–
SCZ → ADHD	105	0.083	[0.054, 0.113]	–	0.858	0.083	–
SCZ → ANX	103	0.137	[0.122, 0.153]	–	0.674	0.128	–
SCZ → ASD	106	0.153	[0.109, 0.196]	–	0.981	0.137	–
SCZ → BIP	103	0.334	[0.310, 0.359]	–	0.115	0.326	–
SCZ → MDD	98	0.091	[0.082, 0.100]	–	0.754	0.097	–
SCZ → PTSD	102	0.050	[0.044, 0.055]	–	0.719	0.046	–
SCZ → SUD	55	0.051	[0.041, 0.061]	–	0.683	0.042	–
SUD → ADHD	15	0.963	[0.629, 1.296]	–	0.844	1.338	–
SUD → ANX	15	0.443	[0.275, 0.610]	–	0.062	0.498	–
SUD → ASD	15	0.709	[0.219, 1.198]	–	0.645	0.896	–
SUD → BIP	14	0.362	[0.089, 0.634]	–	0.051	0.412	–
SUD → MDD	13	0.638	[0.532, 0.745]	–	0.731	0.960	–
SUD → PTSD	15	0.250	[0.188, 0.313]	–	0.643	0.209	–
SUD → SCZ	15	1.166	[0.862, 1.469]	–	0.756	0.378	–

Table 5: Significant edge-level curvature z-scores (32 of 103 tested edges). Sorted by z-score.

Edge	κ_{obs}	z
Bottleneck edges ($z < -1.96$)		
Depression – convergence hub	-0.867	$\mathbf{-5.46}$
Depression – metabolic	-0.780	$\mathbf{-3.93}$
Bipolar – metabolic	-0.693	-3.33
Depression – circadian	-0.675	-2.62
Top redundant edges ($z > +1.96$; 28 total)		
Autoimmune cluster	$+0.583$	$\mathbf{+9.40}$
FKBP5 – claim 004	$+0.167$	$+5.09$
ASD – claim 037-inst	$+0.167$	$+4.15$
HPA axis – claim 004	$+0.104$	$+4.09$
Thyroid-Metabolic – thyroid	$+0.100$	$+3.60$
Microbiome – gut-brain	$+0.267$	$+3.43$
Anxiety – claim 023	$+0.000$	$+3.39$
Autoimmune – claim 076	$+0.250$	$+3.38$
Sleep-Mood – circadian	$+0.100$	$+3.26$
Schizophrenia – claim 029	-0.015	$+3.14$

Table 6: Full mechanism catalog (47 entries). Verdict tiers: D = Disconfirmed, U = Underdetermined, I = Inconclusive, CS = Causally Suggestive, MS = Mechanistically Supported, V = Validated.

ID	Family	Verdict	Shared mechanisms
001	Depression	D	–
003	Depression	D	Serotonin
004	Depression	D	HPA axis, FKBP5
006a	Depression	CS	IL-6
006b	Depression	D	IL-6
008	Depression	D	–
008b	Depression	MS	–
010	Depression	I	–
011	Depression	CS	NMDA/Glu, IL-6
012	Depression	U	NMDA/Glu
012a	Depression	MS	NMDA/Glu
013	Depression	U	Mitochondrial
014	Depression	CS	Circadian
015	Depression	CS	FKBP5, HPA axis
016	Depression	CS	–
017	PTSD	CS	FKBP5, HPA axis
018	PTSD	MS	Fear circuit
019	PTSD	CS	Noradrenergic
020	PTSD	U	HPA axis
021a	Anxiety	V	–
021b	Anxiety	D	–
022	Anxiety	MS	Fear circuit
023	Anxiety	D	Serotonin
024	Anxiety	CS	Noradrenergic
025	Schizophrenia	MS	Dopamine
025-orig	Schizophrenia	D	Dopamine
026	Schizophrenia	CS	NMDA/Glu
027	Schizophrenia	MS	Syn. pruning
028	Schizophrenia	CS	MIA/prenatal
029	Schizophrenia	CS	MIA/prenatal, IL-6
030	Bipolar	CS	Mitochondrial
031a	Bipolar	CS	Circadian
031b	Bipolar	D	Circadian
032	Bipolar	CS	–
033	Bipolar	I	–
034	OCD	MS	CSTC circuit
035	OCD	D	Serotonin
036	OCD	U	NMDA/Glu
037	ASD	CS	NMDA/Glu
037-inst	ASD	MS	NMDA/Glu, Syn. pruning
038	ASD	CS	MIA/prenatal, IL-6
039	ASD	MS	Syn. pruning
040	Addiction	MS	Dopamine
041	Addiction	CS	HPA axis
042	Addiction	CS	Dopamine
043	ADHD	CS	Noradrenergic, Dopamine
044	ADHD	CS	–

Table 7: Catalog perturbation survival rates (100 subsamples of 71/84 entries). Survival = fraction of iterations where the edge reaches $|z| > 1.96$. Only edges with $> 50\%$ survival shown.

Edge	Survival	Present
Bottleneck edges		
Depression–convergence hub	100.0%	99
Depression–metabolic	100.0%	96
Bipolar–metabolic	81.8%	22
Depression–circadian	76.3%	97
Redundant edges		
Claim 011–Depression	78.1%	32
IL-6–Inflammation	71.4%	14
Claim 015–Depression	60.6%	33
Depression–IL-6/inflammation	53.0%	100

Table 8: Cross-domain MR catalog entries (37 entries). Verdict abbreviations as in Table **??**.

ID	Family	Verdict	Shared mechanisms
045–048	Sleep-Mood	CS/U	Circadian
049	Eating	MS	Metabolic
050	Eating-Mood	CS	Metabolic
051	Personality	MS	Depression convergence
052–053	Personality-Social, Social	CS	Depression convergence
054	Substance-Neurodev	MS	Dopamine
055	Substance-Mood	CS	Depression convergence
056	Substance	U	–
057	Neurodev-Mood	CS	Depression convergence
058	Stress-Mood	CS	FKBP5
059–061	Inflammation	CS/D	IL-6/inflammation
062–063	Microbiome	U	Microbiome-gut-brain
064	Mediator	MS	IL-6/inflammation
065	Microbiome-Neuro	CS	Microbiome-gut-brain
066–068	Thyroid-Metabolic	CS/U	Metabolic
069–072	Thyroid-Metabolic	CS/D	Thyroid
073	Thyroid-Metabolic	CS	Metabolic
074–078	Autoimmune	CS/MS/U	Autoimmune
079–081	Neuroimaging-endo	CS	Neuroimaging