CURIOSITYBITS LIVING LAB BOSTON · STUDY #07 TECHNICAL & ECONOMETRIC AUDIT SUPPLEMENT (V1.0)

Clicks Without Conversation? Forensic Technical & Econometric Supplement — Study #07 (`Edition v1.0`)

Authors: CuriosityBits Living Lab Boston (Architecture B Peer-to-Peer Discovery & Red-Team Swarm: `@theorist`, `@data-gatekeeper`, `@falsifier`, `@econometrician`, `@red-team`)

Study ID: `study-07` (`projects/virality-deliberation` | `100%` Pure Communication & Media Science)

Edition: `v1.0` (Layer 3 Forensic Technical & Audit Supplement synchronized with Layer 2 Academic Research Article `PAPER_ARTICLE.md` and Layer 1 Public Interactive Feature `web/virality-deliberation.html`; `SKILL-029`)

Pre-Data Pre-Registration Lock (`STUDY_DESIGN_CARD_V1_PRE_DATA_LOCKED.md` & `STUDY_DESIGN_CARD.md` SHA-256): `7c25286e6b930f06560b29fe50b3bd023c8d55ef66611dfd2990bd70430362ab` (Verified in standalone pre-data Git Commit `131fa8a` / `131fa8ab956ea3b2ecba5617ac082c708693d8a1`)

Verified Literature Corpus (`LITERATURE_SYNTHESIS_VERIFIED.json` SHA-256): `097aa728e6f8ec5c8de277db0e7ded5a3d2878a20972d643b5ea351575073c05` (`64` OpenAlex + Consensus AI works)

Pre-Data Swarm Discovery & Rejected-Directions Ledger (`SWARM_DISCOVERY_AND_FEASIBILITY_LEDGER.md` SHA-256): `aaa9b7ca1f2e1a1d0c9e6c215533c11f66bc98c7c2a8953f1a49693388fb1b5a`

Independent Peer Review & Red-Team Report (`INDEPENDENT_PEER_REVIEW_REPORT.md` SHA-256): `f53870692bb2833014bfc01fad1e1760b1a0abe8de5d23f555b9b1eb57daf1de`

Post-Data Swarm Theoretical Stress-Test (`SWARM_THEORETICAL_STRESS_TEST.md` SHA-256): `2c88c5738d86edb0e7dff389abe6cd56c9b77d9f498fa2403876af21cd3bb569`

Reference Verification Stdout (`results/VERIFY_STUDY_07_STDOUT.txt` SHA-256): `5d1af3e4428adfc76bc776e44474ed012b65862cadbb510186c6ce90552db901`

Section 1. Cryptographic Pre-Registration Lock, Rejected Candidate Directions & Artifact Registry (`SKILL-024` , `SKILL-025` , `SKILL-027` , `SKILL-029`)

1A. Standalone Pre-Data Git Commit, Pure Communication Domain Mandate & Zero-HARKing Verification

Prior to full-panel Algolia Hacker News universe enumeration, recursive `/items/{story_id}` comment-tree harvesting, four-judge blind OpenRouter annotation (`scripts/build_dataset.py`), or confirmatory regression estimation (`scripts/run_econometrics_and_stress_tests.py`), Study #07 locked its research questions (`RQ1` – `RQ4`), four-way strong-inference competing mechanisms (`H1` , `H2` , `H3` , `H0` , plus `H4` Paywall Moderation), deterministic tree-pruning and link-stripping equations, two-way clustering rules, Holm–Bonferroni multiplicity families, and 10 falsification gates in standalone Git commit `131fa8a` (`git cat-file -t 131fa8a` → `commit`). Both `STUDY_DESIGN_CARD_V1_PRE_DATA_LOCKED.md` and `STUDY_DESIGN_CARD.md` remain `100%` byte-for-byte identical (`SHA-256: 7c25286e6b930f06560b29fe50b3bd023c8d55ef66611dfd2990bd70430362ab`).

Study #07 enforces a strict **pure Communication & Media Science domain boundary**: `0` equity-price or market-return metrics, `0` trading/financial columns, `0` corporate regulatory filings, and `0` business/financial outlets across all `M = 85` curated general, investigative, science, technology, and culture newsrooms (`2021-2025`). When confirmatory estimation yielded a joint null on Stage-1 binary `Discussed_Breakthrough` (`num_comments >= 15` : `Neg_z = +0.0158` , t = +1.52, p_{Holm} = 0.3946; `Conflict_z = -0.0059` , t = -0.51, p_{Holm} = 0.6199) alongside a significant positive `Moral_Conflict_Blame` slope on Stage-2 `Dyadic_ABA_Reciprocity_Rate_48h` (`+0.0135` , t = +2.82, p_{Holm} = 0.0240 in the primary `K = 4` family), the pre-registered direction-blind adjudication rule (`Mixed_Or_Boundary_Specific_Equilibrium`) was preserved with `0` post-data hypothesis edits.

1B. Pre-Data Adversarial Swarm Discovery & Rejected Candidate Directions Summary

Before pre-registration, `@theorist` , `@data-gatekeeper` , and `@falsifier` stress-tested five pure communication candidate designs via live server-side API probes and construct-fit/identification audits (complete probe telemetry archived in `SWARM_DISCOVERY_AND_FEASIBILITY_LEDGER.md` , `SHA-256: aaa9b7ca1f2e1a1d0c9e6c215533c11f66bc98c7c2a8953f1a49693388fb1b5a`):

Candidate ID	Proposed Communication Theory Clash	Pre-Registration Gatekeeper & Falsifier Rejection Grounds	Final Pre-Data Disposition
`CAND-02`	Normative Contagion vs. Community Self-Regulation of Opening-Comment	**Hard Rejected (`SKILL-022` , `SKILL-025`):** Public Reddit JSON endpoints returned `HTTP 403` blocks from cloud IP space, and live open-weight LLM calibration on moderated Hacker News opening comments revealed severe zero-floor censoring on	**HARD REJECTED**

Candidate ID	Proposed Communication Theory Clash	Pre-Registration Gatekeeper & Falsifier Rejection Grounds	Final Pre-Data Disposition
	Incivility in Public Discussion Threads	`personal_incivility` alongside survivorship confounding from active moderator pruning (`[dead]` / `[flagged]`).	
`CAND-04`	Intermedia Agenda-Setting vs. Audience Frame Selection Across Multi-Outlet Breaking-News Dyads	**Hard Rejected (`SKILL-013` , `SKILL-025`):** Live multi-year Wikipedia Current Events Portal dyad probes revealed that platform `[dupe]` suppression leaves almost no simultaneous multi-outlet dyads crossing the discussion threshold for the same breaking event, rendering within-event dyad fixed effects statistically underpowered and biased by arrival-order lock-in.	**HARD REJECTED**
`CAND-05`	Episodic Outrage Decay vs. Collaborative Epistemic Memory Repair on Public-Affairs Wikipedia Articles	**Hard Rejected (`SKILL-022`):** Live Wikimedia `Talk:` revision probes showed that a majority of breaking-news entity pages accumulate zero `Talk:` deliberation edits within post-event windows while conflating encyclopedic housekeeping with public deliberation and overlapping methodologically with Study #06's Wikimedia attention infrastructure.	**HARD REJECTED**
`CAND-03`	Epistemic Overstatement ("Hype") vs. Communicative Hedging in Science Journalism & Institutional Press Releases	**Rejected Standalone → Absorbed (`SKILL-020` , `SKILL-022`):** Conditioning on viral science threads without an unconditional Stage-1 denominator created a severe viral-breakthrough collider, while top science domains mixed primary research manuscripts with journalistic reporting. Science/ health/environment journalism (`Beat 3`) and documentary evidence link substantiation (`External_Evidence_Link_Rate_48h`) were absorbed cleanly into `CAND-01` .	**REJECTED STANDALONE (BEAT & `Y3` ABSORBED)**
`CAND-01`	**The Virality–Deliberation Paradox: Affective Negativity vs. Moral-Conflict Framing Across Two-Stage Nested Public Discourse**	**Adopted & Fortified (`6` Forensic Safeguards + `4` Refinements):** Enforced unconditional Stage-1 + conditional Stage-2 two-stage hurdle sampling, `0 < dt <= 48h` ancestor-pruned trees, interpersonal `A != B` `ABA` reciprocity (`0` `A->A` self-replies), `html.unescape()` mirror/echo link stripping (`Paywall_Mirror_Link_Rate_48h` split out for `H4`), 4-knot `log(V_48)` splines, and deterministic URL-slug verbatim headline verification (`Gate 2`).	**LOCKED AS STUDY #07**

1C. Complete Tracked Artifact SHA-256 Registry (`SKILL-025` , `SKILL-027`)

Artifact Path (`projects/virality-deliberation/`)	Full 64-Character SHA-256 Digest
`STUDY_DESIGN_CARD_V1_PRE_DATA_LOCKED.md`	`7c25286e6b930f06560b29fe50b3bd023c8d55ef66611dfd2990bd704`
`STUDY_DESIGN_CARD.md`	`7c25286e6b930f06560b29fe50b3bd023c8d55ef66611dfd2990bd704`
`LITERATURE_SYNTHESIS_VERIFIED.json`	`097aa728e6f8ec5c8de277db0e7ded5a3d2878a20972d643b5ea3515`
`SWARM_DISCOVERY_AND_FEASIBILITY_LEDGER.md`	`aaa9b7ca1f2e1a1d0c9e6c215533c11f66bc98c7c2a8953f1a496933`
`INDEPENDENT_PEER_REVIEW_REPORT.md`	`f53870692bb2833014bfc01fad1e1760b1a0abe8de5d23f555b9b1eb`
`SWARM_THEORETICAL_STRESS_TEST.md`	`2c88c5738d86edb0e7dff389abe6cd56c9b77d9f498fa2403876af21`

Artifact Path (projects/virality-deliberation/)	Full 64-Character SHA-256 Digest
scripts/build_dataset.py	a6d77edb84fa83b466b4efe4cf5db20cc31478df0115621c941e47a07
scripts/run_econometrics_and_stress_tests.py	cf4af2be6144e955b1e8ba816380b12ddfe97f8402b21c68d05b0a64b
scripts/verify_study_07.py	2975721f2bd33dedce6cf70892a2d4364157a3593d70d3523e1b5b2b5
data/stage1_submissions_panel.csv	b456c545440250f853fdba8fec64f72fc17e5afa06a941c5f7cf18728
data/stage2_deliberation_trees_panel.csv	fe66ee05068befcb26cf91360f3c6ba15b87b48b001c3b5d82568d8f7
data/raw_judge_ballots.json	0cd9e66a64460b77484f77b8b04717d4a5f36d2788b4012e559366e7b

Artifact Path (projects/virality-deliberation/)	Full 64-Character SHA-256 Digest
data/headline_lexicons.json	4de109c0b49625c620ce3897b8c5de74fd03848c29a27a3d1099932b5
data/sample_attrition_and_diagnostics.json	a0481d13da01a1ddf6890f707374e8b4d5c16aee1fd169841a0e8a2b3
results/econometric_results.json	2ddf0ff9fd70366331ffc00201ddb70eacfc9a047094575d79fa31b9c
results/chart_series_summary.json	bbffc16596804d748aa257f4f65eab474de92534334a5abd65aa44e40
results/VERIFY_STUDY_07_STDOUT.txt	5d1af3e4428adfc76bc776e44474ed012b65862cadbb510186c6ce905

Section 2. Curated Journalism Universe, Two-Stage Attrition Waterfall & Distribution Diagnostics (`SKILL-011` , `SKILL-013` , `SKILL-017`)

2A. Curated Non-Financial Journalism Universe & Two-Stage Attrition Waterfall (`2021-2025`)

Pipeline Stage	Deterministic Filter / Stratified Sampling Rule	Retained Count (`N` , `G_domain` / `M` , `G_week`)	Dropped / Verification Telemetry
Outlet Roster	Pre-specified `M_probed = 85` non-financial general, investigative, science, tech & culture outlets (`>=5` stories)	`M = 85` qualifying outlets (`M = 58` open-access, `M = 27` paywalled)	`0` financial or business wire outlets included in roster
Algolia Scan	Exhaustive `created_at_i` window partition scan across `2021-2025` (`tags=story`)	`N = 221497` raw hits → **`N = 220696` clean universe**	Dropped `294` host/fin mismatch, `0` duplicates, `0` missing fields, `14` meta-posts (`Show/Ask/Launch HN`), `493` market-headline regex hits (`N = 23453` discussed `>=15` comments across `M = 85` outlets)
Stage 1 (`Y1`)	Stratified `(domain, year)` unconditional crowd gatekeeping sample (`seed = 20261117`)	**`N = 3653`** , `G_domain = 85` (`M = 85`), `G_week = 262`	Exact universe `Sum IPW = 220696.00` ; `N = 398` (`10.90%`) discussed stories (`num_comments >= 15`) entering Stage 2
Stage 2 (`Y2` / `Y3`)	Stratified discussed tree harvest (`2251` candidates → `2251` HTTP 200 recursive JSON trees)	**`N = 2114`** , `G_domain = 85` (`M = 85`), `G_week = 261`	`0` HTTP failures; dropped `137` trees with `<10` valid `48h` comments; `226317` valid `48h` comments (`235510` valid `14d` comments); `562` (`0.25%`) `A->A` self-replies excluded
Live 4-Judge JEV	Blind 4-model OpenRouter headline scoring (`Llama-3.3-70B` , `Qwen-2.5-72B` , `DeepSeek-V3-0324` , `Gemma-3-27B`)	`5398` unique headlines × `4` judges = **`21592` live ballots**	`0` synthetic, cached-template, or dictionary fallback scores (`100%` live API provenance)

2B. Primary Outcome Distributions & Natural 3-Sigma Tail Preservation (`SKILL-017`)

Panel & Outcome Variable	N	Mean	SD	Min	p01	Median
`stage1_log_points` (\log(1 + points))	3653	+1.8228	1.1932	+0.6931	+0.6931	+1.3863
`stage1_log_comments` (\log(1 + num_comments))	3653	+0.8259	1.3806	+0.0000	+0.0000	+0.0000
`stage2_dyadic_aba_reciprocity_rate_48h`	2114	+0.2117	0.1053	+0.0000	+0.0000	+0.2042
`stage2_interpersonal_nested_reply_share_48h`	2114	+0.7982	0.1111	+0.3000	+0.4384	+0.8234
`stage2_branch_depth_p90_48h`	2114	+5.2674	1.6128	+2.0000	+2.1000	+5.0000
`stage2_external_evidence_link_rate_48h`	2114	+0.0991	0.0650	+0.0000	+0.0000	+0.0874
`stage2_speaker_shannon_entropy_48h`	2114	+0.9548	0.0269	+0.7200	+0.8620	+0.9598

Section 3. Shrout & Fleiss (`1979`) 4-Judge Reliability, Per-Judge Calibration & 4-Cell Typology (`SKILL-019` , `SKILL-022`)

3A. Multi-Formula Shrout–Fleiss Intraclass Correlations & Floor/Ceiling Diagnostics

Analysis Sample & Construct	N	ICC(1,4)	ICC(2,4)	ICC(3,4)	ICC(2,1)	Pairwise r	Mean	SD
Stage 1 (`N_stage1`): `headline_negativity`	3653	0.9664	**0.9665**	0.9682	0.8782	0.8886	0.2944	0.273
Stage 1 (`N_stage1`): `moral_conflict_blame`	3653	0.9649	**0.9650**	0.9655	0.8732	0.8768	0.2465	0.274
Stage 2 (`N_trees`): `headline_negativity`	2114	0.9637	**0.9638**	0.9659	0.8694	0.8806	0.3460	0.275
Stage 2 (`N_trees`): `moral_conflict_blame`	2114	0.9623	**0.9623**	0.9629	0.8645	0.8684	0.2903	0.281

Note (`SKILL-022` Zero-Mass Disclosure): Exact zero-floor rates are `<35%` across all four construct-stage cells (`15.14%` to `31.37%`), and ceiling censoring is `<=0.08%` . Because Stage-1 unconditional headlines include routine science, technology, and cultural reporting without interpersonal blame (`Stage 1`

moral_conflict_blame <= 0.05 = 43.33%), [SECTION 6] Check E explicitly verifies two-part Any_Conf > 0.05 hurdle and positive-conflict subsample specifications.

3B. Per-Judge Open-Weight Model Calibration Matrix & 4-Cell Typology Counts

Open-Weight LLM Judge (OpenRouter Live API)	Stage-1 Neg Mean (SD)	Stage-1 Conf Mean (SD)	Stage-2 Neg Mean (SD)	Stage-2 Conf Mean (SD)
j1_llama (meta-llama/llama-3.3-70b-instruct)	0.3043 (0.3129)	0.2381 (0.2934)	0.3589 (0.3157)	0.2805 (0.3027)
j2_qwen (qwen/qwen-2.5-72b-instruct)	0.2611 (0.2576)	0.2371 (0.2781)	0.3091 (0.2653)	0.2815 (0.2927)
j3_deepseek (deepseek/deepseek-chat-v3-0324)	0.2999 (0.2896)	0.2448 (0.2745)	0.3478 (0.2874)	0.2876 (0.2798)
j4_gemma (google/gemma-3-27b-it)	0.3123 (0.2817)	0.2658 (0.3048)	0.3681 (0.2843)	0.3117 (0.3117)

- **Inter-Construct Correlations & Pre-Registered 0.50 vs. Sample-Median 4-Cell Counts:**
- **Stage 1 (N = 3653, Corr(Neg, Conflict) = +0.7533):** At pre-registered 0.50 cutoffs: Cell_0_Neutral_Constructive = 2342, Cell_1_Pure_Threat_Tragedy_Negativity = 441, Cell_2_Moral_Conflict_Blame = 568, Cell_3_NonNegative_Debate = 302. At sample medians (neg_median = 0.25, conf_median = 0.125): Cell_0 = 1547, Cell_1 = 239, Cell_2 = 1591, Cell_3 = 276.
- **Stage 2 (N = 2114, Corr(Neg, Conflict) = +0.7176):** At pre-registered 0.50 cutoffs: Cell_0_Neutral_Constructive = 1198, Cell_1_Pure_Threat_Tragedy_Negativity = 315, Cell_2_Moral_Conflict_Blame = 412, Cell_3_NonNegative_Debate = 189. At sample medians (neg_median = 0.325, conf_median = 0.2): Cell_0 = 820, Cell_1 = 197, Cell_2 = 890, Cell_3 = 207.

Section 4. Confirmatory Econometric Estimation (Models 1–4) & Holm–Bonferroni Adjudication

All primary models absorb four-way fixed effects (domain + year_quarter + utc_dow_6h_slot + url_beat, 0 singletons dropped in both Stage 1 and Stage 2) alongside six structural headline controls (headline_word_count, headline_char_len, has_question_mark, has_colon_or_dash, has_digits, uppercase_acronym_ratio) and two deterministic lexical controls (lexical_negativity_rate from 2444 words, lexical_moral_conflict_rate from 185 words;

`SKILL-020`), with two-way `(domain, year_week)` Cameron–Gelbach–Miller (`CGM`) standard errors (`df_t = 84`).

4A. Strong-Inference Adjudication & Primary `K = 4` Holm–Bonferroni Family

- **Adjudicated Strong-Inference Mechanism:** `Mixed_Or_Boundary_Specific_Equilibrium` (Live spot re-derivation check `< 1e-9` `[PASS]`).
- **Primary `K = 4` Co-Primary Holm–Bonferroni Family:**
- Stage-1 `discussed_breakthrough` $\beta_{1,\text{neg}}$ (`Joint Neg_z`): `coef = +0.015772` (`+0.0158`), `SE = 0.010356` (`0.0104`), `t = +1.5230` (`+1.52`), raw `p = 0.131523` (`0.1315`), **`p_Holm(K=4) = 0.394568` (`0.3946`)**.
- Stage-1 `discussed_breakthrough` $\beta_{1,\text{conf}}$ (`Joint Conflict_z`): `coef = -0.005923` (`-0.0059`), `SE = 0.011502` (`0.0115`), `t = -0.5150` (`-0.51`), raw `p = 0.607913` (`0.6079`), **`p_Holm(K=4) = 0.619933` (`0.6199`)**.
- Stage-2 `dyadic_aba_reciprocity_rate_48h` $\beta_{2,\text{neg}}$ (`Joint Neg_z`): `coef = +0.003460` (`+0.0035`), `SE = 0.003388` (`0.0034`), `t = +1.0215` (`+1.02`), raw `p = 0.309967` (`0.3100`), **`p_Holm(K=4) = 0.619933` (`0.6199`)**.
- Stage-2 `dyadic_aba_reciprocity_rate_48h` $\beta_{2,\text{conf}}$ (`Joint Conflict_z`): **`coef = +0.013542` (`+0.0135`), `SE = 0.004804` (`0.0048`), `t = +2.8193` (`+2.82`), raw `p = 0.006002` (`0.0060`), `p_Holm(K=4) = 0.024010` (`0.0240`)**.

4B. Table 1 — Model 1: Stage-1 Unconditional Crowd Gatekeeping & Viral Breakthrough (`RQ1`, `N = 3653`, `G_domain = 85`, `G_week = 262`)

(Joint `VIFs`: `headline_negativity_z = 2.3527`, `moral_conflict_blame_z = 2.4105`; `within-R2`: `0.0062` on `discussed_breakthrough`, `0.0203` on `log_points`, `0.0125` on `log_comments`)

Stage-1 Outcome & Specification	Coefficient (coef)	Two-Way SE	t-Stat	Raw p-Value	95% CI	Standardized β_{std} (95% CI_std)
`discussed_breakthrough` (`>=15`): Joint `Neg_z`	`+0.015772` (`+0.0158`)	`0.010356` (`0.0104`)	`+1.5230` (`+1.52`)	`0.131523` (`0.1315`)	`[-0.0048, +0.0364]`	`+0.0506 SD` (`[-0.0155, +0.1167]`)
`discussed_breakthrough` (`>=15`): Joint `Conflict_z`	`-0.005923` (`-0.0059`)	`0.011502` (`0.0115`)	`-0.5150` (`-0.51`)	`0.607913` (`0.6079`)	`[-0.0288, +0.0169]`	`-0.0190 SD` (`[-0.0924, +0.0544]`)

Stage-1 Outcome & Specification	Coefficient (coef)	Two-Way SE	t-Stat	Raw p-Value	95% CI	Standardized β_{std} (95% CI_std)
discussed_breakthrough : Joint Raw [0,1] Neg / Conf	+0.057694 / -0.021609	0.037883 / 0.041960	+1.52 / -0.51	0.1315 / 0.6079	[-0.0176, +0.1330] / [-0.1051, +0.0618]	Raw [0,1] scale
discussed_breakthrough : Standalone Neg_z	+0.012033 (+0.0120)	0.007282 (0.0073)	+1.6524 (+1.65)	0.102181 (0.1022)	[-0.0024, +0.0265]	+0.0386 SD ([-0.0079, +0.0851])
discussed_breakthrough : Standalone Conflict_z	+0.005721 (+0.0057)	0.008118 (0.0081)	+0.7047 (+0.70)	0.482929 (0.4829)	[-0.0104, +0.0219]	+0.0184 SD ([-0.0334, +0.0702])
log_points : Joint Neg_z	+0.057924 (+0.0579)	0.034976 (0.0350)	+1.6561 (+1.66)	0.101434 (0.1014)	[-0.0116, +0.1275]	+0.0485 SD ([-0.0097, +0.1068])
log_points : Joint Conflict_z	**+0.087165 (+0.0872)**	0.038071 (0.0381)	**+2.2895 (+2.29)**	**0.024556 (0.0246)**	**[+0.0115, +0.1629]**	**+0.0731 SD ([+0.0096, +0.1365])**
log_points : Joint Raw [0,1] Neg / Conf	+0.211894 / +0.317977	0.127948 / 0.138882	+1.66 / +2.29	0.1014 / 0.0246	[-0.0425, +0.4663] / [+0.0418, +0.5942]	Raw [0,1] scale
log_points : Standalone Neg_z	**+0.112936 (+0.1129)**	0.024190 (0.0242)	**+4.6687 (+4.67)**	**0.000011 (0.0000)**	**[+0.0648, +0.1610]**	+0.0946 SD ([+0.0543, +0.1350])
log_points : Standalone Conflict_z	**+0.129933 (+0.1299)**	0.026641 (0.0266)	**+4.8773 (+4.88)**	**0.000005 (0.0000)**	**[+0.0770, +0.1829]**	+0.1089 SD ([+0.0645, +0.1533])
log_comments : Joint Neg_z	+0.069346 (+0.0693)	0.041786 (0.0418)	+1.6596 (+1.66)	0.100730 (0.1007)	[-0.0137, +0.1524]	+0.0502 SD ([-0.0100, +0.1104])
log_comments : Joint Conflict_z	+0.045888 (+0.0459)	0.040235 (0.0402)	+1.1405 (+1.14)	0.257324 (0.2573)	[-0.0341, +0.1259]	+0.0332 SD ([-0.0247, +0.0912])
					[-0.0503, +0.5576] /	

Stage-1 Outcome & Specification	Coefficient (coef)	Two-Way SE	t-Stat	Raw p-Value	95% CI	Standardized β_{std} (95% CI_std)
log_comments : Joint Raw [0,1] Neg / Conf	+0.253677 / +0.167398	0.152857 / 0.146777	+1.66 / +1.14	0.1007 / 0.2573	[-0.1245, +0.4593]	Raw [0,1] scale
log_comments : Standalone Neg_z	**+0.098307 (+0.0983)**	0.029064 (0.0291)	**+3.3825 (+3.38)**	**0.001093 (0.0011)**	**[+0.0405, +0.1561]**	+0.0712 SD ([+0.0293, +0.1131])
log_comments : Standalone Conflict_z	**+0.097089 (+0.0971)**	0.028046 (0.0280)	**+3.4618 (+3.46)**	**0.000846 (0.0008)**	**[+0.0413, +0.1529]**	+0.0703 SD ([+0.0299, +0.1107])

4C. Table 2 — Model 2: Stage-2 Size-Adjusted Structural Deliberative Architecture (RQ2 , N = 2114 , G_domain = 85 , G_week = 261 , 4-Knot log(V_48) RCS)

(Restricted cubic spline knots of $\log V_{48}$ at quantiles [q05, q35, q65, q95] = [2.772589, 3.663562, 4.543295, 5.834811] , corresponding to comment counts [16.0, 39.0, 94.0, 342.0] ; joint VIFs : headline_negativity_z = 2.1736 , moral_conflict_blame_z = 2.2687)

Stage-2 Outcome (within-R2) & Predictor	Coefficient (coef)	Two-Way SE	t-Stat	Raw p-Value	95% CI	St … p_
dyadic_aba_reciprocity_rate_48h (0.0300): Joint Neg_z	+0.003460 (+0.0035)	0.003388 (0.0034)	+1.0215 (+1.02)	0.309967 (0.3100)	[-0.0033, +0.0102]	+
dyadic_aba_reciprocity_rate_48h (0.0300): Joint Conflict_z	**+0.013542 (+0.0135)**	0.004804 (0.0048)	**+2.8193 (+2.82)**	**0.006002 (0.0060)**	**[+0.0040, +0.0231]**	+
dyadic_aba_reciprocity_rate_48h : Raw [0,1] Neg / Conf	+0.012582 / +0.048070	0.012318 / 0.017051	+1.02 / +2.82	0.3100 / 0.0060	[-0.0119, +0.0371] / [+0.0142, +0.0820]	z
interpersonal_nested_reply_share_48h (0.4244): Joint Neg_z	-0.001840 (-0.0018)	0.003072 (0.0031)	-0.5991 (-0.60)	0.550730 (0.5507)	[-0.0079, +0.0043]	-
interpersonal_nested_reply_share_48h (0.4244): Joint Conflict_z	**+0.012726 (+0.0127)**	0.003383 (0.0034)	**+3.7616 (+3.76)**	**0.000311 (0.0003)**	**[+0.0060, +0.0195]**	

Stage-2 Outcome (within-R2) & Predictor	Coefficient (coef)	Two-Way SE	t-Stat	Raw p-Value	95% CI	St p_ +
interpersonal_nested_reply_share_48h : Raw [0,1] Neg / Conf	-0.006691 / +0.045171	0.011168 / 0.012008	-0.60 / +3.76	0.5507 / 0.0003	[-0.0289, +0.0155] / [+0.0213, +0.0691]	z
branch_depth_p90_48h (0.3327): Joint Neg_z	+0.011655 (+0.0117)	0.043107 (0.0431)	+0.2704 (+0.27)	0.787541 (0.7875)	[-0.0741, +0.0974]	+
branch_depth_p90_48h (0.3327): Joint Conflict_z	**+0.181183 (+0.1812)**	0.047627 (0.0476)	**+3.8042 (+3.80)**	**0.000269 (0.0003)**	**[+0.0865, +0.2759]**	+
branch_depth_p90_48h : Raw [0,1] Neg / Conf	+0.042376 / +0.643127	0.156736 / 0.169058	+0.27 / +3.80	0.7875 / 0.0003	[-0.2693, +0.3541] / [+0.3069, +0.9793]	z
external_evidence_link_rate_48h (0.0317): Joint Neg_z	+0.001732 (+0.0017)	0.002342 (0.0023)	+0.7394 (+0.74)	0.461727 (0.4617)	[-0.0029, +0.0064]	+
external_evidence_link_rate_48h (0.0317): Joint Conflict_z	**-0.007289 (-0.0073)**	0.002475 (0.0025)	**-2.9450 (-2.95)**	**0.004177 (0.0042)**	**[-0.0122, -0.0024]**	-
external_evidence_link_rate_48h : Raw [0,1] Neg / Conf	+0.006297 / -0.025872	0.008516 / 0.008785	+0.74 / -2.95	0.4617 / 0.0042	[-0.0106, +0.0232] / [-0.0433, -0.0084]	z
speaker_shannon_entropy_48h (0.0462): Joint Neg_z	-0.000359 (-0.0004)	0.000913 (0.0009)	-0.3932 (-0.39)	0.695187 (0.6952)	[-0.0022, +0.0015]	-
speaker_shannon_entropy_48h (0.0462): Joint Conflict_z	**-0.002865 (-0.0029)**	0.001185 (0.0012)	**-2.4188 (-2.42)**	**0.017730 (0.0177)**	**[-0.0052, -0.0005]**	

Stage-2 Outcome (within-R2) & Predictor	Coefficient (coef)	Two-Way SE	t-Stat	Raw p-Value	95% CI	St V p_ -
speaker_shannon_entropy_48h : Raw [0,1] Neg / Conf	-0.001305 / -0.010171	0.003320 / 0.004205	-0.39 / -2.42	0.6952 / 0.0177	[-0.0079, +0.0053] / [-0.0185, -0.0018]	S z
Aux dyadic_aba_total_share_48h (0.0695): Neg_z / Conf_z	+0.002198 / **+0.013686**	0.002957 / 0.004168	+0.74 / **+3.28**	0.4595 / **0.0015**	[-0.0037, +0.0081] / [+0.0054, +0.0220]	SD
Aux branch_depth_mean_48h (0.3954): Neg_z / Conf_z	-0.001194 / **+0.100994**	0.020969 / 0.023630	-0.06 / **+4.27**	0.9547 / **0.0001**	[-0.0429, +0.0405] / [+0.0540, +0.1480]	SD
Aux mean_log_comment_words_48h (0.0477): Neg_z / Conf_z	-0.010999 / +0.016614	0.011609 / 0.014429	-0.95 / +1.15	0.3461 / 0.2528	[-0.0341, +0.0121] / [-0.0121, +0.0453]	SD

4D. Table 3 — Model 3: Valence-vs-Conflict Functional Bifurcation (RQ3 / H3 Wald Tests & Pre-Registered 0.50 4-Cell Typology)

Stage & Outcome	Continuous H3 Wald (Conf - Neg)_z (t , p)	Cell_1 Pure Threat (Neg>=0.50, Conf<0.50)	Cell_2 Moral Blame (Neg>=0.50, Conf>=0.50)	Cell_3 Non-Neg Debate (Neg<0.50, Conf>=0.50)	4-Cell (Cell_2 - Cell_1) (t , p)
S1 discussed_breakthrough	-0.0217 (-1.07 , 0.2868)	+0.0231 (+1.24 , 0.2171)	+0.0016 (+0.09 , 0.9282)	+0.0257 (+1.03 , 0.3058)	-0.0216 (-1.17 , 0.2467)
S1 log_points	+0.0292 (+0.43 , 0.6669)	**+0.1277 (+2.29 , 0.0243)**	**+0.2387 (+3.89 , 0.0002)**	**+0.2675 (+2.69 , 0.0086)**	+0.1110 (+1.66 , 0.1014)

Stage & Outcome	Continuous H3 Wald (Conf - Neg)_z (t, p)	Cell_1 Pure Threat (Neg>=0.50, Conf<0.50)	Cell_2 Moral Blame (Neg>=0.50, Conf>=0.50)	Cell_3 Non-Neg Debate (Neg<0.50, Conf>=0.50)	4-Cell (Cell_2 - Cell_1) (t, p)
S1 log_comments	-0.0235 (-0.31, 0.7587)	**+0.1535 (+2.06, 0.0422)**	**+0.1237 (+2.00, 0.0486)**	**+0.2516 (+2.25, 0.0273)**	-0.0298 (-0.39, 0.6956)
S2 dyadic_aba_reciprocity_rate_48h	+0.0101 (+1.34, 0.1829)	+0.0135 (+1.79, 0.0767)	**+0.0246 (+2.73, 0.0076)**	**+0.0212 (+2.20, 0.0302)**	+0.0111 (+0.96, 0.3389)
S2 interpersonal_nested_reply_share_48h	**+0.0146 (+2.46, 0.0158)**	**+0.0133 (+2.77, 0.0070)**	**+0.0204 (+3.16, 0.0022)**	**+0.0179 (+2.70, 0.0083)**	+0.0071 (+0.94, 0.3506)
S2 branch_depth_p90_48h	**+0.1695 (+2.09, 0.0395)**	+0.1595 (+1.70, 0.0935)	**+0.3487 (+3.36, 0.0012)**	**+0.2694 (+2.99, 0.0037)**	+0.1892 (+1.56, 0.1228)
S2 external_evidence_link_rate_48h	**-0.0090 (-1.99, 0.0500)**	+0.0004 (+0.10, 0.9204)	-0.0031 (-0.91, 0.3647)	-0.0054 (-1.18, 0.2412)	-0.0035 (-0.67, 0.5057)
S2 speaker_shannon_entropy_48h	-0.0025 (-1.28, 0.2048)	-0.0029 (-1.35, 0.1802)	-0.0040 (-1.84, 0.0688)	-0.0028 (-0.97, 0.3357)	-0.0011 (-0.37, 0.7138)
S2 dyadic_aba_total_share_48h	+0.0115 (+1.75, 0.0845)	**+0.0135 (+2.09, 0.0399)**	**+0.0238 (+3.05, 0.0031)**	**+0.0207 (+2.43, 0.0173)**	+0.0103 (+1.02, 0.3111)
S2 branch_depth_mean_48h	**+0.1022 (+2.56, 0.0121)**	**+0.0971 (+2.12, 0.0366)**	**+0.1797 (+3.45, 0.0009)**	**+0.1494 (+3.11, 0.0026)**	+0.0826 (+1.33, 0.1856)
S2 mean_log_comment_words_48h	+0.0276 (+1.11, 0.2691)	-0.0331 (-1.59, 0.1165)	+0.0028 (+0.18, 0.8590)	-0.0039 (-0.17, 0.8648)	+0.0359 (+1.28, 0.2039)

4E. Table 4 — Model 4: Paywall Information-Starvation Boundary Condition (RQ4 / H4, N = 2114)

- **First-Stage Behavioral Access-Friction Check (Paywall_Mirror_Link_Rate_48h on Paywalled_Domain, N = 2114):** coef = +0.009202 (+0.0092, +0.92pp), SE = 0.001607

(0.0016), **t = +5.7260 (+5.73), p = 0.000000 (0.0000)**, 95% CI = [+0.0060, +0.0124], beta_std = +0.3106 SD ([+0.2028, +0.4185]).

Stage-2 Outcome (N = 2114)	$\gamma_{pw,neg}$ (Neg_z × Paywalled)	$\gamma_{pw,conf}$ (Conf_z × Paywalled)	Open-Access Conf_z Slope	Paywalled Total Conf_z Slope	p_Holm(Neg Conf
dyadic_aba_reciprocity_rate_48h	+0.0014 (t=+0.23 , p=0.8149)	+0.0052 (t=+0.59 , p=0.5540)	**+0.0117 (t=+2.28 , p=0.0252)**	**+0.0169 (t=+2.11 , p=0.0383)**	1.000 **1.00**
interpersonal_nested_reply_share_48h	+0.0107 (t=+2.29 , p=0.0242)	-0.0046 (t=-0.93 , p=0.3565)	**+0.0151 (t=+3.73 , p=0.0003)**	**+0.0104 (t=+2.48 , p=0.0152)**	0.072 **1.00**
external_evidence_link_rate_48h	+0.0001 (t=+0.02 , p=0.9878)	-0.0012 (t=-0.25 , p=0.8034)	-0.0068 (t=-1.90 , p=0.0608)	**-0.0080 (t=-2.51 , p=0.0140)**	1.000 **1.00**
branch_depth_p90_48h	+0.0769 (t=+0.94 , p=0.3475)	-0.0075 (t=-0.08 , p=0.9368)	**+0.1885 (t=+3.39 , p=0.0011)**	**+0.1811 (t=+2.28 , p=0.0250)**	Second
speaker_shannon_entropy_48h	-0.0003 (t=-0.17 , p=0.8635)	-0.0021 (t=-1.03 , p=0.3043)	-0.0021 (t=-1.51 , p=0.1336)	**-0.0042 (t=-2.42 , p=0.0176)**	Second

Section 5. Mandatory 10-Gate Falsification, Selection, Horizon & Out-of-Sample Suite (Gates 1–10)

5A. Gates 1–3: Per-Judge Standalone Regressions, Verbatim URL-Slug Headlines & Heckman / IPW Selection Correction

- **Gate 1 — Per-Judge Joint Model Re-Estimation:**
- j1_llama : Stage-1 Breakthrough (Neg_z b = +0.0163, t = +1.74, p = 0.0857 ; Conf_z b = -0.0014, t = -0.13, p = 0.8977) | Stage-2 ABA (Neg_z b = +0.0039, t = +1.16, p = 0.2507 ; **Conf_z b = +0.0117, t = +2.92, p = 0.0045**).
- j2_qwen : Stage-1 Breakthrough (Neg_z b = +0.0055, t = +0.67, p = 0.5016 ; Conf_z b = -0.0007, t = -0.07, p = 0.9428) | Stage-2 ABA (Neg_z b = +0.0072, t = +2.50, p =

`0.0144`; `Conf_z b = +0.0051, t = +1.32, p = 0.1889` — audited via 3-judge LOJO in `Check E`, where dropping `j2_qwen` yields `Conf3_z = +0.0157, t = +3.22, p = 0.0018`).

- `j3_deepseek`: Stage-1 Breakthrough (`Neg_z b = +0.0106, t = +1.00, p = 0.3208`; `Conf_z b = -0.0024, t = -0.21, p = 0.8318`) | Stage-2 `ABA` (`Neg_z b = +0.0031, t = +0.82, p = 0.4163`; **`Conf_z b = +0.0144, t = +2.66, p = 0.0093`**).
- `j4_gemma`: Stage-1 Breakthrough (`Neg_z b = +0.0220, t = +1.99, p = 0.0497`; `Conf_z b = -0.0125, t = -1.06, p = 0.2923`) | Stage-2 `ABA` (`Neg_z b = +0.0011, t = +0.38, p = 0.7018`; **`Conf_z b = +0.0164, t = +3.53, p = 0.0007`**).
- **Gate 2 — Unaltered Verbatim Newsroom Headlines (`Has_Lexical_Slug == 1` & `Title_Slug_Token_Overlap >= 0.60`):**
- Verbatim retention rates: **`2167/3172` (`68.32%`)** in Stage 1; **`1182/1797` (`65.78%`)** in Stage 2.
- Stage-1 Breakthrough Unedited (`2167` unedited stories): `Neg_z = +0.009591` (`+0.0096`, `SE = 0.0150`, `t = +0.64`, `p = 0.5243`, `95% CI [-0.0203, +0.0394]`); `Conf_z = +0.003034` (`+0.0030`, `SE = 0.0149`, `t = +0.20`, `p = 0.8394`, `95% CI [-0.0267, +0.0327]`).
- Stage-2 `ABA 48h` Unedited (`1182` unedited trees): `Neg_z = +0.000136` (`+0.0001`, `SE = 0.0043`, `t = +0.03`, `p = 0.9748`, `95% CI [-0.0084, +0.0086]`); **`Conf_z = +0.014019` (`+0.0140`, `SE = 0.0058`, `t = +2.40`, `p = 0.0189`, `95% CI [+0.0024, +0.0257]`, `beta_std = +0.1328 SD`)** — eliminating post-hoc submitter/moderator title neutralization as an alternative explanation.
- **Gate 3 — Stage-1 to Stage-2 Selection Bias Audit (`Heckman IMR` & `IPW`, Stage-1 Probit `pseudo-R2 = 0.0436`):**
- `Heckman IMR ABA 48h`: `Neg_z = +0.004736` (`+0.0047`, `t = +1.35`, `p = 0.1804`); **`Conf_z = +0.013682` (`+0.0137`, `SE = 0.0049`, `t = +2.81`, `p = 0.0061`)**; `IMR lambda = +0.022847` (`+0.0228`, `SE = 0.0335`, `t = +0.68`, `p = 0.4975`, `95% CI [-0.0438, +0.0895]`).
- `IPW Breakthrough ABA`: `Neg_z = +0.000927` (`t = +0.24`, `p = 0.8101`); **`Conf_z = +0.017906` (`+0.0179`, `SE = 0.0052`, `t = +3.47`, `p = 0.0008`)**.
- `Combined Universe+Breakthrough IPW ABA`: `Neg_z = +0.000808` (`t = +0.19`, `p = 0.8484`); **`Conf_z = +0.020185` (`+0.0202`, `SE = 0.0061`, `t = +3.33`, `p = 0.0013`)**.

5B. Gates 4–6: Volume Functional Form, Lifecycle Horizon Invariance (`24h` / `48h` / `14d`) & Mirror-Link Stripping

- **Gate 4 — Stage-2 Volume-Control Functional-Form Decomposition:**
- `dyadic_aba_reciprocity_rate_48h`: `4-Knot RCS` (`Neg b = +0.0035, t = +1.02`; **`Conf b = +0.0135, t = +2.82`**) | `20-Vigintile FE` (`Neg b = +0.0037, t = +1.10`;

Conf b = +0.0135, t = +2.75) | Unadjusted (Neg b = +0.0038, t = +1.08; **Conf b = +0.0130, t = +2.70**).

- interpersonal_nested_reply_share_48h: 4-Knot RCS (Neg b = -0.0018, t = -0.60; **Conf b = +0.0127, t = +3.76**) | 20-Vigintile FE (Neg b = -0.0018, t = -0.58; **Conf b = +0.0132, t = +3.76**) | Unadjusted (Neg b = +0.0014, t = +0.32; Conf b = +0.0091, t = +1.88).
- branch_depth_p90_48h: 4-Knot RCS (Neg b = +0.0117, t = +0.27; **Conf b = +0.1812, t = +3.80**) | 20-Vigintile FE (Neg b = +0.0099, t = +0.23; **Conf b = +0.1851, t = +3.77**) | Unadjusted (Neg b = +0.0468, t = +0.72; **Conf b = +0.1483, t = +2.23**).
- external_evidence_link_rate_48h: 4-Knot RCS (Neg b = +0.0017, t = +0.74; **Conf b = -0.0073, t = -2.95**) | 20-Vigintile FE (Neg b = +0.0018, t = +0.80; **Conf b = -0.0074, t = -2.96**) | Unadjusted (Neg b = +0.0014, t = +0.60; **Conf b = -0.0070, t = -2.81**).
- **Gate 5 — Post-Submission Horizon Invariance (24h N = 1996 vs. 48h N = 2114 vs. Full 14d Lifecycle N = 2114):**
- 24h ABA: Neg_z = +0.0053 (t = +1.50, p = 0.1371) | **Conf_z = +0.0129 (SE = 0.0046, t = +2.82, p = 0.0060)**.
- 48h ABA: Neg_z = +0.0035 (t = +1.02, p = 0.3100) | **Conf_z = +0.0135 (SE = 0.0048, t = +2.82, p = 0.0060)**.
- 14d ABA: Neg_z = +0.0011 (t = +0.37, p = 0.7117) | **Conf_z = +0.0135 (SE = 0.0045, t = +2.98, p = 0.0038)**.
- 14d Nested: Neg_z = -0.0014 (t = -0.48, p = 0.6327) | **Conf_z = +0.0122 (SE = 0.0034, t = +3.57, p = 0.0006)**.
- **Gate 6 — A->A Self-Reply & Paywall-Mirror Positive-Control Audit:**
- Mean Paywall_Mirror_Link_Rate_48h: **1.089% (0.010892)** across M = 27 paywalled outlets vs. **0.202% (0.002023)** across M = 58 open-access outlets (**5.38x ratio**); mean A->A self-reply share = 0.28% (562 of 226317 comments = 0.25% aggregate).
- Paywalled_Domain -> Unstripped All Links: coef = +0.005231 (+0.0052, SE = 0.0046, t = +1.13, p = 0.2629, 95% CI [-0.0040, +0.0145]).
- Paywalled_Domain -> Clean Evidence Links (mirror & echo stripped): coef = -0.004619 (-0.0046, SE = 0.0041, t = -1.13, p = 0.2617, 95% CI [-0.0127, +0.0035]) — proving unstripped URL counts conflate archive.* bypass links with documentary citations.

5C. Gates 7–9: Nonlinear / IPW Stage-1 Estimators, Lexical Deconfounding & Leave-One-Out Stability

- **Gate 7 — Universe-Weighted Stage-1 IPW, Logit & Poisson-PPML (`N = 3653`):**
- `Stage-1 IPW Breakthrough`: `Neg_z = +0.002254` (`+0.0023`, `SE = 0.0107`, `t = +0.21`, `p = 0.8340`); `Conf_z = +0.008215` (`+0.0082`, `SE = 0.0124`, `t = +0.66`, `p = 0.5107`).
- `Stage-1 Logit Breakthrough`: `Neg_z log_odds = +0.1446` (`z = +1.52`, `p = 0.1294`); `Conf_z log_odds = +0.0102` (`z = +0.10`, `p = 0.9237`).
- `Stage-1 Poisson PPML Comments`: `Neg_z = -0.0920` (`z = -0.66`, `p = 0.5104`); **`Conf_z = +0.2869` (`z = +2.16`, `p = 0.0311`)**.
- **Gate 8 — Without Lexical Controls, Sample-Median 4-Cell & `1%/99%` Winsorization:**
- `No-Lexical Controls S1 Breakthrough`: `Neg_z = +0.020324` (`+0.0203`, `SE = 0.0101`, `t = +2.01`, `p = 0.0473`); `Conf_z = -0.007671` (`-0.0077`, `SE = 0.0109`, `t = -0.70`, `p = 0.4842`).
- `No-Lexical Controls S2 ABA 48h`: `Neg_z = +0.001447` (`+0.0014`, `SE = 0.0033`, `t = +0.44`, `p = 0.6580`); **`Conf_z = +0.014321` (`+0.0143`, `SE = 0.0044`, `t = +3.28`, `p = 0.0015`)**.
- `Sample-Median 4-Cell (Cell_2 - Cell_1)`: Stage-1 Breakthrough `-0.022586` (`-0.0226`, `t = -0.81`, `p = 0.4225`); **Stage-2 `ABA 48h` `+0.019926` (`+0.0199`, `SE = 0.0097`, `t = +2.05`, `p = 0.0437`)**.
- `Winsorized 1%/99% S2 ABA 48h`: `Neg_z = +0.003670` (`+0.0037`, `t = +1.08`, `p = 0.2837`); **`Conf_z = +0.013322` (`+0.0133`, `SE = 0.0046`, `t = +2.88`, `p = 0.0051`)**.
- **Gate 9 — Leave-One-Domain-Out (`Top 10`), Leave-One-Beat-Out, Leave-One-Year-Out (`2021`..`2025`) & `Post-2023`:**
- **Top-10 Domains Dropped Individually (`theverge.com`, `reuters.com`, `theguardian.com`, `washingtonpost.com`, `arstechnica.com`, `bbc.com`, `nytimes.com`, `theatlantic.com`, `npr.org`, `newyorker.com`):** Stage-2 `ABA Conf_z` remains positive and significant across all ten single-outlet exclusions (`b = +0.0114` to `+0.0144`, `t = +2.54` to `+2.90`).
- **Leave-One-Beat-Out:** Dropping `beat2` (`+0.0148, t = +2.96`), `beat3` (`+0.0156, t = +3.03`), `beat4` (`+0.0158, t = +2.90`), or `beat5` (`+0.0107, t = +2.04`) preserves significance; dropping `beat1_national_intl_wire` yields `Conf_z = +0.0086` (`t = +1.33` — diagnosed via `Check D`).
- **Leave-One-Year-Out (`2021`..`2025`) & `Post-2023` Regime Check:** Stage-2 `ABA Conf_z` holds across all five years (`2021`: `+0.0149, t = +2.78`; `2022`: `+0.0106, t = +2.02`; `2023`: `+0.0137, t = +2.62`; `2024`: `+0.0156, t = +2.84`; `2025`: `+0.0128, t = +2.43`), with null `Post-2023` interactions (`S1 Neg*Post23 t = +1.13, p = 0.2628`, `Conf*Post23 t = -0.91, p = 0.3679`; `S2 ABA Neg*Post23 t = -0.64, p = 0.5219`, `Conf*Post23 t = +0.32, p = 0.7472`).

5D. Table 5 — Gate 10: Strict Chronological Out-of-Sample Evaluation (`Train 2021-2023` → `Test 2024-2025` , `SKILL-018` , `SKILL-020`)

Held-Out OOS Task (`Train` → `Test` , `W_test` Weeks)	Benchmark / Candidate Model	RMSE	MSE / Brier	MAE	R2_ v Con
S1 `discussed_breakthrough` (`2231` → `1422` , `W=105`)	`constant_training_mean`	`0.307230` (`0.3072`)	`0.094390`	`0.193194`	`+0.0`
S1 `discussed_breakthrough`	`controls_only_no_framing`	`0.307120` (`0.3071`)	`0.094323`	`0.193025`	`+0.0`
S1 `discussed_breakthrough`	`lexical_only_zero_llm`	`0.307004` (`0.3070`)	`0.094252`	`0.192935`	`+0.1`
S1 `discussed_breakthrough`	**`full_4judge_llm_jev_shrunk`**	**`0.306780` (`0.3068`)**	**`0.094114`**	**`0.192535`**	**`+0.2`**
S1 `log_comments` (`2231` → `1422` , `W=105`)	`constant_training_mean`	`1.329077` (`1.3291`)	`1.766445`	`0.939266`	`+0.0`
S1 `log_comments`	`controls_only_no_framing`	`1.323174` (`1.3232`)	`1.750790`	`0.934846`	`+0.8`

Held-Out OOS Task (Train → Test, W_test Weeks)	Benchmark / Candidate Model	RMSE	MSE / Brier	MAE	R2_ v Con
S1 log_comments	lexical_only_zero_llm	1.321978 (1.3220)	1.747626	0.933874	+1.(
S1 log_comments	**full_4judge_llm_jev_shrunk**	**1.318048 (1.3180)**	**1.737250**	**0.928592**	**+1.(
S2 dyadic_aba_reciprocity_48h (1355 → 759, W=104)	constant_training_mean	0.107603 (0.1076)	0.011578	0.081674	+0.(
S2 dyadic_aba_reciprocity_48h	controls_only_no_framing	0.107380 (0.1074)	0.011531	0.081484	+0.4
S2 dyadic_aba_reciprocity_48h	lexical_only_zero_llm	0.107334 (0.1073)	0.011521	0.081439	+0.4
S2 dyadic_aba_reciprocity_48h	**full_4judge_llm_jev_shrunk**	**0.107039 (0.1070)**	**0.011457**	**0.081235**	**+1.(

Held-Out OOS Task (Train → Test , W_test Weeks)	Benchmark / Candidate Model	RMSE	MSE / Brier	MAE	R2 v Con
S2 interpersonal_nested_share_48h (1355 → 759 , W=104)	constant_training_mean	0.109017 (0.1090)	0.011885	0.083596	+0.(
S2 interpersonal_nested_share_48h	controls_only_no_framing	0.088710 (0.0887)	0.007869	0.066309	+33.
S2 interpersonal_nested_share_48h	lexical_only_zero_llm	0.088562 (0.0886)	0.007843	0.066214	+34.
S2 interpersonal_nested_share_48h	**full_4judge_llm_jev_shrunk**	**0.087998 (0.0880)**	**0.007744**	**0.065930**	**+34.**

Section 6. Red-Team Adversarial Stress-Test Addendum (`Checks A–F`, `[SECTION 6]` of `VERIFY_STUDY_07_STDOUT.txt`)

6A. Table 6 — Check A: Exact Additive Dyadic `ABA` Decomposition, Speaker Compression & Within-Turn Evidence (`N = 2114`)

(Row-wise identity `Total_ABA_Rate (0.211650) == Unique_Dyad_ABA_Rate (0.145407) + Repeat_PingPong_ABA_Rate (0.066243)` verified to `< 1e-15`; repeat share of baseline `ABA` = `31.30%`; `Conf_z` effect shares: `80.44%` repeat ping-pong vs. `19.56%` first-time unique dyads)

Graph Decomposition Metric (Sample Mean)	Conf_z coef (SE)	Conf_z t & p (95% CI)	Neg_z coef (SE, t, p)	Structural Interpretation
Total_ABA_Rate (21.16%, 0.211650)	**+0.013542 (0.004804)**	**t = +2.82, p = 0.0060** ([+0.0040, +0.0231])	+0.003460 (0.003388, +1.02, 0.3100)	Aggregate reciprocal ABA share (100%)
Unique_Dyad_First_ABA_Rate (14.54%, 0.145407)	+0.002649 (0.002514)	t = +1.05, p = 0.2950 ([-0.0023, +0.0076])	+0.002506 (0.002164, +1.16, 0.2501)	First reciprocal exchange per pair (19.56% share — **null**)
Repeat_PingPong_ABA_Rate (6.62%, 0.066243)	**+0.010893 (0.003478)**	**t = +3.13, p = 0.0024** ([+0.0040, +0.0178])	+0.000955 (0.002166, +0.44, 0.6605)	Repeat back-and-forth within locked pairs (**80.44% share**)
Unique_ABA_Dyads_Per_Speaker	**+0.013027 (0.004574)**	**t = +2.85, p = 0.0055** ([+0.0039, +0.0221])	+0.001473 (0.003373, +0.44, 0.6634)	Rises mechanically because denominator speaker count contracts
Unique_Speakers_Per_Comment (67.84%, 0.678437)	**-0.017397 (0.005210)**	**t = -3.34, p = 0.0013** ([-0.0278, -0.0070])	-0.001489 (0.004012, -0.37, 0.7114)	Fewer distinct citizens generate a given comment volume
Top-2_Speaker_Comment_Share (17.34%, 0.173396)	**+0.007356 (0.003096)**	**t = +2.38, p = 0.0198** ([+0.0012, +0.0135])	+0.000650 (0.002464, +0.26, 0.7926)	Top-2 combatant dominance rises significantly

Graph Decomposition Metric (Sample Mean)	Conf_z coef (SE)	Conf_z t & p (95% CI)	Neg_z coef (SE, t, p)	Structural Interpretation
Within-ABA_Turns_Ev_Rate (N = 2046, mean 7.77%)	-0.004461 (0.005338)	t = -0.84, p = 0.4057 ([-0.0151, +0.0062])	Co-estimated in Model 2	ABA turns carry lower baseline citing (7.77% vs. 10.28%)
Within-Non-ABA_Turns_Ev_Rate (N = 2114, mean 10.28%)	**-0.007916 (0.002637)**	**t = -3.00, p = 0.0035** ([-0.0132, -0.0027])	Co-estimated in Model 2	Direct non-compositional documentary erosion outside ABA turns

6B. Table 7 — Check B: Stage-1 Multi-Threshold Crowd Gatekeeping Hurdle Curve (N = 3653)

Stage-1 Comment Threshold (Count, Base Rate)	Joint Neg_z coef (SE, t, p)	Joint Conf_z coef (SE, t, p)	Standalone Neg_z (b, t, p)	Standalone Conf_z (b, t, p)
num_comments >= 1 (n = 1580, 43.25%)	**+0.0306 (0.0120, +2.54, 0.0129)**	**+0.0300 (0.0128, +2.35, 0.0211)**	**+0.0495 (+5.47, 0.0000)**	**+0.0526 (+5.41, 0.0000)**
num_comments >= 5 (n = 598, 16.37%)	+0.0053 (0.0105, +0.50, 0.6179)	+0.0185 (0.0101, +1.84, 0.0700)	**+0.0169 (+2.11, 0.0376)**	**+0.0224 (+2.84, 0.0056)**
num_comments >= 10 (n = 451, 12.35%)	+0.0129 (0.0104, +1.24, 0.2188)	+0.0016 (0.0106, +0.15, 0.8810)	+0.0139 (+1.76, 0.0822)	+0.0111 (+1.37, 0.1751)
num_comments >= 15 (n = 398, 10.90%)	+0.0158 (0.0104, +1.52, 0.1315)	-0.0059 (0.0115, -0.51, 0.6079)	+0.0120 (+1.65, 0.1022)	+0.0057 (+0.70, 0.4829)
num_comments >= 30 (n = 314, 8.60%)	+0.0153 (0.0094, +1.62, 0.1085)	-0.0085 (0.0103, -0.83, 0.4113)	+0.0100 (+1.60, 0.1131)	+0.0028 (+0.41, 0.6793)

Stage-1 Comment Threshold (Count, Base Rate)	Joint Neg_z coef (SE, t, p)	Joint Conf_z coef (SE, t, p)	Standalone Neg_z (b, t, p)	Standalone Conf_z (b, t, p)
num_comments >= 50 (n = 241, 6.60%)	+0.0128 (0.0087, +1.48, 0.1417)	-0.0089 (0.0090, -0.99, 0.3258)	+0.0072 (+1.35, 0.1802)	+0.0006 (+0.10, 0.9193)
num_comments >= 100 (n = 138, 3.78%)	+0.0040 (0.0057, +0.70, 0.4875)	+0.0030 (0.0064, +0.47, 0.6405)	+0.0059 (+1.45, 0.1517)	+0.0059 (+1.31, 0.1935)
num_comments >= 150 (n = 88, 2.41%)	-0.0041 (0.0046, -0.89, 0.3770)	+0.0085 (0.0048, +1.80, 0.0761)	+0.0013 (+0.42, 0.6763)	+0.0055 (+1.76, 0.0828)
Intensive log_comments | >= 1 (N_eff = 1579)	+0.0090 (0.0713, +0.13, 0.9000)	-0.0078 (0.0714, -0.11, 0.9137)	+0.0041 (+0.08)	-0.0014 (-0.03)

6C. Table 8 — Check C: H3 4-Cell Cutoff Grid Sensitivity (0.35, 0.40, 0.50, 0.60) & Cell_1 Sub-Threshold Conflict Contamination (N = 2114)

Binary Cutoff (n0, n1, n2, n3)	Cell_1 Mean Conf (% >= 0.30)	Cell_1 Pure Threat (t, p)	Cell_2 Moral Blame (t, p)	Cell_3 Non-Neg Debate (t, p)	Contrast Cell_2 - Cell_1 (SE, t, p)
cut = 0.35 (962, 334, 701, 117)	0.1717 (15.27%)	+0.0029 (+0.46, 0.6476)	**+0.0302 (+4.14, 0.0001)**	+0.0213 (+1.77, 0.0805)	**+0.0273 (0.0087, +3.15, 0.0023)**
cut = 0.40 (1030, 336, 612, 136)	0.1953 (25.60%)	+0.0094 (+1.41, 0.1615)	**+0.0280 (+3.82, 0.0003)**	**+0.0228 (+2.00, 0.0490)**	**+0.0187 (0.0089, +2.10, 0.0385)**
cut = 0.50 (1198, 314, 413, 189)	0.2431 (41.08%)	+0.0135 (+1.78, 0.0783)	**+0.0246 (+2.74, 0.0076)**	**+0.0212 (+2.21, 0.0302)**	+0.0111 (0.0116,

Binary Cutoff (n0, n1, n2, n3)	Cell_1 Mean Conf (% >= 0.30)	Cell_1 Pure Threat (t, p)	Cell_2 Moral Blame (t, p)	Cell_3 Non-Neg Debate (t, p)	Contrast Cell_2 - Cell_1 (SE, t, p)
					+0.96, 0.3382)
cut = 0.60 (1401, 254, 261, 198)	0.2844 (48.03%)	**+0.0202 (+2.63, 0.0101)**	+0.0200 (+1.74, 0.0856)	**+0.0254 (+2.82, 0.0060)**	-0.0002 (0.0142, -0.02, 0.9871)

6D. Check D: Journalistic Beat Domain Boundary & Gelman–Stern Full-Panel Interactions

- **Subsample Estimates:**

- Civic & Public-Affairs (Beats 1+2) (N_eff = 962, M_domain = 41, df_t = 40, SD_Conf = 0.2900): **ABA Conf_z = +0.0165 (SE = 0.0068, t = +2.43, p = 0.0195)**, Neg_z = +0.0037 (t = +0.87); **Nested Conf_z = +0.0132 (t = +2.92, p = 0.0058)**; Evidence Conf_z = -0.0067 (t = -1.56, p = 0.1261).
- Specialist Sci/Tech/Culture (Beats 3+4+5) (N_init = 1152, N_eff = 1151, M_domain = 44, df_t = 43, SD_Conf = 0.2581): ABA Conf_z = +0.0102 (SE = 0.0073, t = +1.40, p = 0.1694), Neg_z = +0.0030 (t = +0.67); **Nested Conf_z = +0.0112 (t = +2.30, p = 0.0267)**; Evidence Conf_z = -0.0060 (t = -1.92, p = 0.0609).
- Beat 1 National/Intl/Wire Only (N_eff = 831, M_domain = 27, df_t = 26, SD_Conf = 0.2881): **ABA Conf_z = +0.0199 (SE = 0.0072, t = +2.76, p = 0.0105)**; **Nested Conf_z = +0.0140 (t = +2.79, p = 0.0097)**; Evidence Conf_z = -0.0067 (t = -1.48).
- Beat 4 Tech-in-Society Only (N_init = 459, N_eff = 458, M_domain = 11, df_t = 10, SD_Conf = 0.2734): ABA Conf_z = -0.0031 (SE = 0.0121, t = -0.26, p = 0.8014, 95% CI [-0.0300, +0.0237]); Nested Conf_z = +0.0122 (t = +1.38); Evidence Conf_z = -0.0064 (t = -1.53).

- **Full-Panel Gelman–Stern Slope Interactions (N = 2114, M_domain = 85):**

- ABA: Specialist Conf_z = +0.0094 (t = +1.44, p = 0.1542), Conf_z × Civic = +0.0066 (SE = 0.0081, t = +0.82, p = 0.4173), **Civic Total = +0.0161 (t = +2.71, p = 0.0081)**; Non-Beat4 Conf_z = +0.0155 (t = +3.02, p = 0.0033), Conf_z × Beat4_Tech = -0.0079 (SE = 0.0086, t = -0.92, p = 0.3599), Beat4 Total = +0.0076 (t = +0.96, p = 0.3406).

- Nested: **Specialist Conf_z = +0.0099 (t = +2.10, p = 0.0384)**, Conf_z × Civic = +0.0049 (t = +0.92, p = 0.3585), **Civic Total = +0.0148 (t = +3.68, p = 0.0004)**.
- Evidence: **Specialist Conf_z = -0.0066 (t = -2.18, p = 0.0320)**, Conf_z × Civic = -0.0014 (t = -0.28, p = 0.7789), **Civic Total = -0.0080 (t = -2.09, p = 0.0395)**.

6E. Check E: Leave-One-Judge-Out (`3-Judge JEV`) & Two-Part Zero-Hurdle Diagnostics

- **3-Judge LOJO on Stage-2 `ABA 48h` (`N = 2114`):**
- drop_j1_llama: Neg3_z = +0.0035 (t = +1.06, p = 0.2940) | **Conf3_z = +0.0134 (SE = 0.0050, t = +2.68, p = 0.0089)**.
- drop_j2_qwen: Neg3_z = +0.0023 (t = +0.66, p = 0.5116) | **Conf3_z = +0.0157 (SE = 0.0049, t = +3.22, p = 0.0018)**.
- drop_j3_deepseek: Neg3_z = +0.0037 (t = +1.14, p = 0.2582) | **Conf3_z = +0.0126 (SE = 0.0045, t = +2.82, p = 0.0060)**.
- drop_j4_gemma: Neg3_z = +0.0047 (t = +1.36, p = 0.1787) | **Conf3_z = +0.0114 (SE = 0.0047, t = +2.45, p = 0.0162)**.
- **Two-Part `Any_Conf > 0.05` Hurdle & Positive-Conflict Subsample:**
- Stage-1 Breakthrough Two-Part Hurdle: Any_Conf > 0.05 = +0.0273 (SE = 0.0169, t = +1.62, p = 0.1100), Conf_z = -0.0120 (t = -1.11, p = 0.2688).
- Stage-2 ABA Two-Part Hurdle: **Any_Conf > 0.05 = +0.0302 (SE = 0.0089, t = +3.39, p = 0.0011)**, Conf_z = +0.0072 (SE = 0.0052, t = +1.39, p = 0.1677).
- Stage-2 ABA Positive-Conflict Subsample (Conf > 0.05, N_eff = 1389 after variance truncation of the 34.30% zero-conflict baseline): Neg_z = -0.0035 (t = -0.70, p = 0.4879), Conf_z = +0.0074 (SE = 0.0055, t = +1.34, p = 0.1848).

6F. Table 9 — Check F: Strictly Ex-Ante `t = 0` Stage-2 Out-of-Sample Evaluation (`N_train = 1355` → `N_test = 759`, `W_test = 104`, Zero Realized `log(V_48)` Splines, Strict `2021-2022` Inner Scaling)

Pure t = 0 Ex-Ante Stage-2 OOS Task	Constant RMSE (MAE)	Controls-Only R2_OOS	Lexical-Only (Zero LLM) R2_OOS	4-Judge LLM JEV RMSE (R2_OOS)	DM vs. Constant (t, p)	DM vs. Controls (t, p)	L (
stage2_aba_ex_ante_t0	0.107603 (0.081674)	+0.4086%	+0.4997%	**0.107034 (+1.0542%)**	-1.79 (0.0761)	**-2.35 (0.0206)**	(

Pure t = 0 Ex-Ante Stage-2 OOS Task	Constant RMSE (MAE)	Controls-Only R2_OOS	Lexical-Only (Zero LLM) R2_OOS	4-Judge LLM JEV RMSE (R2_OOS)	DM vs. Constant (t, p)	DM vs. Controls (t, p)	L
stage2_nested_ex_ante_t0	0.109017 (0.083596)	+1.6946%	+1.8861%	0.107757 (+2.2984%)	-3.49 (0.0007)	-3.21 (0.0018)	(
stage2_evidence_ex_ante_t0	0.067817 (0.050846)	+3.5737%	+3.5924%	0.066372 (+4.2182%)	-4.93 (0.0000)	-2.48 (0.0149)	(

Section 7. Authoritative Reference Verification Stdout (`results/VERIFY_STUDY_07_STDOUT.txt`, `SKILL-027`)

7A. Verbatim Reference Verification Output (`SHA-256: 5d1af3e4428adfc76bc776e44474ed012b65862cadbb510186c6ce90552db901`)

```
 ================================================================================
===========
STUDY #07 DETERMINISTIC END-TO-END VERIFICATION & REFERENCE STDOUT (SKILL-027)
Project: projects/virality-deliberation (study-07 | 100% Pure Communication & Media
Science)
================================================================================
===========

[SECTION 1] CRYPTOGRAPHIC PROVENANCE & FULL 64-HEX ARTIFACT SHA-256 REGISTRY
  Pre-Data Git Commit Verified: 131fa8a (git cat-file -t -> commit) [PASS]
  Immutable Pre-Registration Lock Byte-Identical: True [PASS]
  SHA256(STUDY_DESIGN_CARD_V1_PRE_DATA_LOCKED.md) =
7c25286e6b930f06560b29fe50b3bd023c8d55ef66611dfd2990bd70430362ab
  SHA256(STUDY_DESIGN_CARD.md) =
7c25286e6b930f06560b29fe50b3bd023c8d55ef66611dfd2990bd70430362ab
  SHA256(LITERATURE_SYNTHESIS_VERIFIED.json) =
097aa728e6f8ec5c8de277db0e7ded5a3d2878a20972d643b5ea351575073c05
  SHA256(SWARM_DISCOVERY_AND_FEASIBILITY_LEDGER.md) =
aaa9b7ca1f2e1a1d0c9e6c215533c11f66bc98c7c2a8953f1a49693388fb1b5a
  SHA256(scripts/build_dataset.py) =
a6d77edb84fa83b466b4efe4cf5db20cc31478df0115621c941e47a071e631dd
  SHA256(scripts/run_econometrics_and_stress_tests.py) =
cf4af2be6144e955b1e8ba816380b12ddfe97f8402b21c68d05b0a64b5a2238b
  SHA256(scripts/verify_study_07.py) =
2975721f2bd33dedce6cf70892a2d4364157a3593d70d3523e1b5b2b502e7a0d
```

```
  SHA256(data/stage1_submissions_panel.csv) =
b456c545440250f853fdba8fec64f72fc17e5afa06a941c5f7cf1872831cbd54
  SHA256(data/stage2_deliberation_trees_panel.csv) =
fe66ee05068befcb26cf91360f3c6ba15b87b48b001c3b5d82568d8f76de4457
  SHA256(data/raw_judge_ballots.json) =
0cd9e66a64460b77484f77b8b04717d4a5f36d2788b4012e559366e7bf9c2d66
  SHA256(data/headline_lexicons.json) =
4de109c0b49625c620ce3897b8c5de74fd03848c29a27a3d1099932b5b7e600c
  SHA256(data/sample_attrition_and_diagnostics.json) =
a0481d13da01a1ddf6890f707374e8b4d5c16aee1fd169841a0e8a2b388a1aac
  SHA256(results/econometric_results.json) =
2ddf0ff9fd70366331ffc00201ddb70eacfc9a047094575d79fa31b9c36b5665
  SHA256(results/chart_series_summary.json) =
bbffc16596804d748aa257f4f65eab474de92534334a5abd65aa44e40836d532

[SECTION 2] CURATED JOURNALISM UNIVERSE, ATTRITION WATERFALL & DISTRIBUTIONS
(2021-2025)
  Curated Non-Financial Outlets Probed: M_probed = 85 -> Qualifying (>=5 stories): M
= 85
  Algolia Hits Scanned: N = 221497 -> Drop host/fin=294, dup=0, missing=0, meta=14,
fin_headline_regex=493 -> Clean Universe: N = 220696 (Discussed >=15 comments: N =
23453 across M = 85 outlets)
  Stage-1 Unconditional Crowd Gatekeeping Panel: N = 3653, G_domain = 85 (M = 85),
G_week = 262, Stage-1 Discussed Entering Stage 2: N = 398, Sum IPW = 220696.00 [PASS]
  Stage-2 Deliberative Tree Panel: Candidates = 2251 -> HTTP 200 = 2251 (drop HTTP
fail = 0, drop <10 valid 48h comments = 137) -> Retained Trees: N = 2114, G_domain =
85 (M = 85), G_week = 261, Valid 48h Comments = 226317, Valid 14d Comments = 235510,
A->A Self-Replies Excluded (48h) = 562 (0.25%) [PASS]
  4-Judge Live OpenRouter Ballots: Unique Headlines = 5398, Total Live Ballots =
21592, Synthetic/Fallback = 0 [PASS]
  Dist[stage1_log_points]: mean=+1.8228, sd=1.1932, min=+0.6931, p01=+0.6931,
median=+1.3863, p99=+5.7230, max=+7.2056, max|z|=4.51
  Dist[stage1_log_comments]: mean=+0.8259, sd=1.3806, min=+0.0000, p01=+0.0000,
median=+0.0000, p99=+5.6112, max=+7.4645, max|z|=4.81
  Dist[stage2_dyadic_aba_reciprocity_rate_48h]: mean=+0.2117, sd=0.1053, min=+0.0000,
p01=+0.0000, median=+0.2042, p99=+0.5207, max=+0.8000, max|z|=5.59
  Dist[stage2_interpersonal_nested_reply_share_48h]: mean=+0.7982, sd=0.1111,
min=+0.3000, p01=+0.4384, median=+0.8234, p99=+0.9511, max=+0.9677, max|z|=4.48
  Dist[stage2_branch_depth_p90_48h]: mean=+5.2674, sd=1.6128, min=+2.0000,
p01=+2.1000, median=+5.0000, p99=+9.9000, max=+16.0000, max|z|=6.65
  Dist[stage2_external_evidence_link_rate_48h]: mean=+0.0991, sd=0.0650, min=+0.0000,
p01=+0.0000, median=+0.0874, p99=+0.3169, max=+0.4667, max|z|=5.65
  Dist[stage2_speaker_shannon_entropy_48h]: mean=+0.9548, sd=0.0269, min=+0.7200,
p01=+0.8620, median=+0.9598, p99=+0.9935, max=+1.0000, max|z|=8.71

[SECTION 3] SHROUT & FLEISS (1979) 4-JUDGE RELIABILITY & 4-CELL TYPOLOGY COUNTS
```

```
(SKILL-019 / SKILL-022)
  Stage 1 (N_stage1) [headline_negativity] (N = 3653): ICC(1,4) = 0.9664, ICC(2,4) =
0.9665, ICC(3,4) = 0.9682, ICC(2,1) = 0.8782, pairwise_r = 0.8886, mean = 0.2944, SD
= 0.2734, floor==0 = 20.67%, floor<=0.05 = 33.42%, ceiling>=0.95 = 0.08% [PASS]
    Judge j1_llama (meta-llama/llama-3.3-70b-instruct): mean = 0.3043, SD = 0.3129
    Judge j2_qwen (qwen/qwen-2.5-72b-instruct): mean = 0.2611, SD = 0.2576
    Judge j3_deepseek (deepseek/deepseek-chat-v3-0324): mean = 0.2999, SD = 0.2896
    Judge j4_gemma (google/gemma-3-27b-it): mean = 0.3123, SD = 0.2817
  Stage 1 (N_stage1) [moral_conflict_blame] (N = 3653): ICC(1,4) = 0.9649, ICC(2,4) =
0.9650, ICC(3,4) = 0.9655, ICC(2,1) = 0.8732, pairwise_r = 0.8768, mean = 0.2465, SD
= 0.2741, floor==0 = 31.37%, floor<=0.05 = 43.33%, ceiling>=0.95 = 0.00% [PASS]
    Judge j1_llama (meta-llama/llama-3.3-70b-instruct): mean = 0.2381, SD = 0.2934
    Judge j2_qwen (qwen/qwen-2.5-72b-instruct): mean = 0.2371, SD = 0.2781
    Judge j3_deepseek (deepseek/deepseek-chat-v3-0324): mean = 0.2448, SD = 0.2745
    Judge j4_gemma (google/gemma-3-27b-it): mean = 0.2658, SD = 0.3048
  Stage 1 (N_stage1) Corr(Neg, Conflict): r = +0.7533 | 4-Cell (0.50) Counts =
{'Cell_0_Neutral_Constructive': 2342, 'Cell_2_Moral_Conflict_Blame': 568,
'Cell_3_NonNegative_Debate': 302, 'Cell_1_Pure_Threat_Tragedy_Negativity': 441} |
Median Cutoffs = {'neg_median': 0.25, 'conf_median': 0.125} -> Counts =
{'Cell_0_Neutral_Constructive': 1547, 'Cell_2_Moral_Conflict_Blame': 1591,
'Cell_3_NonNegative_Debate': 276, 'Cell_1_Pure_Threat_Tragedy_Negativity': 239}
  Stage 2 (N_trees) [headline_negativity] (N = 2114): ICC(1,4) = 0.9637, ICC(2,4) =
0.9638, ICC(3,4) = 0.9659, ICC(2,1) = 0.8694, pairwise_r = 0.8806, mean = 0.3460, SD
= 0.2750, floor==0 = 15.14%, floor<=0.05 = 24.41%, ceiling>=0.95 = 0.05% [PASS]
    Judge j1_llama (meta-llama/llama-3.3-70b-instruct): mean = 0.3589, SD = 0.3157
    Judge j2_qwen (qwen/qwen-2.5-72b-instruct): mean = 0.3091, SD = 0.2653
    Judge j3_deepseek (deepseek/deepseek-chat-v3-0324): mean = 0.3478, SD = 0.2874
    Judge j4_gemma (google/gemma-3-27b-it): mean = 0.3681, SD = 0.2843
  Stage 2 (N_trees) [moral_conflict_blame] (N = 2114): ICC(1,4) = 0.9623, ICC(2,4) =
0.9623, ICC(3,4) = 0.9629, ICC(2,1) = 0.8645, pairwise_r = 0.8684, mean = 0.2903, SD
= 0.2817, floor==0 = 23.37%, floor<=0.05 = 34.30%, ceiling>=0.95 = 0.00% [PASS]
    Judge j1_llama (meta-llama/llama-3.3-70b-instruct): mean = 0.2805, SD = 0.3027
    Judge j2_qwen (qwen/qwen-2.5-72b-instruct): mean = 0.2815, SD = 0.2927
    Judge j3_deepseek (deepseek/deepseek-chat-v3-0324): mean = 0.2876, SD = 0.2798
    Judge j4_gemma (google/gemma-3-27b-it): mean = 0.3117, SD = 0.3117
  Stage 2 (N_trees) Corr(Neg, Conflict): r = +0.7176 | 4-Cell (0.50) Counts =
{'Cell_0_Neutral_Constructive': 1198, 'Cell_2_Moral_Conflict_Blame': 412,
'Cell_1_Pure_Threat_Tragedy_Negativity': 315, 'Cell_3_NonNegative_Debate': 189} |
Median Cutoffs = {'neg_median': 0.325, 'conf_median': 0.2} -> Counts =
{'Cell_0_Neutral_Constructive': 820, 'Cell_2_Moral_Conflict_Blame': 890,
'Cell_3_NonNegative_Debate': 207, 'Cell_1_Pure_Threat_Tragedy_Negativity': 197}

[SECTION 4] CONFIRMATORY MODELS 1-4 & HOLM-BONFERRONI MULTIPLICITY ADJUDICATION
  Live Spot Re-Derivation Check (Stage-1 Breakthrough & Stage-2 ABA Reciprocity):
EXACT MATCH (< 1e-9) [PASS]
  Adjudicated Strong-Inference Mechanism: Mixed_Or_Boundary_Specific_Equilibrium
```

```
  Primary Holm-Bonferroni Family (K = 4): p_Holm(S1 beta1_neg) = 0.394568 (0.3946),
p_Holm(S1 beta1_conf) = 0.619933 (0.6199), p_Holm(S2 beta2_neg) = 0.619933 (0.6199),
p_Holm(S2 beta2_conf) = 0.024010 (0.0240)

  --- MODEL 1: Stage-1 Unconditional Crowd Gatekeeping & Viral Breakthrough (RQ1) ---
  Stage-1 Outcome [discussed_breakthrough] (N = 3653, singletons=0, G_domain = 85,
G_week = 262, within-R2 = 0.0062, VIFs = {'headline_negativity_z': 2.3527,
'moral_conflict_blame_z': 2.4105}):
    Joint Neg_z:            coef = +0.015772 (+0.0158, +1.58pp), SE = 0.010356
(0.0104), t = +1.5230 (+1.52), p = 0.131523 (0.1315), 95% CI = [-0.0048, +0.0364] |
beta_std = +0.0506 SD (SE_std = 0.0332, 95% CI_std = [-0.0155, +0.1167])
    Joint Conflict_z:       coef = -0.005923 (-0.0059, -0.59pp), SE = 0.011502
(0.0115), t = -0.5150 (-0.51), p = 0.607913 (0.6079), 95% CI = [-0.0288, +0.0169] |
beta_std = -0.0190 SD (SE_std = 0.0369, 95% CI_std = [-0.0924, +0.0544])
    Joint Raw[0,1] Neg:     coef = +0.057694 (+0.0577, +5.77pp), SE = 0.037883
(0.0379), t = +1.5230 (+1.52), p = 0.131523 (0.1315), 95% CI = [-0.0176, +0.1330]
    Joint Raw[0,1] Conf:    coef = -0.021609 (-0.0216, -2.16pp), SE = 0.041960
(0.0420), t = -0.5150 (-0.51), p = 0.607913 (0.6079), 95% CI = [-0.1051, +0.0618]
    Standalone Neg_z:       coef = +0.012033 (+0.0120, +1.20pp), SE = 0.007282
(0.0073), t = +1.6524 (+1.65), p = 0.102181 (0.1022), 95% CI = [-0.0024, +0.0265] |
beta_std = +0.0386 SD (SE_std = 0.0234, 95% CI_std = [-0.0079, +0.0851])
    Standalone Conflict_z:  coef = +0.005721 (+0.0057, +0.57pp), SE = 0.008118
(0.0081), t = +0.7047 (+0.70), p = 0.482929 (0.4829), 95% CI = [-0.0104, +0.0219] |
beta_std = +0.0184 SD (SE_std = 0.0261, 95% CI_std = [-0.0334, +0.0702])
    H3 Wald (Conf - Neg)_z: coef = -0.021695 (-0.0217, -2.17pp), SE = 0.020237
(0.0202), t = -1.0721 (-1.07), p = 0.286765 (0.2868), 95% CI = [-0.0619, +0.0185]
    4-Cell Cell_1 Threat:   coef = +0.023143 (+0.0231, +2.31pp), SE = 0.018611
(0.0186), t = +1.2435 (+1.24), p = 0.217140 (0.2171), 95% CI = [-0.0139, +0.0602]
    4-Cell Cell_2 Blame:    coef = +0.001588 (+0.0016, +0.16pp), SE = 0.017566
(0.0176), t = +0.0904 (+0.09), p = 0.928200 (0.9282), 95% CI = [-0.0333, +0.0365]
    4-Cell Cell_3 Debate:   coef = +0.025660 (+0.0257, +2.57pp), SE = 0.024904
(0.0249), t = +1.0304 (+1.03), p = 0.305796 (0.3058), 95% CI = [-0.0239, +0.0752]
    4-Cell (Cell_2-Cell_1): coef = -0.021555 (-0.0216, -2.16pp), SE = 0.018477
(0.0185), t = -1.1666 (-1.17), p = 0.246688 (0.2467), 95% CI = [-0.0583, +0.0152]
  Stage-1 Outcome [log_points] (N = 3653, singletons=0, G_domain = 85, G_week = 262,
within-R2 = 0.0203, VIFs = {'headline_negativity_z': 2.3527,
'moral_conflict_blame_z': 2.4105}):
    Joint Neg_z:            coef = +0.057924 (+0.0579, +5.79pp), SE = 0.034976
(0.0350), t = +1.6561 (+1.66), p = 0.101434 (0.1014), 95% CI = [-0.0116, +0.1275] |
beta_std = +0.0485 SD (SE_std = 0.0293, 95% CI_std = [-0.0097, +0.1068])
    Joint Conflict_z:       coef = +0.087165 (+0.0872, +8.72pp), SE = 0.038071
(0.0381), t = +2.2895 (+2.29), p = 0.024556 (0.0246), 95% CI = [+0.0115, +0.1629] |
beta_std = +0.0731 SD (SE_std = 0.0319, 95% CI_std = [+0.0096, +0.1365])
    Joint Raw[0,1] Neg:     coef = +0.211894 (+0.2119, +21.19pp), SE = 0.127948
(0.1279), t = +1.6561 (+1.66), p = 0.101434 (0.1014), 95% CI = [-0.0425, +0.4663]
    Joint Raw[0,1] Conf:    coef = +0.317977 (+0.3180, +31.80pp), SE = 0.138882
```

```
(0.1389), t = +2.2895 (+2.29), p = 0.024556 (0.0246), 95% CI = [+0.0418, +0.5942]
    Standalone Neg_z:       coef = +0.112936 (+0.1129, +11.29pp), SE = 0.024190
(0.0242), t = +4.6687 (+4.67), p = 0.000011 (0.0000), 95% CI = [+0.0648, +0.1610] |
beta_std = +0.0946 SD (SE_std = 0.0203, 95% CI_std = [+0.0543, +0.1350])
    Standalone Conflict_z:  coef = +0.129933 (+0.1299, +12.99pp), SE = 0.026641
(0.0266), t = +4.8773 (+4.88), p = 0.000005 (0.0000), 95% CI = [+0.0770, +0.1829] |
beta_std = +0.1089 SD (SE_std = 0.0223, 95% CI_std = [+0.0645, +0.1533])
    H3 Wald (Conf - Neg)_z: coef = +0.029241 (+0.0292, +2.92pp), SE = 0.067695
(0.0677), t = +0.4319 (+0.43), p = 0.666887 (0.6669), 95% CI = [-0.1054, +0.1639]
    4-Cell Cell_1 Threat:   coef = +0.127735 (+0.1277, +12.77pp), SE = 0.055701
(0.0557), t = +2.2932 (+2.29), p = 0.024333 (0.0243), 95% CI = [+0.0170, +0.2385]
    4-Cell Cell_2 Blame:    coef = +0.238719 (+0.2387, +23.87pp), SE = 0.061364
(0.0614), t = +3.8902 (+3.89), p = 0.000200 (0.0002), 95% CI = [+0.1167, +0.3607]
    4-Cell Cell_3 Debate:   coef = +0.267479 (+0.2675, +26.75pp), SE = 0.099363
(0.0994), t = +2.6919 (+2.69), p = 0.008572 (0.0086), 95% CI = [+0.0699, +0.4651]
    4-Cell (Cell_2-Cell_1): coef = +0.110983 (+0.1110, +11.10pp), SE = 0.067012
(0.0670), t = +1.6562 (+1.66), p = 0.101421 (0.1014), 95% CI = [-0.0223, +0.2442]
  Stage-1 Outcome [log_comments] (N = 3653, singletons=0, G_domain = 85, G_week =
262, within-R2 = 0.0125, VIFs = {'headline_negativity_z': 2.3527,
'moral_conflict_blame_z': 2.4105}):
    Joint Neg_z:            coef = +0.069346 (+0.0693, +6.93pp), SE = 0.041786
(0.0418), t = +1.6596 (+1.66), p = 0.100730 (0.1007), 95% CI = [-0.0137, +0.1524] |
beta_std = +0.0502 SD (SE_std = 0.0303, 95% CI_std = [-0.0100, +0.1104])
    Joint Conflict_z:       coef = +0.045888 (+0.0459, +4.59pp), SE = 0.040235
(0.0402), t = +1.1405 (+1.14), p = 0.257324 (0.2573), 95% CI = [-0.0341, +0.1259] |
beta_std = +0.0332 SD (SE_std = 0.0291, 95% CI_std = [-0.0247, +0.0912])
    Joint Raw[0,1] Neg:     coef = +0.253677 (+0.2537, +25.37pp), SE = 0.152857
(0.1529), t = +1.6596 (+1.66), p = 0.100730 (0.1007), 95% CI = [-0.0503, +0.5576]
    Joint Raw[0,1] Conf:    coef = +0.167398 (+0.1674, +16.74pp), SE = 0.146777
(0.1468), t = +1.1405 (+1.14), p = 0.257324 (0.2573), 95% CI = [-0.1245, +0.4593]
    Standalone Neg_z:       coef = +0.098307 (+0.0983, +9.83pp), SE = 0.029064
(0.0291), t = +3.3825 (+3.38), p = 0.001093 (0.0011), 95% CI = [+0.0405, +0.1561] |
beta_std = +0.0712 SD (SE_std = 0.0211, 95% CI_std = [+0.0293, +0.1131])
    Standalone Conflict_z:  coef = +0.097089 (+0.0971, +9.71pp), SE = 0.028046
(0.0280), t = +3.4618 (+3.46), p = 0.000846 (0.0008), 95% CI = [+0.0413, +0.1529] |
beta_std = +0.0703 SD (SE_std = 0.0203, 95% CI_std = [+0.0299, +0.1107])
    H3 Wald (Conf - Neg)_z: coef = -0.023458 (-0.0235, -2.35pp), SE = 0.076130
(0.0761), t = -0.3081 (-0.31), p = 0.758739 (0.7587), 95% CI = [-0.1749, +0.1279]
    4-Cell Cell_1 Threat:   coef = +0.153498 (+0.1535, +15.35pp), SE = 0.074408
(0.0744), t = +2.0629 (+2.06), p = 0.042209 (0.0422), 95% CI = [+0.0055, +0.3015]
    4-Cell Cell_2 Blame:    coef = +0.123740 (+0.1237, +12.37pp), SE = 0.061842
(0.0618), t = +2.0009 (+2.00), p = 0.048634 (0.0486), 95% CI = [+0.0008, +0.2467]
    4-Cell Cell_3 Debate:   coef = +0.251624 (+0.2516, +25.16pp), SE = 0.111982
(0.1120), t = +2.2470 (+2.25), p = 0.027261 (0.0273), 95% CI = [+0.0289, +0.4743]
    4-Cell (Cell_2-Cell_1): coef = -0.029759 (-0.0298, -2.98pp), SE = 0.075803
(0.0758), t = -0.3926 (-0.39), p = 0.695625 (0.6956), 95% CI = [-0.1805, +0.1210]
```

```
  --- MODEL 2: Stage-2 Size-Adjusted Structural Deliberative Architecture (RQ2, 4-
Knot log(V_48) RCS) ---
  Spline Knots log(V_48) [q05, q35, q65, q95]: [2.772589, 3.663562, 4.543295,
5.834811] (~counts [16.0, 39.0, 94.0, 342.0])
  Stage-2 Outcome [dyadic_aba_reciprocity_rate_48h] (N = 2114, singletons=0, G_domain
= 85, G_week = 261, within-R2 = 0.0300, VIFs = {'headline_negativity_z': 2.1736,
'moral_conflict_blame_z': 2.2687}):
    Joint Neg_z:            coef = +0.003460 (+0.0035, +0.35pp), SE = 0.003388
(0.0034), t = +1.0215 (+1.02), p = 0.309967 (0.3100), 95% CI = [-0.0033, +0.0102] |
beta_std = +0.0329 SD (SE_std = 0.0322, 95% CI_std = [-0.0311, +0.0969]) |
p_Holm(K=5) = 1.0000
    Joint Conflict_z:       coef = +0.013542 (+0.0135, +1.35pp), SE = 0.004804
(0.0048), t = +2.8193 (+2.82), p = 0.006002 (0.0060), 95% CI = [+0.0040, +0.0231] |
beta_std = +0.1286 SD (SE_std = 0.0456, 95% CI_std = [+0.0379, +0.2194]) |
p_Holm(K=5) = 0.0125
    Joint Raw[0,1] Neg:     coef = +0.012582 (+0.0126, +1.26pp), SE = 0.012318
(0.0123), t = +1.0215 (+1.02), p = 0.309967 (0.3100), 95% CI = [-0.0119, +0.0371]
    Joint Raw[0,1] Conf:    coef = +0.048070 (+0.0481, +4.81pp), SE = 0.017051
(0.0171), t = +2.8193 (+2.82), p = 0.006002 (0.0060), 95% CI = [+0.0142, +0.0820]
    Standalone Neg_z:       coef = +0.011834 (+0.0118, +1.18pp), SE = 0.002621
(0.0026), t = +4.5148 (+4.51), p = 0.000020 (0.0000), 95% CI = [+0.0066, +0.0170] |
beta_std = +0.1124 SD (SE_std = 0.0249, 95% CI_std = [+0.0629, +0.1619])
    Standalone Conflict_z:  coef = +0.016002 (+0.0160, +1.60pp), SE = 0.003656
(0.0037), t = +4.3774 (+4.38), p = 0.000034 (0.0000), 95% CI = [+0.0087, +0.0233] |
beta_std = +0.1520 SD (SE_std = 0.0347, 95% CI_std = [+0.0830, +0.2211])
    H3 Wald (Conf - Neg)_z: coef = +0.010082 (+0.0101, +1.01pp), SE = 0.007507
(0.0075), t = +1.3429 (+1.34), p = 0.182908 (0.1829), 95% CI = [-0.0048, +0.0250]
    4-Cell Cell_1 Threat:   coef = +0.013527 (+0.0135, +1.35pp), SE = 0.007549
(0.0075), t = +1.7919 (+1.79), p = 0.076749 (0.0767), 95% CI = [-0.0015, +0.0285]
    4-Cell Cell_2 Blame:    coef = +0.024617 (+0.0246, +2.46pp), SE = 0.009003
(0.0090), t = +2.7342 (+2.73), p = 0.007625 (0.0076), 95% CI = [+0.0067, +0.0425]
    4-Cell Cell_3 Debate:   coef = +0.021200 (+0.0212, +2.12pp), SE = 0.009617
(0.0096), t = +2.2043 (+2.20), p = 0.030237 (0.0302), 95% CI = [+0.0021, +0.0403]
    4-Cell (Cell_2-Cell_1): coef = +0.011090 (+0.0111, +1.11pp), SE = 0.011529
(0.0115), t = +0.9619 (+0.96), p = 0.338872 (0.3389), 95% CI = [-0.0118, +0.0340]
  Stage-2 Outcome [interpersonal_nested_reply_share_48h] (N = 2114, singletons=0,
G_domain = 85, G_week = 261, within-R2 = 0.4244, VIFs = {'headline_negativity_z':
2.1736, 'moral_conflict_blame_z': 2.2687}):
    Joint Neg_z:            coef = -0.001840 (-0.0018, -0.18pp), SE = 0.003072
(0.0031), t = -0.5991 (-0.60), p = 0.550730 (0.5507), 95% CI = [-0.0079, +0.0043] |
beta_std = -0.0166 SD (SE_std = 0.0276, 95% CI_std = [-0.0715, +0.0384]) |
p_Holm(K=5) = 1.0000
    Joint Conflict_z:       coef = +0.012726 (+0.0127, +1.27pp), SE = 0.003383
(0.0034), t = +3.7616 (+3.76), p = 0.000311 (0.0003), 95% CI = [+0.0060, +0.0195] |
beta_std = +0.1145 SD (SE_std = 0.0305, 95% CI_std = [+0.0540, +0.1751]) |
```

```
p_Holm(K=5) = 0.0013
    Joint Raw[0,1] Neg:     coef = -0.006691 (-0.0067, -0.67pp), SE = 0.011168
(0.0112), t = -0.5991 (-0.60), p = 0.550730 (0.5507), 95% CI = [-0.0289, +0.0155]
    Joint Raw[0,1] Conf:    coef = +0.045171 (+0.0452, +4.52pp), SE = 0.012008
(0.0120), t = +3.7616 (+3.76), p = 0.000311 (0.0003), 95% CI = [+0.0213, +0.0691]
    Standalone Neg_z:       coef = +0.006029 (+0.0060, +0.60pp), SE = 0.002282
(0.0023), t = +2.6416 (+2.64), p = 0.009839 (0.0098), 95% CI = [+0.0015, +0.0106] |
beta_std = +0.0543 SD (SE_std = 0.0205, 95% CI_std = [+0.0134, +0.0951])
    Standalone Conflict_z:  coef = +0.011418 (+0.0114, +1.14pp), SE = 0.002491
(0.0025), t = +4.5842 (+4.58), p = 0.000016 (0.0000), 95% CI = [+0.0065, +0.0164] |
beta_std = +0.1028 SD (SE_std = 0.0224, 95% CI_std = [+0.0582, +0.1474])
    H3 Wald (Conf - Neg)_z: coef = +0.014566 (+0.0146, +1.46pp), SE = 0.005911
(0.0059), t = +2.4640 (+2.46), p = 0.015778 (0.0158), 95% CI = [+0.0028, +0.0263]
    4-Cell Cell_1 Threat:   coef = +0.013317 (+0.0133, +1.33pp), SE = 0.004814
(0.0048), t = +2.7661 (+2.77), p = 0.006975 (0.0070), 95% CI = [+0.0037, +0.0229]
    4-Cell Cell_2 Blame:    coef = +0.020397 (+0.0204, +2.04pp), SE = 0.006464
(0.0065), t = +3.1557 (+3.16), p = 0.002223 (0.0022), 95% CI = [+0.0075, +0.0333]
    4-Cell Cell_3 Debate:   coef = +0.017923 (+0.0179, +1.79pp), SE = 0.006628
(0.0066), t = +2.7043 (+2.70), p = 0.008286 (0.0083), 95% CI = [+0.0047, +0.0311]
    4-Cell (Cell_2-Cell_1): coef = +0.007081 (+0.0071, +0.71pp), SE = 0.007543
(0.0075), t = +0.9387 (+0.94), p = 0.350560 (0.3506), 95% CI = [-0.0079, +0.0221]
  Stage-2 Outcome [branch_depth_p90_48h] (N = 2114, singletons=0, G_domain = 85,
G_week = 261, within-R2 = 0.3327, VIFs = {'headline_negativity_z': 2.1736,
'moral_conflict_blame_z': 2.2687}):
    Joint Neg_z:            coef = +0.011655 (+0.0117, +1.17pp), SE = 0.043107
(0.0431), t = +0.2704 (+0.27), p = 0.787541 (0.7875), 95% CI = [-0.0741, +0.0974] |
beta_std = +0.0072 SD (SE_std = 0.0267, 95% CI_std = [-0.0459, +0.0604]) |
p_Holm(K=5) = 1.0000
    Joint Conflict_z:       coef = +0.181183 (+0.1812, +18.12pp), SE = 0.047627
(0.0476), t = +3.8042 (+3.80), p = 0.000269 (0.0003), 95% CI = [+0.0865, +0.2759] |
beta_std = +0.1123 SD (SE_std = 0.0295, 95% CI_std = [+0.0536, +0.1711]) |
p_Holm(K=5) = 0.0013
    Joint Raw[0,1] Neg:     coef = +0.042376 (+0.0424, +4.24pp), SE = 0.156736
(0.1567), t = +0.2704 (+0.27), p = 0.787541 (0.7875), 95% CI = [-0.2693, +0.3541]
    Joint Raw[0,1] Conf:    coef = +0.643127 (+0.6431, +64.31pp), SE = 0.169058
(0.1691), t = +3.8042 (+3.80), p = 0.000269 (0.0003), 95% CI = [+0.3069, +0.9793]
    Standalone Neg_z:       coef = +0.123686 (+0.1237, +12.37pp), SE = 0.035750
(0.0357), t = +3.4598 (+3.46), p = 0.000852 (0.0009), 95% CI = [+0.0526, +0.1948] |
beta_std = +0.0767 SD (SE_std = 0.0222, 95% CI_std = [+0.0326, +0.1208])
    Standalone Conflict_z:  coef = +0.189467 (+0.1895, +18.95pp), SE = 0.038349
(0.0383), t = +4.9406 (+4.94), p = 0.000004 (0.0000), 95% CI = [+0.1132, +0.2657] |
beta_std = +0.1175 SD (SE_std = 0.0238, 95% CI_std = [+0.0702, +0.1648])
    H3 Wald (Conf - Neg)_z: coef = +0.169528 (+0.1695, +16.95pp), SE = 0.081047
(0.0810), t = +2.0917 (+2.09), p = 0.039484 (0.0395), 95% CI = [+0.0084, +0.3307]
    4-Cell Cell_1 Threat:   coef = +0.159460 (+0.1595, +15.95pp), SE = 0.093994
(0.0940), t = +1.6965 (+1.70), p = 0.093496 (0.0935), 95% CI = [-0.0275, +0.3464]
```

```
    4-Cell Cell_2 Blame:    coef = +0.348681 (+0.3487, +34.87pp), SE = 0.103652
(0.1037), t = +3.3640 (+3.36), p = 0.001159 (0.0012), 95% CI = [+0.1426, +0.5548]
    4-Cell Cell_3 Debate:   coef = +0.269436 (+0.2694, +26.94pp), SE = 0.090164
(0.0902), t = +2.9883 (+2.99), p = 0.003678 (0.0037), 95% CI = [+0.0901, +0.4487]
    4-Cell (Cell_2-Cell_1): coef = +0.189221 (+0.1892, +18.92pp), SE = 0.121387
(0.1214), t = +1.5588 (+1.56), p = 0.122796 (0.1228), 95% CI = [-0.0522, +0.4306]
  Stage-2 Outcome [external_evidence_link_rate_48h] (N = 2114, singletons=0, G_domain
= 85, G_week = 261, within-R2 = 0.0317, VIFs = {'headline_negativity_z': 2.1736,
'moral_conflict_blame_z': 2.2687}):
    Joint Neg_z:            coef = +0.001732 (+0.0017, +0.17pp), SE = 0.002342
(0.0023), t = +0.7394 (+0.74), p = 0.461727 (0.4617), 95% CI = [-0.0029, +0.0064] |
beta_std = +0.0266 SD (SE_std = 0.0360, 95% CI_std = [-0.0450, +0.0983]) |
p_Holm(K=5) = 1.0000
    Joint Conflict_z:       coef = -0.007289 (-0.0073, -0.73pp), SE = 0.002475
(0.0025), t = -2.9450 (-2.95), p = 0.004177 (0.0042), 95% CI = [-0.0122, -0.0024] |
beta_std = -0.1121 SD (SE_std = 0.0381, 95% CI_std = [-0.1878, -0.0364]) |
p_Holm(K=5) = 0.0125
    Joint Raw[0,1] Neg:     coef = +0.006297 (+0.0063, +0.63pp), SE = 0.008516
(0.0085), t = +0.7394 (+0.74), p = 0.461727 (0.4617), 95% CI = [-0.0106, +0.0232]
    Joint Raw[0,1] Conf:    coef = -0.025872 (-0.0259, -2.59pp), SE = 0.008785
(0.0088), t = -2.9450 (-2.95), p = 0.004177 (0.0042), 95% CI = [-0.0433, -0.0084]
    Standalone Neg_z:       coef = -0.002775 (-0.0028, -0.28pp), SE = 0.001553
(0.0016), t = -1.7871 (-1.79), p = 0.077526 (0.0775), 95% CI = [-0.0059, +0.0003] |
beta_std = -0.0427 SD (SE_std = 0.0239, 95% CI_std = [-0.0902, +0.0048])
    Standalone Conflict_z:  coef = -0.006058 (-0.0061, -0.61pp), SE = 0.001592
(0.0016), t = -3.8043 (-3.80), p = 0.000269 (0.0003), 95% CI = [-0.0092, -0.0029] |
beta_std = -0.0932 SD (SE_std = 0.0245, 95% CI_std = [-0.1419, -0.0445])
    H3 Wald (Conf - Neg)_z: coef = -0.009021 (-0.0090, -0.90pp), SE = 0.004536
(0.0045), t = -1.9886 (-1.99), p = 0.050002 (0.0500), 95% CI = [-0.0180, +0.0000]
    4-Cell Cell_1 Threat:   coef = +0.000398 (+0.0004, +0.04pp), SE = 0.003969
(0.0040), t = +0.1003 (+0.10), p = 0.920353 (0.9204), 95% CI = [-0.0075, +0.0083]
    4-Cell Cell_2 Blame:    coef = -0.003123 (-0.0031, -0.31pp), SE = 0.003427
(0.0034), t = -0.9114 (-0.91), p = 0.364673 (0.3647), 95% CI = [-0.0099, +0.0037]
    4-Cell Cell_3 Debate:   coef = -0.005375 (-0.0054, -0.54pp), SE = 0.004554
(0.0046), t = -1.1802 (-1.18), p = 0.241234 (0.2412), 95% CI = [-0.0144, +0.0037]
    4-Cell (Cell_2-Cell_1): coef = -0.003521 (-0.0035, -0.35pp), SE = 0.005267
(0.0053), t = -0.6685 (-0.67), p = 0.505663 (0.5057), 95% CI = [-0.0140, +0.0070]
  Stage-2 Outcome [speaker_shannon_entropy_48h] (N = 2114, singletons=0, G_domain =
85, G_week = 261, within-R2 = 0.0462, VIFs = {'headline_negativity_z': 2.1736,
'moral_conflict_blame_z': 2.2687}):
    Joint Neg_z:            coef = -0.000359 (-0.0004, -0.04pp), SE = 0.000913
(0.0009), t = -0.3932 (-0.39), p = 0.695187 (0.6952), 95% CI = [-0.0022, +0.0015] |
beta_std = -0.0133 SD (SE_std = 0.0339, 95% CI_std = [-0.0807, +0.0541]) |
p_Holm(K=5) = 1.0000
    Joint Conflict_z:       coef = -0.002865 (-0.0029, -0.29pp), SE = 0.001185
(0.0012), t = -2.4188 (-2.42), p = 0.017730 (0.0177), 95% CI = [-0.0052, -0.0005] |
```

```
beta_std = -0.1063 SD (SE_std = 0.0440, 95% CI_std = [-0.1937, -0.0189]) |
p_Holm(K=5) = 0.0177
    Joint Raw[0,1] Neg:     coef = -0.001305 (-0.0013, -0.13pp), SE = 0.003320
(0.0033), t = -0.3932 (-0.39), p = 0.695187 (0.6952), 95% CI = [-0.0079, +0.0053]
    Joint Raw[0,1] Conf:    coef = -0.010171 (-0.0102, -1.02pp), SE = 0.004205
(0.0042), t = -2.4188 (-2.42), p = 0.017730 (0.0177), 95% CI = [-0.0185, -0.0018]
    Standalone Neg_z:       coef = -0.002131 (-0.0021, -0.21pp), SE = 0.000611
(0.0006), t = -3.4859 (-3.49), p = 0.000782 (0.0008), 95% CI = [-0.0033, -0.0009] |
beta_std = -0.0791 SD (SE_std = 0.0227, 95% CI_std = [-0.1242, -0.0340])
    Standalone Conflict_z:  coef = -0.003121 (-0.0031, -0.31pp), SE = 0.000825
(0.0008), t = -3.7829 (-3.78), p = 0.000290 (0.0003), 95% CI = [-0.0048, -0.0015] |
beta_std = -0.1158 SD (SE_std = 0.0306, 95% CI_std = [-0.1767, -0.0549])
    H3 Wald (Conf - Neg)_z: coef = -0.002506 (-0.0025, -0.25pp), SE = 0.001961
(0.0020), t = -1.2780 (-1.28), p = 0.204772 (0.2048), 95% CI = [-0.0064, +0.0014]
    4-Cell Cell_1 Threat:   coef = -0.002924 (-0.0029, -0.29pp), SE = 0.002164
(0.0022), t = -1.3513 (-1.35), p = 0.180213 (0.1802), 95% CI = [-0.0072, +0.0014]
    4-Cell Cell_2 Blame:    coef = -0.003991 (-0.0040, -0.40pp), SE = 0.002165
(0.0022), t = -1.8433 (-1.84), p = 0.068813 (0.0688), 95% CI = [-0.0083, +0.0003]
    4-Cell Cell_3 Debate:   coef = -0.002774 (-0.0028, -0.28pp), SE = 0.002865
(0.0029), t = -0.9683 (-0.97), p = 0.335693 (0.3357), 95% CI = [-0.0085, +0.0029]
    4-Cell (Cell_2-Cell_1): coef = -0.001068 (-0.0011, -0.11pp), SE = 0.002901
(0.0029), t = -0.3680 (-0.37), p = 0.713804 (0.7138), 95% CI = [-0.0068, +0.0047]
  Stage-2 Outcome [dyadic_aba_total_share_48h] (N = 2114, singletons=0, G_domain =
85, G_week = 261, within-R2 = 0.0695, VIFs = {'headline_negativity_z': 2.1736,
'moral_conflict_blame_z': 2.2687}):
    Joint Neg_z:            coef = +0.002198 (+0.0022, +0.22pp), SE = 0.002957
(0.0030), t = +0.7431 (+0.74), p = 0.459503 (0.4595), 95% CI = [-0.0037, +0.0081] |
beta_std = +0.0241 SD (SE_std = 0.0325, 95% CI_std = [-0.0405, +0.0888])
    Joint Conflict_z:       coef = +0.013686 (+0.0137, +1.37pp), SE = 0.004168
(0.0042), t = +3.2835 (+3.28), p = 0.001496 (0.0015), 95% CI = [+0.0054, +0.0220] |
beta_std = +0.1504 SD (SE_std = 0.0458, 95% CI_std = [+0.0593, +0.2414])
    Joint Raw[0,1] Neg:     coef = +0.007991 (+0.0080, +0.80pp), SE = 0.010753
(0.0108), t = +0.7431 (+0.74), p = 0.459503 (0.4595), 95% CI = [-0.0134, +0.0294]
    Joint Raw[0,1] Conf:    coef = +0.048579 (+0.0486, +4.86pp), SE = 0.014795
(0.0148), t = +3.2835 (+3.28), p = 0.001496 (0.0015), 95% CI = [+0.0192, +0.0780]
    Standalone Neg_z:       coef = +0.010660 (+0.0107, +1.07pp), SE = 0.002195
(0.0022), t = +4.8574 (+4.86), p = 0.000005 (0.0000), 95% CI = [+0.0063, +0.0150] |
beta_std = +0.1171 SD (SE_std = 0.0241, 95% CI_std = [+0.0692, +0.1651])
    Standalone Conflict_z:  coef = +0.015248 (+0.0152, +1.52pp), SE = 0.003093
(0.0031), t = +4.9304 (+4.93), p = 0.000004 (0.0000), 95% CI = [+0.0091, +0.0214] |
beta_std = +0.1675 SD (SE_std = 0.0340, 95% CI_std = [+0.1000, +0.2351])
    H3 Wald (Conf - Neg)_z: coef = +0.011488 (+0.0115, +1.15pp), SE = 0.006581
(0.0066), t = +1.7455 (+1.75), p = 0.084550 (0.0845), 95% CI = [-0.0016, +0.0246]
    4-Cell Cell_1 Threat:   coef = +0.013532 (+0.0135, +1.35pp), SE = 0.006482
(0.0065), t = +2.0875 (+2.09), p = 0.039871 (0.0399), 95% CI = [+0.0006, +0.0264]
    4-Cell Cell_2 Blame:    coef = +0.023849 (+0.0238, +2.38pp), SE = 0.007820
```

```
(0.0078), t = +3.0498 (+3.05), p = 0.003063 (0.0031), 95% CI = [+0.0083, +0.0394]
    4-Cell Cell_3 Debate:   coef = +0.020654 (+0.0207, +2.07pp), SE = 0.008504
(0.0085), t = +2.4286 (+2.43), p = 0.017291 (0.0173), 95% CI = [+0.0037, +0.0376]
    4-Cell (Cell_2-Cell_1): coef = +0.010317 (+0.0103, +1.03pp), SE = 0.010123
(0.0101), t = +1.0191 (+1.02), p = 0.311059 (0.3111), 95% CI = [-0.0098, +0.0304]
  Stage-2 Outcome [branch_depth_mean_48h] (N = 2114, singletons=0, G_domain = 85,
G_week = 261, within-R2 = 0.3954, VIFs = {'headline_negativity_z': 2.1736,
'moral_conflict_blame_z': 2.2687}):
    Joint Neg_z:            coef = -0.001194 (-0.0012, -0.12pp), SE = 0.020969
(0.0210), t = -0.0570 (-0.06), p = 0.954714 (0.9547), 95% CI = [-0.0429, +0.0405] |
beta_std = -0.0015 SD (SE_std = 0.0262, 95% CI_std = [-0.0535, +0.0506])
    Joint Conflict_z:       coef = +0.100994 (+0.1010, +10.10pp), SE = 0.023630
(0.0236), t = +4.2739 (+4.27), p = 0.000050 (0.0001), 95% CI = [+0.0540, +0.1480] |
beta_std = +0.1261 SD (SE_std = 0.0295, 95% CI_std = [+0.0674, +0.1847])
    Joint Raw[0,1] Neg:     coef = -0.004343 (-0.0043, -0.43pp), SE = 0.076244
(0.0762), t = -0.0570 (-0.06), p = 0.954714 (0.9547), 95% CI = [-0.1560, +0.1473]
    Joint Raw[0,1] Conf:    coef = +0.358489 (+0.3585, +35.85pp), SE = 0.083879
(0.0839), t = +4.2739 (+4.27), p = 0.000050 (0.0001), 95% CI = [+0.1917, +0.5253]
    Standalone Neg_z:       coef = +0.061254 (+0.0613, +6.13pp), SE = 0.017589
(0.0176), t = +3.4825 (+3.48), p = 0.000791 (0.0008), 95% CI = [+0.0263, +0.0962] |
beta_std = +0.0765 SD (SE_std = 0.0220, 95% CI_std = [+0.0328, +0.1201])
    Standalone Conflict_z:  coef = +0.100146 (+0.1001, +10.01pp), SE = 0.019000
(0.0190), t = +5.2708 (+5.27), p = 0.000001 (0.0000), 95% CI = [+0.0624, +0.1379] |
beta_std = +0.1250 SD (SE_std = 0.0237, 95% CI_std = [+0.0778, +0.1722])
    H3 Wald (Conf - Neg)_z: coef = +0.102189 (+0.1022, +10.22pp), SE = 0.039850
(0.0398), t = +2.5643 (+2.56), p = 0.012116 (0.0121), 95% CI = [+0.0229, +0.1814]
    4-Cell Cell_1 Threat:   coef = +0.097069 (+0.0971, +9.71pp), SE = 0.045689
(0.0457), t = +2.1245 (+2.12), p = 0.036564 (0.0366), 95% CI = [+0.0062, +0.1879]
    4-Cell Cell_2 Blame:    coef = +0.179691 (+0.1797, +17.97pp), SE = 0.052097
(0.0521), t = +3.4492 (+3.45), p = 0.000882 (0.0009), 95% CI = [+0.0761, +0.2833]
    4-Cell Cell_3 Debate:   coef = +0.149427 (+0.1494, +14.94pp), SE = 0.048074
(0.0481), t = +3.1083 (+3.11), p = 0.002568 (0.0026), 95% CI = [+0.0538, +0.2450]
    4-Cell (Cell_2-Cell_1): coef = +0.082622 (+0.0826, +8.26pp), SE = 0.061907
(0.0619), t = +1.3346 (+1.33), p = 0.185609 (0.1856), 95% CI = [-0.0405, +0.2057]
  Stage-2 Outcome [mean_log_comment_words_48h] (N = 2114, singletons=0, G_domain =
85, G_week = 261, within-R2 = 0.0477, VIFs = {'headline_negativity_z': 2.1736,
'moral_conflict_blame_z': 2.2687}):
    Joint Neg_z:            coef = -0.010999 (-0.0110, -1.10pp), SE = 0.011609
(0.0116), t = -0.9474 (-0.95), p = 0.346147 (0.3461), 95% CI = [-0.0341, +0.0121] |
beta_std = -0.0395 SD (SE_std = 0.0417, 95% CI_std = [-0.1225, +0.0434])
    Joint Conflict_z:       coef = +0.016614 (+0.0166, +1.66pp), SE = 0.014429
(0.0144), t = +1.1515 (+1.15), p = 0.252813 (0.2528), 95% CI = [-0.0121, +0.0453] |
beta_std = +0.0597 SD (SE_std = 0.0519, 95% CI_std = [-0.0434, +0.1628])
    Joint Raw[0,1] Neg:     coef = -0.039991 (-0.0400, -4.00pp), SE = 0.042211
(0.0422), t = -0.9474 (-0.95), p = 0.346147 (0.3461), 95% CI = [-0.1239, +0.0439]
    Joint Raw[0,1] Conf:    coef = +0.058973 (+0.0590, +5.90pp), SE = 0.051216
```

```
(0.0512), t = +1.1515 (+1.15), p = 0.252813 (0.2528), 95% CI = [-0.0429, +0.1608]
    Standalone Neg_z:       coef = -0.000726 (-0.0007, -0.07pp), SE = 0.006723
(0.0067), t = -0.1079 (-0.11), p = 0.914295 (0.9143), 95% CI = [-0.0141, +0.0126] |
beta_std = -0.0026 SD (SE_std = 0.0242, 95% CI_std = [-0.0507, +0.0454])
    Standalone Conflict_z:  coef = +0.008797 (+0.0088, +0.88pp), SE = 0.009037
(0.0090), t = +0.9735 (+0.97), p = 0.333121 (0.3331), 95% CI = [-0.0092, +0.0268] |
beta_std = +0.0316 SD (SE_std = 0.0325, 95% CI_std = [-0.0330, +0.0962])
    H3 Wald (Conf - Neg)_z: coef = +0.027613 (+0.0276, +2.76pp), SE = 0.024822
(0.0248), t = +1.1124 (+1.11), p = 0.269122 (0.2691), 95% CI = [-0.0217, +0.0770]
    4-Cell Cell_1 Threat:   coef = -0.033103 (-0.0331, -3.31pp), SE = 0.020871
(0.0209), t = -1.5861 (-1.59), p = 0.116478 (0.1165), 95% CI = [-0.0746, +0.0084]
    4-Cell Cell_2 Blame:    coef = +0.002804 (+0.0028, +0.28pp), SE = 0.015739
(0.0157), t = +0.1781 (+0.18), p = 0.859049 (0.8590), 95% CI = [-0.0285, +0.0341]
    4-Cell Cell_3 Debate:   coef = -0.003866 (-0.0039, -0.39pp), SE = 0.022635
(0.0226), t = -0.1708 (-0.17), p = 0.864798 (0.8648), 95% CI = [-0.0489, +0.0411]
    4-Cell (Cell_2-Cell_1): coef = +0.035907 (+0.0359, +3.59pp), SE = 0.028044
(0.0280), t = +1.2804 (+1.28), p = 0.203931 (0.2039), 95% CI = [-0.0199, +0.0917]

  --- MODEL 4: Paywall Information-Starvation Boundary Condition (RQ4 / H4) ---
  First-Stage Behavioral Check (Paywall_Mirror_Link_Rate_48h on Paywalled_Domain, N =
2114): coef = +0.009202 (+0.0092, +0.92pp), SE = 0.001607 (0.0016), t = +5.7260
(+5.73), p = 0.000000 (0.0000), 95% CI = [+0.0060, +0.0124] | beta_std = +0.3106 SD
(SE_std = 0.0543, 95% CI_std = [+0.2028, +0.4185])
  H4 Outcome [dyadic_aba_reciprocity_rate_48h] (N = 2114):
    gamma_pw_neg (Neg_z x Paywalled):   coef = +0.001426 (+0.0014, +0.14pp), SE =
0.006074 (0.0061), t = +0.2348 (+0.23), p = 0.814906 (0.8149), 95% CI = [-0.0107,
+0.0135] | p_Holm(K=3) = 1.0000
    gamma_pw_conf (Conf_z x Paywalled): coef = +0.005219 (+0.0052, +0.52pp), SE =
0.008784 (0.0088), t = +0.5941 (+0.59), p = 0.554014 (0.5540), 95% CI = [-0.0122,
+0.0227] | p_Holm(K=3) = 1.0000
    Open-Access Neg_z slope:            coef = +0.002940 (+0.0029, +0.29pp), SE =
0.003772 (0.0038), t = +0.7796 (+0.78), p = 0.437818 (0.4378), 95% CI = [-0.0046,
+0.0104]
    Open-Access Conf_z slope:           coef = +0.011697 (+0.0117, +1.17pp), SE =
0.005134 (0.0051), t = +2.2783 (+2.28), p = 0.025249 (0.0252), 95% CI = [+0.0015,
+0.0219]
    Paywalled Total Neg_z slope:        coef = +0.004367 (+0.0044, +0.44pp), SE =
0.005350 (0.0054), t = +0.8162 (+0.82), p = 0.416721 (0.4167), 95% CI = [-0.0063,
+0.0150]
    Paywalled Total Conf_z slope:       coef = +0.016915 (+0.0169, +1.69pp), SE =
0.008035 (0.0080), t = +2.1053 (+2.11), p = 0.038255 (0.0383), 95% CI = [+0.0009,
+0.0329]
  H4 Outcome [interpersonal_nested_reply_share_48h] (N = 2114):
    gamma_pw_neg (Neg_z x Paywalled):   coef = +0.010708 (+0.0107, +1.07pp), SE =
0.004667 (0.0047), t = +2.2946 (+2.29), p = 0.024249 (0.0242), 95% CI = [+0.0014,
+0.0200] | p_Holm(K=3) = 0.0727
```

```
    gamma_pw_conf (Conf_z x Paywalled): coef = -0.004648 (-0.0046, -0.46pp), SE =
0.005013 (0.0050), t = -0.9271 (-0.93), p = 0.356508 (0.3565), 95% CI = [-0.0146,
+0.0053] | p_Holm(K=3) = 1.0000
    Open-Access Neg_z slope:            coef = -0.006478 (-0.0065, -0.65pp), SE =
0.003853 (0.0039), t = -1.6811 (-1.68), p = 0.096453 (0.0965), 95% CI = [-0.0141,
+0.0012]
    Open-Access Conf_z slope:           coef = +0.015084 (+0.0151, +1.51pp), SE =
0.004044 (0.0040), t = +3.7299 (+3.73), p = 0.000347 (0.0003), 95% CI = [+0.0070,
+0.0231]
    Paywalled Total Neg_z slope:        coef = +0.004231 (+0.0042, +0.42pp), SE =
0.003539 (0.0035), t = +1.1953 (+1.20), p = 0.235321 (0.2353), 95% CI = [-0.0028,
+0.0113]
    Paywalled Total Conf_z slope:       coef = +0.010437 (+0.0104, +1.04pp), SE =
0.004211 (0.0042), t = +2.4787 (+2.48), p = 0.015185 (0.0152), 95% CI = [+0.0021,
+0.0188]
  H4 Outcome [external_evidence_link_rate_48h] (N = 2114):
    gamma_pw_neg (Neg_z x Paywalled):   coef = +0.000076 (+0.0001, +0.01pp), SE =
0.004955 (0.0050), t = +0.0153 (+0.02), p = 0.987826 (0.9878), 95% CI = [-0.0098,
+0.0099] | p_Holm(K=3) = 1.0000
    gamma_pw_conf (Conf_z x Paywalled): coef = -0.001214 (-0.0012, -0.12pp), SE =
0.004861 (0.0049), t = -0.2497 (-0.25), p = 0.803420 (0.8034), 95% CI = [-0.0109,
+0.0085] | p_Holm(K=3) = 1.0000
    Open-Access Neg_z slope:            coef = +0.001679 (+0.0017, +0.17pp), SE =
0.003357 (0.0034), t = +0.5002 (+0.50), p = 0.618275 (0.6183), 95% CI = [-0.0050,
+0.0084]
    Open-Access Conf_z slope:           coef = -0.006835 (-0.0068, -0.68pp), SE =
0.003597 (0.0036), t = -1.9001 (-1.90), p = 0.060848 (0.0608), 95% CI = [-0.0140,
+0.0003]
    Paywalled Total Neg_z slope:        coef = +0.001755 (+0.0018, +0.18pp), SE =
0.003474 (0.0035), t = +0.5052 (+0.51), p = 0.614721 (0.6147), 95% CI = [-0.0052,
+0.0087]
    Paywalled Total Conf_z slope:       coef = -0.008049 (-0.0080, -0.80pp), SE =
0.003206 (0.0032), t = -2.5107 (-2.51), p = 0.013967 (0.0140), 95% CI = [-0.0144,
-0.0017]
  H4 Outcome [branch_depth_p90_48h] (N = 2114):
    gamma_pw_neg (Neg_z x Paywalled):   coef = +0.076895 (+0.0769, +7.69pp), SE =
0.081398 (0.0814), t = +0.9447 (+0.94), p = 0.347536 (0.3475), 95% CI = [-0.0850,
+0.2388]
    gamma_pw_conf (Conf_z x Paywalled): coef = -0.007459 (-0.0075, -0.75pp), SE =
0.093804 (0.0938), t = -0.0795 (-0.08), p = 0.936813 (0.9368), 95% CI = [-0.1940,
+0.1791]
    Open-Access Neg_z slope:            coef = -0.021213 (-0.0212, -2.12pp), SE =
0.052820 (0.0528), t = -0.4016 (-0.40), p = 0.688995 (0.6890), 95% CI = [-0.1263,
+0.0838]
    Open-Access Conf_z slope:           coef = +0.188533 (+0.1885, +18.85pp), SE =
0.055600 (0.0556), t = +3.3909 (+3.39), p = 0.001064 (0.0011), 95% CI = [+0.0780,
```

```
+0.2991]
    Paywalled Total Neg_z slope:        coef = +0.055682 (+0.0557, +5.57pp), SE =
0.066449 (0.0664), t = +0.8380 (+0.84), p = 0.404429 (0.4044), 95% CI = [-0.0765,
+0.1878]
    Paywalled Total Conf_z slope:       coef = +0.181075 (+0.1811, +18.11pp), SE =
0.079346 (0.0793), t = +2.2821 (+2.28), p = 0.025012 (0.0250), 95% CI = [+0.0233,
+0.3389]
  H4 Outcome [speaker_shannon_entropy_48h] (N = 2114):
    gamma_pw_neg (Neg_z x Paywalled):   coef = -0.000270 (-0.0003, -0.03pp), SE =
0.001565 (0.0016), t = -0.1724 (-0.17), p = 0.863534 (0.8635), 95% CI = [-0.0034,
+0.0028]
    gamma_pw_conf (Conf_z x Paywalled): coef = -0.002074 (-0.0021, -0.21pp), SE =
0.002007 (0.0020), t = -1.0336 (-1.03), p = 0.304297 (0.3043), 95% CI = [-0.0061,
+0.0019]
    Open-Access Neg_z slope:            coef = -0.000279 (-0.0003, -0.03pp), SE =
0.001095 (0.0011), t = -0.2546 (-0.25), p = 0.799645 (0.7996), 95% CI = [-0.0025,
+0.0019]
    Open-Access Conf_z slope:           coef = -0.002114 (-0.0021, -0.21pp), SE =
0.001396 (0.0014), t = -1.5147 (-1.51), p = 0.133603 (0.1336), 95% CI = [-0.0049,
+0.0007]
    Paywalled Total Neg_z slope:        coef = -0.000549 (-0.0005, -0.05pp), SE =
0.001339 (0.0013), t = -0.4098 (-0.41), p = 0.682982 (0.6830), 95% CI = [-0.0032,
+0.0021]
    Paywalled Total Conf_z slope:       coef = -0.004188 (-0.0042, -0.42pp), SE =
0.001730 (0.0017), t = -2.4206 (-2.42), p = 0.017650 (0.0176), 95% CI = [-0.0076,
-0.0007]

[SECTION 5] MANDATORY 10-GATE FALSIFICATION, RELIABILITY & OUT-OF-SAMPLE SUITE
  [Gate 1 — Independent Ballot ICC Re-Computation & Per-Judge Standalone Regressions]
    Judge j1_llama (meta-llama/llama-3.3-70b-instruct): S1_Brk(Neg t=+1.74,
b=+0.0163, p=0.0857; Conf t=-0.13, b=-0.0014, p=0.8977) | S2_ABA(Neg t=+1.16,
b=+0.0039, p=0.2507; Conf t=+2.92, b=+0.0117, p=0.0045)
    Judge j2_qwen (qwen/qwen-2.5-72b-instruct): S1_Brk(Neg t=+0.67, b=+0.0055,
p=0.5016; Conf t=-0.07, b=-0.0007, p=0.9428) | S2_ABA(Neg t=+2.50, b=+0.0072,
p=0.0144; Conf t=+1.32, b=+0.0051, p=0.1889)
    Judge j3_deepseek (deepseek/deepseek-chat-v3-0324): S1_Brk(Neg t=+1.00,
b=+0.0106, p=0.3208; Conf t=-0.21, b=-0.0024, p=0.8318) | S2_ABA(Neg t=+0.82,
b=+0.0031, p=0.4163; Conf t=+2.66, b=+0.0144, p=0.0093)
    Judge j4_gemma (google/gemma-3-27b-it): S1_Brk(Neg t=+1.99, b=+0.0220, p=0.0497;
Conf t=-1.06, b=-0.0125, p=0.2923) | S2_ABA(Neg t=+0.38, b=+0.0011, p=0.7018; Conf
t=+3.53, b=+0.0164, p=0.0007)
  [Gate 2 — Unaltered Verbatim Newsroom Headlines (Has_Lexical_Slug == 1 & Overlap >=
0.60)]
    Stage 1: 2167/3172 (68.32%) unedited | Stage 2: 1182/1797 (65.78%) unedited
    S1 Breakthrough Unedited: Neg_z -> coef = +0.009591 (+0.0096, +0.96pp), SE =
0.014998 (0.0150), t = +0.6395 (+0.64), p = 0.524316 (0.5243), 95% CI = [-0.0203,
```

```
+0.0394] | beta_std = +0.0308 SD (SE_std = 0.0481, 95% CI_std = [-0.0650, +0.1265]) |
Conf_z -> coef = +0.003034 (+0.0030, +0.30pp), SE = 0.014921 (0.0149), t = +0.2033
(+0.20), p = 0.839386 (0.8394), 95% CI = [-0.0267, +0.0327] | beta_std = +0.0100 SD
(SE_std = 0.0491, 95% CI_std = [-0.0877, +0.1076])
    S2 ABA 48h Unedited:      Neg_z -> coef = +0.000136 (+0.0001, +0.01pp), SE =
0.004278 (0.0043), t = +0.0317 (+0.03), p = 0.974776 (0.9748), 95% CI = [-0.0084,
+0.0086] | beta_std = +0.0013 SD (SE_std = 0.0394, 95% CI_std = [-0.0772, +0.0797]) |
Conf_z -> coef = +0.014019 (+0.0140, +1.40pp), SE = 0.005847 (0.0058), t = +2.3978
(+2.40), p = 0.018853 (0.0189), 95% CI = [+0.0024, +0.0257] | beta_std = +0.1328 SD
(SE_std = 0.0554, 95% CI_std = [+0.0226, +0.2431])
  [Gate 3 — Stage-1 to Stage-2 Heckman IMR & IPW Selection Correction (Probit pseudo-
R2 = 0.0436)]
    Heckman IMR ABA 48h: Neg_z -> coef = +0.004736 (+0.0047, +0.47pp), SE = 0.003506
(0.0035), t = +1.3507 (+1.35), p = 0.180408 (0.1804), 95% CI = [-0.0022, +0.0117] |
Conf_z -> coef = +0.013682 (+0.0137, +1.37pp), SE = 0.004861 (0.0049), t = +2.8144
(+2.81), p = 0.006086 (0.0061), 95% CI = [+0.0040, +0.0233] | IMR lambda -> coef =
+0.022847 (+0.0228, +2.28pp), SE = 0.033530 (0.0335), t = +0.6814 (+0.68), p =
0.497497 (0.4975), 95% CI = [-0.0438, +0.0895]
    IPW Breakthrough ABA: Neg_z -> coef = +0.000927 (+0.0009, +0.09pp), SE = 0.003845
(0.0038), t = +0.2411 (+0.24), p = 0.810054 (0.8101), 95% CI = [-0.0067, +0.0086] |
Conf_z -> coef = +0.017906 (+0.0179, +1.79pp), SE = 0.005155 (0.0052), t = +3.4736
(+3.47), p = 0.000814 (0.0008), 95% CI = [+0.0077, +0.0282]
    Combined Universe+Breakthrough IPW ABA: Neg_z -> coef = +0.000808 (+0.0008,
+0.08pp), SE = 0.004214 (0.0042), t = +0.1918 (+0.19), p = 0.848399 (0.8484), 95% CI
= [-0.0076, +0.0092] | Conf_z -> coef = +0.020185 (+0.0202, +2.02pp), SE = 0.006062
(0.0061), t = +3.3299 (+3.33), p = 0.001292 (0.0013), 95% CI = [+0.0081, +0.0322]
  [Gate 4 — Stage-2 Volume-Control Functional-Form Decomposition]
    dyadic_aba_reciprocity_rate_48h: 4-Knot_RCS(Neg b=+0.0035, t=+1.02; Conf
b=+0.0135, t=+2.82) | 20_Vigintile_FE(Neg b=+0.0037, t=+1.10; Conf b=+0.0135,
t=+2.75) | Unadjusted(Neg b=+0.0038, t=+1.08; Conf b=+0.0130, t=+2.70)
    interpersonal_nested_reply_share_48h: 4-Knot_RCS(Neg b=-0.0018, t=-0.60; Conf
b=+0.0127, t=+3.76) | 20_Vigintile_FE(Neg b=-0.0018, t=-0.58; Conf b=+0.0132,
t=+3.76) | Unadjusted(Neg b=+0.0014, t=+0.32; Conf b=+0.0091, t=+1.88)
    branch_depth_p90_48h: 4-Knot_RCS(Neg b=+0.0117, t=+0.27; Conf b=+0.1812, t=+3.80)
| 20_Vigintile_FE(Neg b=+0.0099, t=+0.23; Conf b=+0.1851, t=+3.77) | Unadjusted(Neg
b=+0.0468, t=+0.72; Conf b=+0.1483, t=+2.23)
    external_evidence_link_rate_48h: 4-Knot_RCS(Neg b=+0.0017, t=+0.74; Conf
b=-0.0073, t=-2.95) | 20_Vigintile_FE(Neg b=+0.0018, t=+0.80; Conf b=-0.0074,
t=-2.96) | Unadjusted(Neg b=+0.0014, t=+0.60; Conf b=-0.0070, t=-2.81)
  [Gate 5 — Post-Submission Horizon Invariance: 24h (N=1996) vs 48h vs Full 14d
Lifecycle]
    24h ABA: Neg_z -> coef = +0.005257 (+0.0053, +0.53pp), SE = 0.003503 (0.0035), t
= +1.5009 (+1.50), p = 0.137138 (0.1371), 95% CI = [-0.0017, +0.0122] | Conf_z ->
coef = +0.012924 (+0.0129, +1.29pp), SE = 0.004586 (0.0046), t = +2.8185 (+2.82), p =
0.006016 (0.0060), 95% CI = [+0.0038, +0.0220]
    48h ABA: Neg_z -> coef = +0.003460 (+0.0035, +0.35pp), SE = 0.003388 (0.0034), t
```

```
= +1.0215 (+1.02), p = 0.309967 (0.3100), 95% CI = [-0.0033, +0.0102] | Conf_z ->
coef = +0.013542 (+0.0135, +1.35pp), SE = 0.004804 (0.0048), t = +2.8193 (+2.82), p =
0.006002 (0.0060), 95% CI = [+0.0040, +0.0231]
    14d ABA: Neg_z -> coef = +0.001095 (+0.0011, +0.11pp), SE = 0.002952 (0.0030), t
= +0.3709 (+0.37), p = 0.711668 (0.7117), 95% CI = [-0.0048, +0.0070] | Conf_z ->
coef = +0.013531 (+0.0135, +1.35pp), SE = 0.004544 (0.0045), t = +2.9778 (+2.98), p =
0.003793 (0.0038), 95% CI = [+0.0045, +0.0226]
    14d Nested: Neg_z -> coef = -0.001431 (-0.0014, -0.14pp), SE = 0.002983 (0.0030),
t = -0.4797 (-0.48), p = 0.632710 (0.6327), 95% CI = [-0.0074, +0.0045] | Conf_z ->
coef = +0.012185 (+0.0122, +1.22pp), SE = 0.003409 (0.0034), t = +3.5742 (+3.57), p =
0.000585 (0.0006), 95% CI = [+0.0054, +0.0190]
  [Gate 6 — A->A Self-Reply & Paywall-Mirror Positive-Control Audit]
    Open-Access Outlets (M=58): Mean Mirror Rate = 0.202% (0.002023) vs Paywalled
Outlets (M=27): 1.089% (0.010892, ratio = 5.38x); Mean Self-Reply Share = 0.28%
    Paywalled_Domain -> Unstripped All Links:  coef = +0.005231 (+0.0052, +0.52pp),
SE = 0.004641 (0.0046), t = +1.1270 (+1.13), p = 0.262934 (0.2629), 95% CI =
[-0.0040, +0.0145]
    Paywalled_Domain -> Clean Evidence Links:  coef = -0.004619 (-0.0046, -0.46pp),
SE = 0.004088 (0.0041), t = -1.1300 (-1.13), p = 0.261710 (0.2617), 95% CI =
[-0.0127, +0.0035]
  [Gate 7 — Universe-Weighted Stage-1 IPW, Logit & Poisson-PPML]
    Stage-1 IPW Breakthrough: Neg_z -> coef = +0.002254 (+0.0023, +0.23pp), SE =
0.010720 (0.0107), t = +0.2102 (+0.21), p = 0.833990 (0.8340), 95% CI = [-0.0191,
+0.0236] | Conf_z -> coef = +0.008215 (+0.0082, +0.82pp), SE = 0.012436 (0.0124), t =
+0.6606 (+0.66), p = 0.510654 (0.5107), 95% CI = [-0.0165, +0.0329]
    Stage-1 Logit Breakthrough: Neg_z log_odds=+0.1446 (z=+1.52, p=0.1294) | Conf_z
log_odds=+0.0102 (z=+0.10, p=0.9237)
    Stage-1 Poisson PPML Comments: Neg_z=-0.0920 (z=-0.66, p=0.5104) | Conf_z=+0.2869
(z=+2.16, p=0.0311)
  [Gate 8 — Without Lexical Controls (No Over-Partialling), Sample-Median 4-Cell &
1%/99% Winsorization]
    No-Lex S1 Breakthrough: Neg_z -> coef = +0.020324 (+0.0203, +2.03pp), SE =
0.010096 (0.0101), t = +2.0131 (+2.01), p = 0.047308 (0.0473), 95% CI = [+0.0002,
+0.0404] | Conf_z -> coef = -0.007671 (-0.0077, -0.77pp), SE = 0.010918 (0.0109), t =
-0.7026 (-0.70), p = 0.484227 (0.4842), 95% CI = [-0.0294, +0.0140]
    No-Lex S2 ABA 48h:      Neg_z -> coef = +0.001447 (+0.0014, +0.14pp), SE =
0.003258 (0.0033), t = +0.4442 (+0.44), p = 0.658013 (0.6580), 95% CI = [-0.0050,
+0.0079] | Conf_z -> coef = +0.014321 (+0.0143, +1.43pp), SE = 0.004364 (0.0044), t =
+3.2818 (+3.28), p = 0.001503 (0.0015), 95% CI = [+0.0056, +0.0230]
    Median 4-Cell (Cell2 - Cell1): S1 Breakthrough -> coef = -0.022586 (-0.0226,
-2.26pp), SE = 0.028023 (0.0280), t = -0.8060 (-0.81), p = 0.422516 (0.4225), 95% CI
= [-0.0783, +0.0331] | S2 ABA 48h -> coef = +0.019926 (+0.0199, +1.99pp), SE =
0.009731 (0.0097), t = +2.0477 (+2.05), p = 0.043711 (0.0437), 95% CI = [+0.0006,
+0.0393]
    Winsorized 1%/99% S2 ABA 48h: Neg_z -> coef = +0.003670 (+0.0037, +0.37pp), SE =
0.003402 (0.0034), t = +1.0789 (+1.08), p = 0.283739 (0.2837), 95% CI = [-0.0031,
```

```
+0.0104] | Conf_z -> coef = +0.013322 (+0.0133, +1.33pp), SE = 0.004630 (0.0046), t =
+2.8771 (+2.88), p = 0.005087 (0.0051), 95% CI = [+0.0041, +0.0225]
  [Gate 9 — Leave-One-Domain-Out (Top 10), Leave-One-Beat-Out, Leave-One-Year-Out
(2021..2025) & Post-2023]
    Drop Domain [theverge.com]: S1(Neg t=+1.54, Conf t=-0.68) | S2_ABA(Neg b=+0.0040,
t=+1.16; Conf b=+0.0144, t=+2.90)
    Drop Domain [reuters.com]: S1(Neg t=+1.68, Conf t=-0.76) | S2_ABA(Neg b=+0.0041,
t=+1.13; Conf b=+0.0137, t=+2.71)
    Drop Domain [theguardian.com]: S1(Neg t=+1.68, Conf t=-0.68) | S2_ABA(Neg
b=+0.0033, t=+0.97; Conf b=+0.0127, t=+2.61)
    Drop Domain [washingtonpost.com]: S1(Neg t=+1.65, Conf t=-0.61) | S2_ABA(Neg
b=+0.0024, t=+0.70; Conf b=+0.0141, t=+2.79)
    Drop Domain [arstechnica.com]: S1(Neg t=+1.40, Conf t=-0.42) | S2_ABA(Neg
b=+0.0033, t=+0.96; Conf b=+0.0136, t=+2.80)
    Drop Domain [bbc.com]: S1(Neg t=+1.67, Conf t=-0.21) | S2_ABA(Neg b=+0.0042,
t=+1.23; Conf b=+0.0125, t=+2.54)
    Drop Domain [nytimes.com]: S1(Neg t=+1.60, Conf t=-0.51) | S2_ABA(Neg b=+0.0043,
t=+1.26; Conf b=+0.0114, t=+2.54)
    Drop Domain [theatlantic.com]: S1(Neg t=+1.43, Conf t=-0.58) | S2_ABA(Neg
b=+0.0029, t=+0.84; Conf b=+0.0138, t=+2.75)
    Drop Domain [npr.org]: S1(Neg t=+1.90, Conf t=-0.76) | S2_ABA(Neg b=+0.0036,
t=+1.03; Conf b=+0.0135, t=+2.72)
    Drop Domain [newyorker.com]: S1(Neg t=+1.33, Conf t=-0.58) | S2_ABA(Neg
b=+0.0045, t=+1.36; Conf b=+0.0119, t=+2.55)
    Drop Beat [beat1_national_intl_wire]: S1(Neg t=+1.66, Conf t=+0.48) | S2_ABA(Neg
b=+0.0029, t=+0.68; Conf b=+0.0086, t=+1.33)
    Drop Beat [beat2_investigative_legal_policy]: S1(Neg t=+1.33, Conf t=-0.38) |
S2_ABA(Neg b=+0.0034, t=+0.94; Conf b=+0.0148, t=+2.96)
    Drop Beat [beat3_science_env_health]: S1(Neg t=+0.97, Conf t=-0.39) | S2_ABA(Neg
b=+0.0003, t=+0.08; Conf b=+0.0156, t=+3.03)
    Drop Beat [beat4_tech_society]: S1(Neg t=+1.36, Conf t=-0.58) | S2_ABA(Neg
b=+0.0049, t=+1.31; Conf b=+0.0158, t=+2.90)
    Drop Beat [beat5_culture_ideas_magazines]: S1(Neg t=+1.51, Conf t=-1.04) |
S2_ABA(Neg b=+0.0055, t=+1.63; Conf b=+0.0107, t=+2.04)
    Drop Year [2021]: S1(Neg b=+0.0171, t=+1.55; Conf b=-0.0090, t=-0.69) |
S2_ABA(Neg b=+0.0033, t=+0.81; Conf b=+0.0149, t=+2.78)
    Drop Year [2022]: S1(Neg b=+0.0208, t=+1.73; Conf b=-0.0111, t=-0.84) |
S2_ABA(Neg b=+0.0025, t=+0.68; Conf b=+0.0106, t=+2.02)
    Drop Year [2023]: S1(Neg b=+0.0127, t=+1.04; Conf b=+0.0015, t=+0.12) |
S2_ABA(Neg b=+0.0049, t=+1.26; Conf b=+0.0137, t=+2.62)
    Drop Year [2024]: S1(Neg b=+0.0179, t=+1.72; Conf b=-0.0099, t=-0.86) |
S2_ABA(Neg b=+0.0039, t=+0.89; Conf b=+0.0156, t=+2.84)
    Drop Year [2025]: S1(Neg b=+0.0117, t=+1.06; Conf b=-0.0021, t=-0.16) |
S2_ABA(Neg b=+0.0031, t=+0.87; Conf b=+0.0128, t=+2.43)
    Post-2023 Interactions: S1(Neg*Post23 t=+1.13, p=0.2628; Conf*Post23 t=-0.91,
p=0.3679) | S2_ABA(Neg*Post23 t=-0.64, p=0.5219; Conf*Post23 t=+0.32, p=0.7472)
```

```
  [Gate 10 — Strict Chronological Out-of-Sample Evaluation (Train 2021-2023 -> Test
2024-2025)]
    OOS Task [stage1_discussed_breakthrough] (N_train = 2231, N_test = 1422, W_test =
105):
      Model [constant_training_mean]: RMSE = 0.307230 (0.3072), MSE/Brier = 0.094390,
MAE = 0.193194, R2_OOS_vs_Const = +0.0000%
      Model [controls_only_no_framing]: RMSE = 0.307120 (0.3071), MSE/Brier =
0.094323, MAE = 0.193025, R2_OOS_vs_Const = +0.0714%
      Model [lexical_only_zero_llm]: RMSE = 0.307004 (0.3070), MSE/Brier = 0.094252,
MAE = 0.192935, R2_OOS_vs_Const = +0.1469%
      Model [full_4judge_llm_jev_shrunk]: RMSE = 0.306780 (0.3068), MSE/Brier =
0.094114, MAE = 0.192535, R2_OOS_vs_Const = +0.2925%
      Model [full_4judge_llm_jev_unshrunk_ols]: RMSE = 0.319036 (0.3190), MSE/Brier =
0.101784, MAE = 0.201325, R2_OOS_vs_Const = -7.8330%
      DM [llm_jev_vs_constant]: DM_t = -0.5849 (-0.58), p = 0.559885 (0.5599)
      DM [llm_jev_vs_controls_only]: DM_t = -2.0228 (-2.02), p = 0.045655 (0.0457)
      DM [llm_jev_vs_lexical_only]: DM_t = -1.7524 (-1.75), p = 0.082660 (0.0827)
      DM [llm_unshrunk_ols_vs_constant]: DM_t = +4.2580 (+4.26), p = 0.000045
(0.0000)
      DM [llm_unshrunk_ols_vs_controls_only]: DM_t = +5.4605 (+5.46), p = 0.000000
(0.0000)
    OOS Task [stage1_log_comments] (N_train = 2231, N_test = 1422, W_test = 105):
      Model [constant_training_mean]: RMSE = 1.329077 (1.3291), MSE/Brier = 1.766445,
MAE = 0.939266, R2_OOS_vs_Const = +0.0000%
      Model [controls_only_no_framing]: RMSE = 1.323174 (1.3232), MSE/Brier =
1.750790, MAE = 0.934846, R2_OOS_vs_Const = +0.8863%
      Model [lexical_only_zero_llm]: RMSE = 1.321978 (1.3220), MSE/Brier = 1.747626,
MAE = 0.933874, R2_OOS_vs_Const = +1.0654%
      Model [full_4judge_llm_jev_shrunk]: RMSE = 1.318048 (1.3180), MSE/Brier =
1.737250, MAE = 0.928592, R2_OOS_vs_Const = +1.6528%
      Model [full_4judge_llm_jev_unshrunk_ols]: RMSE = 1.357600 (1.3576), MSE/Brier =
1.843078, MAE = 0.969319, R2_OOS_vs_Const = -4.3383%
      DM [llm_jev_vs_constant]: DM_t = -3.0688 (-3.07), p = 0.002741 (0.0027)
      DM [llm_jev_vs_controls_only]: DM_t = -3.7430 (-3.74), p = 0.000298 (0.0003)
      DM [llm_jev_vs_lexical_only]: DM_t = -3.3902 (-3.39), p = 0.000988 (0.0010)
      DM [llm_unshrunk_ols_vs_constant]: DM_t = +2.3818 (+2.38), p = 0.019046
(0.0190)
      DM [llm_unshrunk_ols_vs_controls_only]: DM_t = +3.7528 (+3.75), p = 0.000288
(0.0003)
    OOS Task [stage2_dyadic_aba_reciprocity_48h] (N_train = 1355, N_test = 759,
W_test = 104):
      Model [constant_training_mean]: RMSE = 0.107603 (0.1076), MSE/Brier = 0.011578,
MAE = 0.081674, R2_OOS_vs_Const = +0.0000%
      Model [controls_only_no_framing]: RMSE = 0.107380 (0.1074), MSE/Brier =
0.011531, MAE = 0.081484, R2_OOS_vs_Const = +0.4129%
      Model [lexical_only_zero_llm]: RMSE = 0.107334 (0.1073), MSE/Brier = 0.011521,
```

```
MAE = 0.081439, R2_OOS_vs_Const = +0.4995%
      Model [full_4judge_llm_jev_shrunk]: RMSE = 0.107039 (0.1070), MSE/Brier =
0.011457, MAE = 0.081235, R2_OOS_vs_Const = +1.0461%
      Model [full_4judge_llm_jev_unshrunk_ols]: RMSE = 0.110755 (0.1108), MSE/Brier =
0.012267, MAE = 0.084402, R2_OOS_vs_Const = -5.9442%
      DM [llm_jev_vs_constant]: DM_t = -1.7466 (-1.75), p = 0.083695 (0.0837)
      DM [llm_jev_vs_controls_only]: DM_t = -2.3305 (-2.33), p = 0.021730 (0.0217)
      DM [llm_jev_vs_lexical_only]: DM_t = -2.1458 (-2.15), p = 0.034239 (0.0342)
      DM [llm_unshrunk_ols_vs_constant]: DM_t = +2.4157 (+2.42), p = 0.017466
(0.0175)
      DM [llm_unshrunk_ols_vs_controls_only]: DM_t = +3.1487 (+3.15), p = 0.002146
(0.0021)
    OOS Task [stage2_interpersonal_nested_share_48h] (N_train = 1355, N_test = 759,
W_test = 104):
      Model [constant_training_mean]: RMSE = 0.109017 (0.1090), MSE/Brier = 0.011885,
MAE = 0.083596, R2_OOS_vs_Const = +0.0000%
      Model [controls_only_no_framing]: RMSE = 0.088710 (0.0887), MSE/Brier =
0.007869, MAE = 0.066309, R2_OOS_vs_Const = +33.7855%
      Model [lexical_only_zero_llm]: RMSE = 0.088562 (0.0886), MSE/Brier = 0.007843,
MAE = 0.066214, R2_OOS_vs_Const = +34.0058%
      Model [full_4judge_llm_jev_shrunk]: RMSE = 0.087998 (0.0880), MSE/Brier =
0.007744, MAE = 0.065930, R2_OOS_vs_Const = +34.8447%
      Model [full_4judge_llm_jev_unshrunk_ols]: RMSE = 0.088442 (0.0884), MSE/Brier =
0.007822, MAE = 0.066437, R2_OOS_vs_Const = +34.1855%
      DM [llm_jev_vs_constant]: DM_t = -6.6387 (-6.64), p = 0.000000 (0.0000)
      DM [llm_jev_vs_controls_only]: DM_t = -1.6614 (-1.66), p = 0.099667 (0.0997)
      DM [llm_jev_vs_lexical_only]: DM_t = -1.3250 (-1.32), p = 0.188120 (0.1881)
      DM [llm_unshrunk_ols_vs_constant]: DM_t = -5.7731 (-5.77), p = 0.000000
(0.0000)
      DM [llm_unshrunk_ols_vs_controls_only]: DM_t = -0.5494 (-0.55), p = 0.583891
(0.5839)

[SECTION 6] RED-TEAM ADVERSARIAL STRESS-TEST ADDENDUM (CHECKS A-F)
  [Check A — Exact Additive Dyadic ABA Decomposition, Speaker Compression & Within-
Turn Evidence]
    Sample Means (N = 2114): Total_ABA_Rate = 0.211650 (21.16%) =
Unique_Dyad_ABA_Rate 0.145407 (14.54%) + Repeat_PingPong_ABA_Rate 0.066243 (6.62%,
repeat share of baseline ABA = 31.30%) | Unique_Speakers_Per_Comment = 0.678437
(67.84%) | Top2_Speaker_Share = 0.173396 (17.34%) | Mean Evidence Link Rate: within
ABA turns = 0.077696 (7.77%) vs within non-ABA turns = 0.102814 (10.28%)
    Conf_z Effect Share Decomposition: Unique_Dyad Share = 19.56% (19.6%) +
Repeat_PingPong Share = 80.44% (80.4%)
    Total ABA Rate (Conf_z):               coef = +0.013542 (+0.0135, +1.35pp), SE =
0.004804 (0.0048), t = +2.8193 (+2.82), p = 0.006002 (0.0060), 95% CI = [+0.0040,
+0.0231]
    Unique-Dyad First ABA Rate (Conf_z):    coef = +0.002649 (+0.0026, +0.26pp), SE =
```

```
0.002514 (0.0025), t = +1.0538 (+1.05), p = 0.294978 (0.2950), 95% CI = [-0.0023,
+0.0076]
    Unique-Dyad First ABA Rate (Neg_z):     coef = +0.002506 (+0.0025, +0.25pp), SE =
0.002164 (0.0022), t = +1.1581 (+1.16), p = 0.250107 (0.2501), 95% CI = [-0.0018,
+0.0068]
    Repeat-Dyad Ping-Pong Rate (Conf_z):    coef = +0.010893 (+0.0109, +1.09pp), SE =
0.003478 (0.0035), t = +3.1321 (+3.13), p = 0.002389 (0.0024), 95% CI = [+0.0040,
+0.0178]
    Repeat-Dyad Ping-Pong Rate (Neg_z):     coef = +0.000955 (+0.0010, +0.10pp), SE =
0.002166 (0.0022), t = +0.4407 (+0.44), p = 0.660529 (0.6605), 95% CI = [-0.0034,
+0.0053]
    Unique ABA Dyads Per Speaker (Conf_z):  coef = +0.013027 (+0.0130, +1.30pp), SE =
0.004574 (0.0046), t = +2.8483 (+2.85), p = 0.005525 (0.0055), 95% CI = [+0.0039,
+0.0221]
    Unique ABA Dyads Per Speaker (Neg_z):   coef = +0.001473 (+0.0015, +0.15pp), SE =
0.003373 (0.0034), t = +0.4368 (+0.44), p = 0.663377 (0.6634), 95% CI = [-0.0052,
+0.0082]
    Unique Speakers Per Comment (Conf_z):   coef = -0.017397 (-0.0174, -1.74pp), SE =
0.005210 (0.0052), t = -3.3394 (-3.34), p = 0.001254 (0.0013), 95% CI = [-0.0278,
-0.0070]
    Unique Speakers Per Comment (Neg_z):    coef = -0.001489 (-0.0015, -0.15pp), SE =
0.004012 (0.0040), t = -0.3712 (-0.37), p = 0.711440 (0.7114), 95% CI = [-0.0095,
+0.0065]
    Top-2 Speaker Comment Share (Conf_z):   coef = +0.007356 (+0.0074, +0.74pp), SE =
0.003096 (0.0031), t = +2.3755 (+2.38), p = 0.019802 (0.0198), 95% CI = [+0.0012,
+0.0135]
    Top-2 Speaker Comment Share (Neg_z):    coef = +0.000650 (+0.0006, +0.06pp), SE =
0.002464 (0.0025), t = +0.2638 (+0.26), p = 0.792587 (0.7926), 95% CI = [-0.0042,
+0.0055]
    Within-ABA Turns Ev Rate (N=2046, Conf_z):     coef = -0.004461 (-0.0045,
-0.45pp), SE = 0.005338 (0.0053), t = -0.8356 (-0.84), p = 0.405748 (0.4057), 95% CI
= [-0.0151, +0.0062]
    Within-Non-ABA Turns Ev Rate (N=2114, Conf_z): coef = -0.007916 (-0.0079,
-0.79pp), SE = 0.002637 (0.0026), t = -3.0025 (-3.00), p = 0.003526 (0.0035), 95% CI
= [-0.0132, -0.0027]
  [Check B — Stage-1 Multi-Threshold Crowd Gatekeeping Hurdle Curve (N = 3653)]
    Hurdle num_comments >= 1 (n = 1580, base = 43.25%): Joint Neg_z -> coef =
+0.030572 (+0.0306, +3.06pp), SE = 0.012027 (0.0120), t = +2.5419 (+2.54), p =
0.012861 (0.0129), 95% CI = [+0.0067, +0.0545] | Joint Conf_z -> coef = +0.029992
(+0.0300, +3.00pp), SE = 0.012764 (0.0128), t = +2.3498 (+2.35), p = 0.021128
(0.0211), 95% CI = [+0.0046, +0.0554] | Standalone(Neg_z b=+0.0495, t=+5.47,
p=0.0000; Conf_z b=+0.0526, t=+5.41, p=0.0000)
    Hurdle num_comments >= 5 (n = 598, base = 16.37%): Joint Neg_z -> coef =
+0.005253 (+0.0053, +0.53pp), SE = 0.010490 (0.0105), t = +0.5007 (+0.50), p =
0.617873 (0.6179), 95% CI = [-0.0156, +0.0261] | Joint Conf_z -> coef = +0.018497
(+0.0185, +1.85pp), SE = 0.010080 (0.0101), t = +1.8351 (+1.84), p = 0.070026
```

```
(0.0700), 95% CI = [-0.0015, +0.0385] | Standalone(Neg_z b=+0.0169, t=+2.11,
p=0.0376; Conf_z b=+0.0224, t=+2.84, p=0.0056)
    Hurdle num_comments >= 10 (n = 451, base = 12.35%): Joint Neg_z -> coef =
+0.012863 (+0.0129, +1.29pp), SE = 0.010381 (0.0104), t = +1.2390 (+1.24), p =
0.218796 (0.2188), 95% CI = [-0.0078, +0.0335] | Joint Conf_z -> coef = +0.001591
(+0.0016, +0.16pp), SE = 0.010592 (0.0106), t = +0.1502 (+0.15), p = 0.880973
(0.8810), 95% CI = [-0.0195, +0.0227] | Standalone(Neg_z b=+0.0139, t=+1.76,
p=0.0822; Conf_z b=+0.0111, t=+1.37, p=0.1751)
    Hurdle num_comments >= 15 (n = 398, base = 10.90%): Joint Neg_z -> coef =
+0.015772 (+0.0158, +1.58pp), SE = 0.010356 (0.0104), t = +1.5230 (+1.52), p =
0.131523 (0.1315), 95% CI = [-0.0048, +0.0364] | Joint Conf_z -> coef = -0.005923
(-0.0059, -0.59pp), SE = 0.011502 (0.0115), t = -0.5150 (-0.51), p = 0.607913
(0.6079), 95% CI = [-0.0288, +0.0169] | Standalone(Neg_z b=+0.0120, t=+1.65,
p=0.1022; Conf_z b=+0.0057, t=+0.70, p=0.4829)
    Hurdle num_comments >= 30 (n = 314, base = 8.60%): Joint Neg_z -> coef =
+0.015314 (+0.0153, +1.53pp), SE = 0.009439 (0.0094), t = +1.6224 (+1.62), p =
0.108472 (0.1085), 95% CI = [-0.0035, +0.0341] | Joint Conf_z -> coef = -0.008496
(-0.0085, -0.85pp), SE = 0.010289 (0.0103), t = -0.8257 (-0.83), p = 0.411334
(0.4113), 95% CI = [-0.0290, +0.0120] | Standalone(Neg_z b=+0.0100, t=+1.60,
p=0.1131; Conf_z b=+0.0028, t=+0.41, p=0.6793)
    Hurdle num_comments >= 50 (n = 241, base = 6.60%): Joint Neg_z -> coef =
+0.012835 (+0.0128, +1.28pp), SE = 0.008653 (0.0087), t = +1.4833 (+1.48), p =
0.141726 (0.1417), 95% CI = [-0.0044, +0.0300] | Joint Conf_z -> coef = -0.008923
(-0.0089, -0.89pp), SE = 0.009028 (0.0090), t = -0.9883 (-0.99), p = 0.325840
(0.3258), 95% CI = [-0.0269, +0.0090] | Standalone(Neg_z b=+0.0072, t=+1.35,
p=0.1802; Conf_z b=+0.0006, t=+0.10, p=0.9193)
    Hurdle num_comments >= 100 (n = 138, base = 3.78%): Joint Neg_z -> coef =
+0.003992 (+0.0040, +0.40pp), SE = 0.005725 (0.0057), t = +0.6973 (+0.70), p =
0.487537 (0.4875), 95% CI = [-0.0074, +0.0154] | Joint Conf_z -> coef = +0.002984
(+0.0030, +0.30pp), SE = 0.006366 (0.0064), t = +0.4687 (+0.47), p = 0.640479
(0.6405), 95% CI = [-0.0097, +0.0156] | Standalone(Neg_z b=+0.0059, t=+1.45,
p=0.1517; Conf_z b=+0.0059, t=+1.31, p=0.1935)
    Hurdle num_comments >= 150 (n = 88, base = 2.41%): Joint Neg_z -> coef =
-0.004111 (-0.0041, -0.41pp), SE = 0.004629 (0.0046), t = -0.8881 (-0.89), p =
0.377041 (0.3770), 95% CI = [-0.0133, +0.0051] | Joint Conf_z -> coef = +0.008541
(+0.0085, +0.85pp), SE = 0.004756 (0.0048), t = +1.7961 (+1.80), p = 0.076081
(0.0761), 95% CI = [-0.0009, +0.0180] | Standalone(Neg_z b=+0.0013, t=+0.42,
p=0.6763; Conf_z b=+0.0055, t=+1.76, p=0.0828)
    Intensive log_comments | num_comments >= 1 (N_initial = 1580, N_eff = 1579):
Joint Neg_z -> coef = +0.008988 (+0.0090, +0.90pp), SE = 0.071288 (0.0713), t =
+0.1261 (+0.13), p = 0.899970 (0.9000), 95% CI = [-0.1328, +0.1508] | Joint Conf_z ->
coef = -0.007769 (-0.0078, -0.78pp), SE = 0.071434 (0.0714), t = -0.1088 (-0.11), p =
0.913662 (0.9137), 95% CI = [-0.1498, +0.1343] | Standalone(Neg_z b=+0.0041, t=+0.08;
Conf_z b=-0.0014, t=-0.03)
  [Check C — H3 4-Cell Cutoff Grid Sensitivity (0.35, 0.40, 0.50, 0.60) & Cell_1 Sub-
Threshold Conflict Diagnostic]
```

```
    Cutoff [cut_0.35] (n0=962, n1=334, n2=701, n3=117; Cell_1 mean Conf = 0.1717
(0.172), Cell_1 Conf>=0.30 = 15.27% (15.3%)):
      Cell_1 Pure Threat:     coef = +0.002922 (+0.0029, +0.29pp), SE = 0.006370
(0.0064), t = +0.4587 (+0.46), p = 0.647613 (0.6476), 95% CI = [-0.0097, +0.0156]
      Cell_2 Moral Blame:     coef = +0.030183 (+0.0302, +3.02pp), SE = 0.007294
(0.0073), t = +4.1383 (+4.14), p = 0.000083 (0.0001), 95% CI = [+0.0157, +0.0447]
      Cell_3 NonNeg Debate:   coef = +0.021320 (+0.0213, +2.13pp), SE = 0.012051
(0.0121), t = +1.7692 (+1.77), p = 0.080487 (0.0805), 95% CI = [-0.0026, +0.0453]
      Contrast (Cell2-Cell1): coef = +0.027261 (+0.0273, +2.73pp), SE = 0.008653
(0.0087), t = +3.1505 (+3.15), p = 0.002259 (0.0023), 95% CI = [+0.0101, +0.0445]
    Cutoff [cut_0.40] (n0=1030, n1=336, n2=612, n3=136; Cell_1 mean Conf = 0.1953
(0.195), Cell_1 Conf>=0.30 = 25.60% (25.6%)):
      Cell_1 Pure Threat:     coef = +0.009360 (+0.0094, +0.94pp), SE = 0.006626
(0.0066), t = +1.4126 (+1.41), p = 0.161465 (0.1615), 95% CI = [-0.0038, +0.0225]
      Cell_2 Moral Blame:     coef = +0.028010 (+0.0280, +2.80pp), SE = 0.007337
(0.0073), t = +3.8175 (+3.82), p = 0.000257 (0.0003), 95% CI = [+0.0134, +0.0426]
      Cell_3 NonNeg Debate:   coef = +0.022777 (+0.0228, +2.28pp), SE = 0.011401
(0.0114), t = +1.9979 (+2.00), p = 0.048963 (0.0490), 95% CI = [+0.0001, +0.0454]
      Contrast (Cell2-Cell1): coef = +0.018651 (+0.0187, +1.87pp), SE = 0.008873
(0.0089), t = +2.1020 (+2.10), p = 0.038545 (0.0385), 95% CI = [+0.0010, +0.0363]
    Cutoff [cut_0.50] (n0=1198, n1=314, n2=413, n3=189; Cell_1 mean Conf = 0.2431
(0.243), Cell_1 Conf>=0.30 = 41.08% (41.1%)):
      Cell_1 Pure Threat:     coef = +0.013503 (+0.0135, +1.35pp), SE = 0.007576
(0.0076), t = +1.7823 (+1.78), p = 0.078308 (0.0783), 95% CI = [-0.0016, +0.0286]
      Cell_2 Moral Blame:     coef = +0.024639 (+0.0246, +2.46pp), SE = 0.009000
(0.0090), t = +2.7377 (+2.74), p = 0.007550 (0.0076), 95% CI = [+0.0067, +0.0425]
      Cell_3 NonNeg Debate:   coef = +0.021214 (+0.0212, +2.12pp), SE = 0.009619
(0.0096), t = +2.2055 (+2.21), p = 0.030154 (0.0302), 95% CI = [+0.0021, +0.0403]
      Contrast (Cell2-Cell1): coef = +0.011135 (+0.0111, +1.11pp), SE = 0.011560
(0.0116), t = +0.9633 (+0.96), p = 0.338155 (0.3382), 95% CI = [-0.0119, +0.0341]
    Cutoff [cut_0.60] (n0=1401, n1=254, n2=261, n3=198; Cell_1 mean Conf = 0.2844
(0.284), Cell_1 Conf>=0.30 = 48.03% (48.0%)):
      Cell_1 Pure Threat:     coef = +0.020219 (+0.0202, +2.02pp), SE = 0.007685
(0.0077), t = +2.6310 (+2.63), p = 0.010127 (0.0101), 95% CI = [+0.0049, +0.0355]
      Cell_2 Moral Blame:     coef = +0.019990 (+0.0200, +2.00pp), SE = 0.011491
(0.0115), t = +1.7397 (+1.74), p = 0.085583 (0.0856), 95% CI = [-0.0029, +0.0428]
      Cell_3 NonNeg Debate:   coef = +0.025408 (+0.0254, +2.54pp), SE = 0.009003
(0.0090), t = +2.8221 (+2.82), p = 0.005953 (0.0060), 95% CI = [+0.0075, +0.0433]
      Contrast (Cell2-Cell1): coef = -0.000229 (-0.0002, -0.02pp), SE = 0.014160
(0.0142), t = -0.0162 (-0.02), p = 0.987129 (0.9871), 95% CI = [-0.0284, +0.0279]
  [Check D — Journalistic Beat Domain Boundary Condition & Gelman-Stern Interactions]
    Subsample [civic_public_affairs_beats_1_2] (N_init = 962, N_eff = 962, M_domain =
41, df_t = 40, SD_Conf = 0.2900 (0.290)):
      ABA 48h Conf_z:      coef = +0.016507 (+0.0165, +1.65pp), SE = 0.006784
(0.0068), t = +2.4334 (+2.43), p = 0.019519 (0.0195), 95% CI = [+0.0028, +0.0302] |
Neg_z: coef = +0.003718 (+0.0037, +0.37pp), SE = 0.004273 (0.0043), t = +0.8702
```

```
(+0.87), p = 0.389403 (0.3894), 95% CI = [-0.0049, +0.0124]
      Nested 48h Conf_z:   coef = +0.013169 (+0.0132, +1.32pp), SE = 0.004515
(0.0045), t = +2.9171 (+2.92), p = 0.005770 (0.0058), 95% CI = [+0.0040, +0.0223]
      Evidence 48h Conf_z: coef = -0.006703 (-0.0067, -0.67pp), SE = 0.004291
(0.0043), t = -1.5623 (-1.56), p = 0.126096 (0.1261), 95% CI = [-0.0154, +0.0020]
    Subsample [specialist_sci_tech_culture_beats_3_4_5] (N_init = 1152, N_eff = 1151,
M_domain = 44, df_t = 43, SD_Conf = 0.2581 (0.258)):
      ABA 48h Conf_z:      coef = +0.010160 (+0.0102, +1.02pp), SE = 0.007270
(0.0073), t = +1.3976 (+1.40), p = 0.169407 (0.1694), 95% CI = [-0.0045, +0.0248] |
Neg_z: coef = +0.003050 (+0.0030, +0.30pp), SE = 0.004580 (0.0046), t = +0.6659
(+0.67), p = 0.509048 (0.5090), 95% CI = [-0.0062, +0.0123]
      Nested 48h Conf_z:   coef = +0.011196 (+0.0112, +1.12pp), SE = 0.004878
(0.0049), t = +2.2952 (+2.30), p = 0.026660 (0.0267), 95% CI = [+0.0014, +0.0210]
      Evidence 48h Conf_z: coef = -0.005980 (-0.0060, -0.60pp), SE = 0.003107
(0.0031), t = -1.9248 (-1.92), p = 0.060893 (0.0609), 95% CI = [-0.0122, +0.0003]
    Subsample [beat1_national_intl_wire_only] (N_init = 831, N_eff = 831, M_domain =
27, df_t = 26, SD_Conf = 0.2881 (0.288)):
      ABA 48h Conf_z:      coef = +0.019941 (+0.0199, +1.99pp), SE = 0.007227
(0.0072), t = +2.7594 (+2.76), p = 0.010466 (0.0105), 95% CI = [+0.0051, +0.0348] |
Neg_z: coef = +0.002976 (+0.0030, +0.30pp), SE = 0.004727 (0.0047), t = +0.6297
(+0.63), p = 0.534386 (0.5344), 95% CI = [-0.0067, +0.0127]
      Nested 48h Conf_z:   coef = +0.014045 (+0.0140, +1.40pp), SE = 0.005029
(0.0050), t = +2.7930 (+2.79), p = 0.009668 (0.0097), 95% CI = [+0.0037, +0.0244]
      Evidence 48h Conf_z: coef = -0.006683 (-0.0067, -0.67pp), SE = 0.004526
(0.0045), t = -1.4765 (-1.48), p = 0.151818 (0.1518), 95% CI = [-0.0160, +0.0026]
    Subsample [beat4_tech_society_only] (N_init = 459, N_eff = 458, M_domain = 11,
df_t = 10, SD_Conf = 0.2734 (0.273)):
      ABA 48h Conf_z:      coef = -0.003112 (-0.0031, -0.31pp), SE = 0.012051
(0.0121), t = -0.2583 (-0.26), p = 0.801434 (0.8014), 95% CI = [-0.0300, +0.0237] |
Neg_z: coef = +0.006288 (+0.0063, +0.63pp), SE = 0.007642 (0.0076), t = +0.8229
(+0.82), p = 0.429774 (0.4298), 95% CI = [-0.0107, +0.0233]
      Nested 48h Conf_z:   coef = +0.012181 (+0.0122, +1.22pp), SE = 0.008849
(0.0088), t = +1.3765 (+1.38), p = 0.198705 (0.1987), 95% CI = [-0.0075, +0.0319]
      Evidence 48h Conf_z: coef = -0.006426 (-0.0064, -0.64pp), SE = 0.004195
(0.0042), t = -1.5319 (-1.53), p = 0.156544 (0.1565), 95% CI = [-0.0158, +0.0029]
    Full-Panel Civic (Beats 1+2) Interaction [dyadic_aba_reciprocity_rate_48h]:
      Specialist (Beats 3+4+5) Conf_z: coef = +0.009421 (+0.0094, +0.94pp), SE =
0.006552 (0.0066), t = +1.4378 (+1.44), p = 0.154200 (0.1542), 95% CI = [-0.0036,
+0.0225]
      Conf_z x Civic Interaction:      coef = +0.006643 (+0.0066, +0.66pp), SE =
0.008149 (0.0081), t = +0.8152 (+0.82), p = 0.417281 (0.4173), 95% CI = [-0.0096,
+0.0228]
      Civic Total Conf_z Slope:        coef = +0.016064 (+0.0161, +1.61pp), SE =
0.005922 (0.0059), t = +2.7128 (+2.71), p = 0.008092 (0.0081), 95% CI = [+0.0043,
+0.0278]
    Full-Panel Civic (Beats 1+2) Interaction [interpersonal_nested_reply_share_48h]:
```

```
      Specialist (Beats 3+4+5) Conf_z: coef = +0.009927 (+0.0099, +0.99pp), SE =
0.004720 (0.0047), t = +2.1032 (+2.10), p = 0.038441 (0.0384), 95% CI = [+0.0005,
+0.0193]
      Conf_z x Civic Interaction:      coef = +0.004905 (+0.0049, +0.49pp), SE =
0.005312 (0.0053), t = +0.9234 (+0.92), p = 0.358458 (0.3585), 95% CI = [-0.0057,
+0.0155]
      Civic Total Conf_z Slope:        coef = +0.014832 (+0.0148, +1.48pp), SE =
0.004027 (0.0040), t = +3.6829 (+3.68), p = 0.000407 (0.0004), 95% CI = [+0.0068,
+0.0228]
    Full-Panel Civic (Beats 1+2) Interaction [external_evidence_link_rate_48h]:
      Specialist (Beats 3+4+5) Conf_z: coef = -0.006643 (-0.0066, -0.66pp), SE =
0.003046 (0.0030), t = -2.1810 (-2.18), p = 0.031980 (0.0320), 95% CI = [-0.0127,
-0.0006]
      Conf_z x Civic Interaction:      coef = -0.001381 (-0.0014, -0.14pp), SE =
0.004904 (0.0049), t = -0.2817 (-0.28), p = 0.778865 (0.7789), 95% CI = [-0.0111,
+0.0084]
      Civic Total Conf_z Slope:        coef = -0.008024 (-0.0080, -0.80pp), SE =
0.003836 (0.0038), t = -2.0919 (-2.09), p = 0.039465 (0.0395), 95% CI = [-0.0157,
-0.0004]
    Full-Panel Beat 4 (Tech-in-Society) Interaction
[dyadic_aba_reciprocity_rate_48h]:
      Non-Beat4 Conf_z:         coef = +0.015534 (+0.0155, +1.55pp), SE = 0.005140
(0.0051), t = +3.0221 (+3.02), p = 0.003327 (0.0033), 95% CI = [+0.0053, +0.0258]
      Conf_z x Beat4_Tech:      coef = -0.007906 (-0.0079, -0.79pp), SE = 0.008588
(0.0086), t = -0.9205 (-0.92), p = 0.359938 (0.3599), 95% CI = [-0.0250, +0.0092]
      Beat4 Total Conf_z Slope: coef = +0.007628 (+0.0076, +0.76pp), SE = 0.007958
(0.0080), t = +0.9585 (+0.96), p = 0.340555 (0.3406), 95% CI = [-0.0082, +0.0235]
  [Check E — Leave-One-Judge-Out (3-Judge JEV) & SKILL-022 Zero-Mass / Hurdle
Diagnostics]
    Zero-Mass Audit: S1 Conf==0 = 31.37%, S1 Conf<=0.05 = 43.33% | S2 Conf==0 =
23.37%, S2 Conf<=0.05 = 34.30%
    LOJO [drop_j1_llama] S2 ABA 48h: Neg3_z -> coef = +0.003517 (+0.0035, +0.35pp),
SE = 0.003330 (0.0033), t = +1.0561 (+1.06), p = 0.293966 (0.2940), 95% CI =
[-0.0031, +0.0101] | Conf3_z -> coef = +0.013418 (+0.0134, +1.34pp), SE = 0.005010
(0.0050), t = +2.6784 (+2.68), p = 0.008898 (0.0089), 95% CI = [+0.0035, +0.0234]
    LOJO [drop_j2_qwen] S2 ABA 48h: Neg3_z -> coef = +0.002267 (+0.0023, +0.23pp), SE
= 0.003439 (0.0034), t = +0.6591 (+0.66), p = 0.511635 (0.5116), 95% CI = [-0.0046,
+0.0091] | Conf3_z -> coef = +0.015692 (+0.0157, +1.57pp), SE = 0.004879 (0.0049), t
= +3.2166 (+3.22), p = 0.001843 (0.0018), 95% CI = [+0.0060, +0.0254]
    LOJO [drop_j3_deepseek] S2 ABA 48h: Neg3_z -> coef = +0.003683 (+0.0037,
+0.37pp), SE = 0.003235 (0.0032), t = +1.1384 (+1.14), p = 0.258203 (0.2582), 95% CI
= [-0.0028, +0.0101] | Conf3_z -> coef = +0.012633 (+0.0126, +1.26pp), SE = 0.004480
(0.0045), t = +2.8198 (+2.82), p = 0.005992 (0.0060), 95% CI = [+0.0037, +0.0215]
    LOJO [drop_j4_gemma] S2 ABA 48h: Neg3_z -> coef = +0.004716 (+0.0047, +0.47pp),
SE = 0.003478 (0.0035), t = +1.3559 (+1.36), p = 0.178750 (0.1787), 95% CI =
[-0.0022, +0.0116] | Conf3_z -> coef = +0.011443 (+0.0114, +1.14pp), SE = 0.004665
```

```
(0.0047), t = +2.4531 (+2.45), p = 0.016233 (0.0162), 95% CI = [+0.0022, +0.0207]
    S1 Breakthrough Two-Part Hurdle: Any_Conf>0.05 -> coef = +0.027334 (+0.0273,
+2.73pp), SE = 0.016924 (0.0169), t = +1.6151 (+1.62), p = 0.110038 (0.1100), 95% CI
= [-0.0063, +0.0610] | Conf_z -> coef = -0.012003 (-0.0120, -1.20pp), SE = 0.010783
(0.0108), t = -1.1132 (-1.11), p = 0.268802 (0.2688), 95% CI = [-0.0334, +0.0094]
    S2 ABA Two-Part Hurdle:          Any_Conf>0.05 -> coef = +0.030183 (+0.0302,
+3.02pp), SE = 0.008907 (0.0089), t = +3.3886 (+3.39), p = 0.001072 (0.0011), 95% CI
= [+0.0125, +0.0479] | Conf_z -> coef = +0.007183 (+0.0072, +0.72pp), SE = 0.005161
(0.0052), t = +1.3918 (+1.39), p = 0.167663 (0.1677), 95% CI = [-0.0031, +0.0174]
    S2 ABA Positive-Conflict Subsample (Conf > 0.05, N_init = 1389, N_eff = 1389):
Neg_z -> coef = -0.003494 (-0.0035, -0.35pp), SE = 0.005015 (0.0050), t = -0.6967
(-0.70), p = 0.487891 (0.4879), 95% CI = [-0.0135, +0.0065] | Conf_z -> coef =
+0.007377 (+0.0074, +0.74pp), SE = 0.005517 (0.0055), t = +1.3372 (+1.34), p =
0.184774 (0.1848), 95% CI = [-0.0036, +0.0183]
  [Check F — Pure Ex-Ante t=0 Stage-2 OOS (Zero Realized log(V_48) Splines, Strict
2021-2022 Inner Scaling)]
    Ex-Ante OOS Task [stage2_aba_ex_ante_t0] (N_train = 1355, N_test = 759, W_test =
104):
      Model [constant_training_mean]: RMSE = 0.107603 (0.1076), MAE = 0.081674
(0.0817), R2_OOS_vs_Const = +0.0000%
      Model [controls_only_no_framing]: RMSE = 0.107383 (0.1074), MAE = 0.081620
(0.0816), R2_OOS_vs_Const = +0.4086%
      Model [lexical_only_zero_llm]: RMSE = 0.107334 (0.1073), MAE = 0.081572
(0.0816), R2_OOS_vs_Const = +0.4997%
      Model [full_4judge_llm_jev_shrunk]: RMSE = 0.107034 (0.1070), MAE = 0.081354
(0.0814), R2_OOS_vs_Const = +1.0542%
      DM [llm_vs_const]: DM_t = -1.7917 (-1.79), p = 0.076115 (0.0761)
      DM [llm_vs_ctrl]: DM_t = -2.3515 (-2.35), p = 0.020603 (0.0206)
      DM [llm_vs_lex]: DM_t = -2.1548 (-2.15), p = 0.033509 (0.0335)
    Ex-Ante OOS Task [stage2_nested_ex_ante_t0] (N_train = 1355, N_test = 759, W_test
= 104):
      Model [constant_training_mean]: RMSE = 0.109017 (0.1090), MAE = 0.083596
(0.0836), R2_OOS_vs_Const = +0.0000%
      Model [controls_only_no_framing]: RMSE = 0.108090 (0.1081), MAE = 0.082695
(0.0827), R2_OOS_vs_Const = +1.6946%
      Model [lexical_only_zero_llm]: RMSE = 0.107984 (0.1080), MAE = 0.082589
(0.0826), R2_OOS_vs_Const = +1.8861%
      Model [full_4judge_llm_jev_shrunk]: RMSE = 0.107757 (0.1078), MAE = 0.082341
(0.0823), R2_OOS_vs_Const = +2.2984%
      DM [llm_vs_const]: DM_t = -3.4893 (-3.49), p = 0.000714 (0.0007)
      DM [llm_vs_ctrl]: DM_t = -3.2065 (-3.21), p = 0.001790 (0.0018)
      DM [llm_vs_lex]: DM_t = -2.6097 (-2.61), p = 0.010411 (0.0104)
    Ex-Ante OOS Task [stage2_evidence_ex_ante_t0] (N_train = 1355, N_test = 759,
W_test = 104):
      Model [constant_training_mean]: RMSE = 0.067817 (0.0678), MAE = 0.050846
(0.0508), R2_OOS_vs_Const = +0.0000%
```

```
      Model [controls_only_no_framing]: RMSE = 0.066595 (0.0666), MAE = 0.049613
(0.0496), R2_OOS_vs_Const = +3.5737%
      Model [lexical_only_zero_llm]: RMSE = 0.066588 (0.0666), MAE = 0.049603
(0.0496), R2_OOS_vs_Const = +3.5924%
      Model [full_4judge_llm_jev_shrunk]: RMSE = 0.066372 (0.0664), MAE = 0.049401
(0.0494), R2_OOS_vs_Const = +4.2182%
      DM [llm_vs_const]: DM_t = -4.9340 (-4.93), p = 0.000003 (0.0000)
      DM [llm_vs_ctrl]: DM_t = -2.4753 (-2.48), p = 0.014942 (0.0149)
      DM [llm_vs_lex]: DM_t = -2.5873 (-2.59), p = 0.011067 (0.0111)

================================================================================
===========
ALL STUDY #07 VERIFICATION ASSERTIONS PASSED (0 DISCREPANCIES)
================================================================================
===========
```