

Supplementary Figures




Figure S1. Scatter plot of kNN semantic distance (x-axis) versus maximum co-clustering membership (y-axis) for all 15 non-capped scenarios on each corpus. Each point represents one scenario-corpus cell; the y-value is the largest number of that scenario's articles simultaneously present in one micro-cluster incarnation under burst injection (seed 42). GN and ST panels are shown separately. Spearman correlations are $\rho = +0.34$ (GN, $p = 0.219$) and $\rho = +0.19$ (ST, $p = 0.503$), neither significant at $n = 15$. The scatter confirms that the tier-grouped pattern visible in Table 3 does not resolve into a clean per-scenario gradient at this sample size.




Figure S2. Two panels. Left: DBSTREAM detection count as a function of the clustering_threshold (ct) parameter at values 0.8, 1.0, and 1.2, extended from the submitted 6 cells (3 scenarios, 1 seed) to 22 cells (5 scenarios, 2 corpora, 3 seeds). The cliff replicates: 0/22 at ct = 0.8, 15/22 at the sensitive midpoint ct = 1.0, and 22/22 at ct = 1.2. One borderline flip at ct = 1.0 (aethelgard/ST, seed 456). Per-bar counts shown as n/N. Right: cosine novelty detector detection counts (out of 32 scenario-corpus cells: 16 scenarios x 2 corpora, burst injection) across 9 parameter combinations (3 percentile thresholds x 3 window sizes). All combinations produce 0 or 1 detections under burst, confirming that burst-blindness is architectural and not a consequence of the chosen operating point.




Figure S3. Null distribution of maximum co-occurrence for topic-incoherent noise articles ($k = 12$ subsamples of the Control B noise set, 20,000 draws per corpus, gray histogram), with the observed scenario outcomes overlaid as triangles: green for scenarios whose maximum co-occurrence is permutation-significant at $p < 0.05$, orange for non-significant. The dashed line marks $N = 5$, the detection threshold used throughout the paper, which sits at approximately the null's 95th percentile: $P(\text{null max} \geq 5) = 0.058$ on GN and 0.056 on ST. Detected scenarios separate cleanly from the null mass on both corpora, confirming that co-occurrence at $N = 5$ is unlikely to arise from incidental embedding proximity among topically unrelated articles.

Detection and absorption across four detector arms

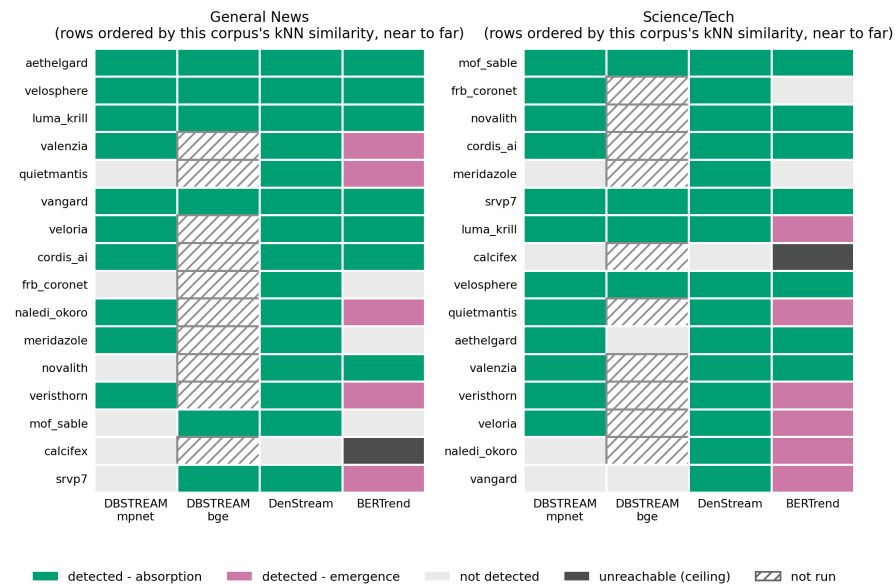


Figure S4. Cross-paradigm detection grid: 16 scenarios × 4 detector arms × 2 corpora under burst injection (seed 42, N = 5 for the three clustering detectors; BERTrend at lineage level). Each cell is marked detected or undetected. DenStream detects 15/16 on both corpora (missing only calcifex), subsuming the other three arms. BERTrend detects 12/16 GN and 13/16 ST. DBSTREAM detects 10/16 GN and 12/16 ST. Cosine (burst) detects 0/16 on both corpora. Scenarios ordered by increasing ST kNN similarity.

Supplementary Tables

Table S1. Full per-scenario DBSTREAM results under burst injection (seed 42 unless noted). Weak signal = first step at which 2 or more scenario articles co-occupied one micro-cluster incarnation. Recall = fraction of the scenario's articles injected up to the detection step that are present in the detecting incarnation (recall = 1.00 indicates detection at the earliest possible step). All detections are absorption into a pre-existing real cluster; no scenario formed a new cluster. A dash in the delay or recall columns indicates the scenario was not detected at N = 5; weak-signal values may still appear for undetected scenarios (e.g., quietmantis/GN weak signal at 1.09 days reflects a transient 2-article co-occurrence that never reached the N = 5 threshold).

Scenario	Corpus	kNN	n	Tier	Det. (s42)	Seeds	Abs. delay (d)	Weak sig. (d)	Recall
aethelgard	GN	0.779	12	near	+	3/3	1.45	0.36	0.56
velosphere	GN	0.728	12	near	+	3/3	2.18	0.36	0.45
luma_krill	GN	0.685	12	medium	+	2/3	0.00	0.36	1.00
valenzia	GN	0.657	12	far	+	2/3	1.45	0.73	0.56
quietmantis	GN	0.656	12	medium	–	2/3	–	1.09	–
vanguard	GN	0.653	12	far	+	3/3	0.00	0.36	1.00
veloria	GN	0.644	12	far	+	3/3	0.00	0.36	1.00
cordis_ai	GN	0.640	12	near	+	3/3	0.36	0.36	0.83
frb_coronet	GN	0.634	5	near	–	0/3	–	–	–
naledi_okoro	GN	0.633	12	far	+	3/3	0.00	0.36	1.00
meridazole	GN	0.626	6	medium	+	2/3	0.00	0.80	1.00
novalith	GN	0.625	12	near	–	1/3	–	–	–
veristhorn	GN	0.588	12	far	+	3/3	0.00	0.36	1.00
mof_sable	GN	0.574	12	near	–	1/3	–	2.18	–
calcifex	GN	0.556	4	medium	–	0/3	–	4.00	–
srp7	GN	0.527	12	medium	–	0/3	–	3.27	–
mof_sable	ST	0.732	12	near	+	3/3	0.36	0.36	0.83
frb_coronet	ST	0.727	5	near	+	3/3	0.00	1.00	1.00
novalith	ST	0.696	12	near	+	3/3	0.00	0.36	1.00
cordis_ai	ST	0.676	12	near	+	3/3	0.00	0.36	1.00
meridazole	ST	0.655	6	medium	–	2/3	–	1.60	–
srp7	ST	0.651	12	medium	+	3/3	0.00	0.36	1.00
luma_krill	ST	0.646	12	medium	+	3/3	0.00	0.36	1.00
calcifex	ST	0.624	4	medium	–	0/3	–	–	–
velosphere	ST	0.599	12	near	+	3/3	0.00	0.36	1.00
quietmantis	ST	0.589	12	medium	+	3/3	0.00	0.36	1.00
aethelgard	ST	0.547	12	near	+	2/3	0.00	0.36	1.00
valenzia	ST	0.532	12	far	+	2/3	0.00	0.36	1.00
veristhorn	ST	0.512	12	far	+	3/3	0.00	0.36	1.00
veloria	ST	0.479	12	far	+	2/3	0.73	0.73	0.71
naledi_okoro	ST	0.467	12	far	–	0/3	–	–	–
vanguard	ST	0.416	12	far	–	0/3	–	–	–

Table S2. The pre-existing real cluster that absorbed each detected scenario (seed 42). Cluster size = total members of the detecting incarnation at the moment of detection. Enrichment = the scenario's density in the cluster relative to its base rate in the stream. Specificity bins: low < 5×, moderate 5 to 20×, high > 20× enrichment; the low-specificity threshold is harmonized with the Figure 4 binning in the main text. velosphere (mc 247) is the only high-specificity (> 20×) GN detection. Topic tags are derived from sampled article titles. Cluster sizes differ between scenarios sharing a micro-cluster id because absorption occurred at different points in the stream; the size is a snapshot at each scenario's detection step.

Scenario	Corpus	mc id	Cluster topic	Size at detection	Enrichment	Low spec.	Shared by
naledi_okoro	GN	184	background news B	5,806	2.2×	yes	2
vanguard	GN	184	background news B	5,343	2.3×	yes	2
luma_krill	GN	229	background news A	2,642	4.7×	yes	6
meridazole	GN	229	background news A	2,589	9.7×	no	6
valenzia	GN	229	background news A	2,391	5.2×	no	6
veristhorn	GN	229	background news A	2,376	5.3×	no	6
veloria	GN	229	background news A	1,882	6.6×	no	6
aethelgard	GN	64	markets/shares	1,334	9.4×	no	1
cordis_ai	GN	229	background news A	1,199	10.4×	no	6
velosphere	GN	247	US-N.Korea politics	238	52.5×	no	1
mof_sable	ST	460	EV/engine tech	1,722	4.9×	yes	2
novalith	ST	460	EV/engine tech	1,328	6.3×	no	2
luma_krill	ST	297	paleontology	1,287	6.5×	no	4
veloria	ST	297	paleontology	1,261	6.7×	no	4
velosphere	ST	637	defense/drones	1,228	6.9×	no	2
frb_coronet	ST	472	astronomy	1,126	17.9×	no	1
srvp7	ST	552	quantum	1,089	7.7×	no	1
cordis_ai	ST	160	AI	914	9.2×	no	1
aethelgard	ST	297	paleontology	850	9.9×	no	4
veristhorn	ST	297	paleontology	536	15.7×	no	4
quietmantis	ST	637	defense/drones	391	21.5×	no	2
valenzia	ST	447	archaeology	126	66.8×	no	1

Table S3. The 10 undetected scenario-corpus cells at N = 5 (seed 42), classified by failure mechanism. Scatter: articles never co-clustered (maximum 1 per incarnation). Fragmentation: partial co-clustering (2 to 4 articles) that never reached the threshold. Structural cap: fewer than 5 articles exist in the scenario. Max co-cluster = largest number of the scenario's articles simultaneously in one incarnation.

Scenario	Corpus	n	kNN	Failure mode	Max co-cluster	Detected (other seeds)	Note
quietmantis	GN	12	0.656	fragmentation	3	2/3	seed-sensitive
frb_coronet	GN	5	0.634	scatter	1	0/3	structurally tight (exactly N = 5)
novalith	GN	12	0.625	scatter	1	1/3	seed-sensitive
mof_sable	GN	12	0.574	fragmentation	2	1/3	seed-sensitive
calcifex	GN	4	0.556	structural cap	2	0/3	
srvp7	GN	12	0.527	fragmentation	2	0/3	
meridazole	ST	6	0.655	fragmentation	4	2/3	seed-sensitive
calcifex	ST	4	0.624	structural cap	1	0/3	
naledi_okoro	ST	12	0.467	scatter	1	0/3	
vanguard	ST	12	0.416	scatter	1	0/3	

Table S4. Absolute delay (days beyond the theoretical minimum) for the Tier 2 DBSTREAM runs that detected, across injection schedules and both generators (3 scenarios: aethelgard, srvp7, vangard; seed 42). Undetected runs are omitted. Achieved latency is dominated by the injection span; the detector itself adds near-zero delay in most cells. Flash-then-silence outcome across all 12 DBSTREAM cells: 8 undetected, including all 6 GN cells; the 4 detected cells are all ST (aethelgard and srvp7, both generators) at zero delay, where all five articles co-clustered within the flash window.

Scenario	Corpus	Schedule	Generator	Abs. delay (d)	Achieved (d)
aethelgard	ST	accelerating	claude	77.28	231.85
aethelgard	ST	accelerating	gemini	29.44	184.01
aethelgard	GN	accelerating	gemini	20.36	44.78
srvp7	ST	accelerating	claude	0.00	154.57
srvp7	ST	accelerating	gemini	0.00	154.57
aethelgard	GN	burst	claude	1.45	2.91
aethelgard	GN	burst	gemini	0.00	1.45
aethelgard	ST	burst	claude	0.00	1.45
aethelgard	ST	burst	gemini	0.00	1.45
srvp7	ST	burst	claude	0.00	1.45
srvp7	ST	burst	gemini	0.00	1.45
vangard	GN	burst	claude	0.00	1.45
vangard	GN	burst	gemini	0.00	1.45
aethelgard	ST	flash	claude	0.00	7.00
aethelgard	ST	flash	gemini	0.00	7.00
srvp7	ST	flash	claude	0.00	7.00
srvp7	ST	flash	gemini	0.00	7.00
aethelgard	ST	steady	claude	25.76	128.81
aethelgard	ST	steady	gemini	25.76	128.81
srvp7	ST	steady	claude	0.00	103.04
srvp7	ST	steady	gemini	0.00	103.04
vangard	GN	steady	claude	0.00	16.28
vangard	GN	steady	gemini	0.00	16.28

Table S5. Detection outcome and maximum co-cluster membership ($N = 5$, burst, DBSTREAM) when scenario intra-cluster tightness is degraded from the as-generated baseline (approximately 0.91 mean pairwise similarity) toward looser targets (0.70, 0.60, 0.50) by adding calibrated Gaussian noise to the article embeddings. The grid covers the 4 robustness scenarios on both corpora at seed 42; cells show maximum co-cluster membership, with detection at $N = 5$ marked + or -. quietmantis/GN is undetected at baseline and therefore uninformative for degradation analysis. The multi-seed block tests seed stability on General News across seeds 42, 123, and 456 at the two most aggressive targets; informative conditions exclude rows undetected at baseline for that seed.

Full degradation grid (seed 42): maximum co-cluster membership at each tightness target

Scenario	Corpus	0.91 (baseline)	0.70	0.60	0.50
veristhorn	GN	12 (+)	11 (+)	11 (+)	6 (+)
luma_krill	GN	12 (+)	11 (+)	10 (+)	9 (+)
velosphere	GN	6 (+)	5 (+)	6 (+)	6 (+)
quietmantis	GN	3 (–)	1 (–)	1 (–)	1 (–)
veristhorn	ST	12 (+)	12 (+)	11 (+)	11 (+)
luma_krill	ST	12 (+)	12 (+)	12 (+)	12 (+)
velosphere	ST	12 (+)	12 (+)	12 (+)	11 (+)
quietmantis	ST	12 (+)	12 (+)	12 (+)	12 (+)

Detection is 7/8 at every target (quietmantis/GN the sole miss throughout). Multi-seed stability (GN, seeds 42/123/456): 10/10 informative conditions survived at target 0.60; 9/10 at target 0.50 (single failure: quietmantis/GN/seed 456, max co-membership 3). Detection is robust to substantial tightness degradation: loosening intra-scenario similarity from 0.91 to 0.50, below the natural-cluster median, eliminates only 1 of 10 informative multi-seed conditions. Source: table_s5_variance_robustness.csv in the released artifact.

Table S6. Per-scenario DBSTREAM detection ($N = 5$, burst, seed 42) for Claude-generated versus Gemini-generated articles across all 16 scenarios on both corpora (32 pairs). This extends the Tier 2 generator agreement ($\kappa = 0.917$, 24 pairs) reported in the main text to the full distance gradient. The only discordant cells are meridazole/GN and meridazole/ST, a small borderline scenario (6 articles, medium tier). Generator choice is immaterial across the full distance gradient, not only the 3 Tier 2 scenarios. Source: generator_agreement_burst16 in statistical_tests.json.

Set	n pairs	Cohen's kappa	Agreement
All 16 scenarios x 2 corpora	32	0.855	30/32
Excluding calcifex (structural cap)	30	0.83	28/30

Table S7. Summary of statistical tests. Spearman correlations computed on the 15 non-capped scenarios; bootstrap confidence intervals from 10,000 resamples. At $n = 15$, correlations are reported with wide intervals and interpreted descriptively; the cosine correlations are the only distance relationships significant at conventional thresholds. Headline detection rates with Wilson 95% confidence intervals: DBSTREAM GN 10/16 [0.39, 0.82] (seed range 9 to 12); DBSTREAM ST 12/16 [0.51, 0.90] (range 11 to 12); cosine burst 0/32 [0.00, 0.11]; cosine steady 6/32 [0.09, 0.35]; cosine accelerating 5/32 [0.07, 0.32].

Test	Statistic	p	95% CI	Interpretation
GN vs ST detection gap (McNemar exact, 15 pairs)	8 discordant (5 ST-only, 3 GN-only)	0.727	—	Not significant; corpora differ in mechanism, not rate
Enrichment GN vs ST (Wilcoxon, 7 scenarios detected on both)	W = 9.0	0.469	—	Not significant
Generator agreement (Cohen's kappa, 24 Tier 2 pairs)	kappa = 0.917 (95.8% agreement)	—	—	Generator choice irrelevant; 1 discordant cell
Distance vs detection, DBSTREAM binary, GN	rho = +0.49	0.063	[0.00, 0.81]	Not significant; wide CI
Distance vs detection, DBSTREAM binary, ST	rho = +0.39	0.156	[-0.25, 0.78]	Not significant
Distance vs co-clustering strength, GN	rho = +0.34	0.219	[-0.24, 0.71]	Not significant
Distance vs co-clustering strength, ST	rho = +0.19	0.503	[-0.56, 0.75]	Not significant
Distance vs cosine flags, steady, GN	rho = -0.61	0.016	[-0.88, -0.13]	Significant negative gradient
Distance vs cosine flags, steady, ST	rho = -0.88	<0.001	[-0.97, -0.60]	Strongest result in the study
Distance vs cosine flags, accelerating, GN	rho = -0.56	0.031	[-0.86, -0.05]	Significant
Distance vs cosine flags, accelerating, ST	rho = -0.88	<0.001	[-0.97, -0.60]	Significant
Chance co-occurrence of 5 articles (permutation null, 20,000 draws)	P = 0.058 GN / 0.056 ST	—	—	N = 5 aligns with the null's 95th percentile
Generator agreement, full gradient (Cohen's kappa, 32 burst pairs)	kappa = 0.855 (30/32); 0.83 excl. calcifex	—	—	Generator choice immaterial; discordant: meridazole/GN, /ST. See Table B6
Detector score vs global distance (measure independence, n = 5,000/corpus)	Pearson r = 0.76 GN / 0.61 ST; Spearman 0.76 / 0.65	—	—	Related but distinct (r-squared 0.37 to 0.58); distance-detection not circular

Table S8. Encoder comparison: DBSTREAM detection and absorption under bge-base-en-v1.5 on the 6-scenario subset, beside the mpnet results on the same scenarios. Block 1: within-space gradient (Spearman rank correlation between kNN similarity and the mpnet design-time ordering; mpnet 0.879 GN / 0.985 ST, bge 0.893 / 0.968). Block 2: per-scenario detection agreement (9/12 agree, 3 disagree; flips run in both directions). Block 3: absorption classification (all 10 bge-detected cells = absorption, synthetic fraction 0.001 to 0.007). The bge clustering threshold was re-derived per corpus under the pre-stated protocol (Section 2.2); all other pipeline parameters are identical.

Corpus	Scenario	Tier	mpnet det.	bge det.	Agreement	mpnet enrich.	bge enrich.
gn	aethelgard	near	True	True	agree	9.4×	1.6×
gn	luma_krill	medium	True	True	agree	4.7×	2.9×
gn	mof_sable	near	False	True	flip-to-detected	n/a (not detected)	1.7×
gn	srvp7	medium	False	True	flip-to-detected	n/a (not detected)	7.1×
gn	vanguard	far	True	True	agree	2.3×	2.2×
gn	velosphere	near	True	True	agree	52.5×	7.1×
st	aethelgard	near	True	False	flip-to-undetected	9.9×	n/a (not detected)
st	luma_krill	medium	True	True	agree	6.5×	11.2×
st	mof_sable	near	True	True	agree	4.9×	0.9×
st	srvp7	medium	True	True	agree	7.7×	6.2×
st	vanguard	far	False	False	agree	n/a (not detected)	n/a (not detected)
st	velosphere	near	True	True	agree	6.9×	0.9×

Table S9. Cross-detector overlap between DenStream lineages and DBSTREAM micro-cluster absorbers for the 22 cells detected by both (10 GN, 12 ST; seed 42, N = 5). Jaccard index computed on real-article membership at the detection snapshot. Median Jaccard 0.15; only 1 of 22 cells reaches 0.5. The DenStream lineage is consistently larger (median size ratio 2.26x GN, 2.86x ST), and the median share of the DBSTREAM absorber contained inside it is 40.8% on GN and 85.5% on ST. Disagreements are one-directional: GN 5 of 16, ST 3 of 16, every disagreement a DenStream detection with no DBSTREAM counterpart.

Corpus	Scenario	DBSTREAM mem- bers	DenStream members	Jaccard	DBSTREAM share contained	Size ratio
gn	aethelgard	1329	559	0.353	0.371	0.420
gn	cordis_ai	1194	2374	0.015	0.044	1.990
gn	luma_krill	2637	6664	0.156	0.476	2.530
gn	meridazole	2584	6607	0.153	0.472	2.560
gn	nal- edi_okoro	5801	877	0.056	0.060	0.150
gn	valenzia	2386	6262	0.126	0.404	2.620
gn	vangard	5338	3171	0.348	0.412	0.590
gn	veloria	1877	705	0.036	0.049	0.380
gn	velosphere	233	1104	0.090	0.476	4.740
gn	veristhorn	2371	6501	0.142	0.467	2.740
st	aethelgard	845	4392	0.004	0.026	5.200
st	cordis_ai	909	881	0.525	0.678	0.970
st	frb_coro- net	1121	3741	0.226	0.799	3.340
st	luma_krill	1282	3677	0.306	0.906	2.870
st	mof_sable	1717	3150	0.443	0.870	1.830
st	novalith	1323	2067	0.487	0.840	1.560
st	quietman- tis	386	2555	0.135	0.904	6.620
st	srvp7	1084	5103	0.101	0.525	4.710
st	valenzia	121	1010	0.103	0.876	8.350
st	veloria	1256	3580	0.307	0.904	2.850
st	velosphere	1223	3145	0.063	0.212	2.570
st	veristhorn	531	988	0.480	0.928	1.860

Table S10. BERTrend classification of the 25 lineage-detected cells (lin_stratum = met). Each cell carries its single-window purity (sw_purity) and lineage purity (lin_purity), both classification labels (lin_label_provisional and window_snapshot_label), both stratum indicators (sw_stratum and lin_stratum), and the knife-edge flags (sw_knife_edge, lin_knife_edge). Of the 24 cells detected at both levels, two flip from emergence to absorption between the snapshot and lineage views (gn/cordis_ai and gn/luma_krill); these are the two downward bound-crossings and represent one finding, not two independent results. The 25th cell (gn/novalith) carries different labels across levels but is not comparable because it does not reach $N \geq 5$ within any single window. Tie-split cells (st/vanguard and st/veloria) move from emergence to absorption under the alternative share-last rule, giving a lineage emergence range of 11 (share-first, adopted) to 9.

Corpus	Scenario	SW purity	Lin. purity	Snapshot label	Lineage label	SW stratum	Lin. stratum	SW knife-edge
gn	aethelgard	0.556	0.042	absorption_analog	absorption_analog	met	met	False
gn	cordis_ai	0.545	0.421	emergence_analog	absorption_analog	met	met	False
gn	luma_krill	0.625	0.258	emergence_analog	absorption_analog	met	met	False
gn	naledi_okoro	0.857	0.786	emergence_analog	emergence_analog	met	met	False
gn	novalith	0.500	0.050	emergence_analog	absorption_analog	not_met	met	True
gn	quietmantis	0.857	0.625	emergence_analog	emergence_analog	met	met	False
gn	srvp7	1.000	1.000	emergence_analog	emergence_analog	met	met	False
gn	valenzia	0.857	0.857	emergence_analog	emergence_analog	met	met	False
gn	vanguard	0.250	0.022	absorption_analog	absorption_analog	met	met	False
gn	veloria	0.462	0.297	absorption_analog	absorption_analog	met	met	False
gn	velosphere	0.667	0.036	absorption_analog	absorption_analog	met	met	False
gn	veristhorn	0.625	0.625	emergence_analog	emergence_analog	met	met	False
st	aethelgard	0.238	0.321	absorption_analog	absorption_analog	met	met	False
st	cordis_ai	0.185	0.010	absorption_analog	absorption_analog	met	met	False
st	luma_krill	1.000	1.000	emergence_analog	emergence_analog	met	met	False
st	mof_sable	0.240	0.045	absorption_analog	absorption_analog	met	met	False
st	naledi_okoro	0.833	0.833	emergence_analog	emergence_analog	met	met	False
st	novalith	0.417	0.175	absorption_analog	absorption_analog	met	met	False
st	quietmantis	0.714	0.692	emergence_analog	emergence_analog	met	met	False
st	srvp7	0.500	0.014	absorption_analog	absorption_analog	met	met	True
st	valenzia	0.417	0.310	absorption_analog	absorption_analog	met	met	False
st	vanguard	0.833	0.833	emergence_analog	emergence_analog	met	met	False
st	veloria	0.714	0.714	emergence_analog	emergence_analog	met	met	False
st	velosphere	0.429	0.077	absorption_analog	absorption_analog	met	met	False
st	veristhorn	1.000	1.000	emergence_analog	emergence_analog	met	met	False

Table S11. DenStream detection and absorption attribution per scenario per corpus (seed 42, N = 5), with the corresponding DBSTREAM verdict. All 30 DenStream detections classify as absorption (synthetic fraction < 0.5 in every case). The DBSTREAM column shows that DBSTREAM's 22 detections are a strict subset of DenStream's 30; the 8 DenStream-only cells are listed. Enrichment ratios and the low-specificity flag (below the 5x heuristic or below DenStream's native p95 floor of 3.60x GN / 5.348x ST) are reported per cell. Disagreement cells between the native floor and the 5x heuristic are named in both directions.

Scenario	Corpus	Enrichment	Low spec.	Classification	DBSTREAM det.	DenStream det.
aethelgard	gn	22.171	False	absorbed	True	True
cordis_ai	gn	5.256	False	absorbed	True	True
frb_coronet	gn	9.218	False	absorbed	False	True
luma_krill	gn	1.875	True	absorbed	True	True
meridazole	gn	3.782	True	absorbed	True	True
mof_sable	gn	8.314	False	absorbed	False	True
naledi_okoro	gn	14.178	False	absorbed	True	True
novalith	gn	7.253	False	absorbed	False	True
quietmantis	gn	6.533	False	absorbed	False	True
srvp7	gn	5.439	False	absorbed	False	True
valenzia	gn	1.995	True	absorbed	True	True
vanguard	gn	3.937	True	absorbed	True	True
veloria	gn	17.612	False	absorbed	True	True
velosphere	gn	11.276	False	absorbed	True	True
veristhorn	gn	1.922	True	absorbed	True	True
aethelgard	st	1.915	True	absorbed	True	True
cordis_ai	st	9.506	False	absorbed	True	True
frb_coronet	st	5.394	False	absorbed	True	True
luma_krill	st	2.287	True	absorbed	True	True
meridazole	st	2.857	True	absorbed	False	True
mof_sable	st	2.669	True	absorbed	True	True
naledi_okoro	st	1.880	True	absorbed	False	True
novalith	st	4.065	True	absorbed	True	True
quietmantis	st	3.290	True	absorbed	True	True
srvp7	st	1.649	True	absorbed	True	True
valenzia	st	8.298	False	absorbed	True	True
vanguard	st	1.695	True	absorbed	False	True
veloria	st	2.349	True	absorbed	True	True
velosphere	st	2.674	True	absorbed	True	True
veristhorn	st	8.481	False	absorbed	True	True