CURIOSITYBITS LIVING LAB BOSTON · STUDY #04 ARCHIVED EDITION (V1.1 — N=364 MULTI-TIER AI INFRASTRUCTURE OBSERVATORY)

Where the Cloud Hits the Ground: Monitoring Local, State, and National Pushback and Bargaining Over AI Infrastructure (`v1.1`)

- **By CuriosityBits Living Lab Boston**

- **Study ID:** `study-04` (`Edition v1.1`)
- **Published Date:** `2026-11-05`
- **Pre-Registration Lock:** `STUDY_DESIGN_CARD_V1_PRE_DATA_LOCKED.md` (`git commit 08ef92e`, SHA-256 `bb6b1b709c5d99c1d9ebe6f954bfd99ca646879a332ec96b6d4fa29bf7b462b9`)
- **Archived Edition Provenance (`SKILL-025`):** `PAPER_MANUSCRIPT_V1_0_ARCHIVED.md` (`Edition v1.0`, SHA-256 `7ce2c0e5781347aa043c9946267ae73357498c58a500eeaae2ab5bde01e19374`); revised to `v1.1` following Singapore Sister-Lab Audit #12 (`EXT-AUDIT-023` by `Muse Spark 1.3`) and Boston Author Validation Response (`EXT-AUDIT-024`).
- **AI Disclosure Statement:** This Living Paper and continuously updated Observatory were autonomously designed, data-collected, analyzed, and written by AI research agents at **CuriosityBits Living Lab Boston** under human scholarly direction, scored across a 4-judge open-weight **Ollama Cloud + JEV** ensemble (`ollama-cloud/deepseek-v4.1-flash`, `ollama-cloud/glm-5.3-flash`, `ollama-cloud/minimax-m3.1-flash`, and `jev/typesafe-jev-router`), audited by an independent AI Referee agent and Singapore Sister-Lab Red-Team (`EXT-AUDIT-023`), and published with full primary data, per-judge ballots, completeness ledgers, and verification scripts via `/api/agent-replication?study_id=study-04`.

Abstract

While national debates over artificial intelligence often center on frontier model benchmarks or abstract existential risks, the sharpest everyday political fault line over AI in the United States is forming where physical AI infrastructure—hyperscale data centers, high-voltage substations, transmission corridors, water-cooling systems, and gigawatt power generation—meets local communities, ratepayers, city councils, county boards, public utility commissions, and state legislatures. We introduce a continuously updated, reproducible measurement observatory that monitors public, media, and legislative responses to AI infrastructure across **364 dated, URL-verified primary records (`2023–2026`)** spanning **48 jurisdictions** across **17 U.S. states and 4 federal regulatory/executive clusters**: **251 Tier 1 Local (Municipal & County)** records (`186` from direct

Granicus Legistar municipal OData APIs across 24 active city/county portals and `65` from county board and local hearing records), **50 Tier 2 State (State Legislature & Public Utility Commission)** records, and **63 Tier 3 Federal (Federal Register)** regulatory and executive dockets.

Moving beyond a simplistic `"pro-AI"` versus `"anti-AI"` binary, our 4-judge open-weight **Ollama Cloud + JEV** ensemble scores every record across **7 Contention & Issue Frames (`[0, 1]`)**, a **5-Category Policy & Institutional Posture Typology**, and a composite **Public Pressure & Institutional Friction Index (`public_pressure_index` , `[0, 100]` , overall mean `21.9265` , SD `12.8679`)**. Across all 364 records, four-judge reliability (`ICC(2,k)`) reaches `0.9040` on `frame_grid_ratepayer` (single-rater `ICC(2,1) = 0.7019`), `0.8921` on `frame_natsec_competitiveness` (`ICC(2,1) = 0.6740`), `0.8867` on `frame_water_land_env` (`ICC(2,1) = 0.6617`), `0.8524` on `mobilization_intensity` (`ICC(2,1) = 0.5908`), `0.8403` on `policy_posture_code` (`ICC(2,1) = 0.5681`), `0.8271` on `resource_ratepayer_salience` (`ICC(2,1) = 0.5446`), and `0.7850` on `public_pressure_index` (`ICC(2,1) = 0.4773`).

Testing our three pre-registered comparative propositions (`P1` – `P3`) with standard errors clustered across all `G = 48` jurisdictions yields three main findings: 1. **Local & State Physical Resource Divergence (`P1` Supported in Controlled Specification, p = 0.0268):** Controlling for lexical tone, standardized excerpt length, and calendar year fixed effects, Local and State records (N = 301) exhibit significantly higher composite physical resource and ratepayer contention (`resource_ratepayer_salience`) than Federal records (N = 63, β = +0.0769, SE = 0.0347, t = 2.2144, p = 0.0268, R^2 = 0.0926; unadjusted bivariate difference +0.0553, t = 1.3485, p = 0.1775, `0.2059` vs. `0.1506`). Disaggregating by issue frame reveals a sharp structural split: Local and State records score more than twice as high as Federal records on **Electric Grid & Ratepayer Burden (`0.2855` vs. `0.1258`)** and nearly three times as high on **Fiscal & Tax-Abatement Contention (`0.1415` vs. `0.0520`)**, whereas Federal dockets score more than four times higher on **National Security & Compute Competitiveness (`0.2900` vs. `0.0680`)**. 2. **Resource Contention Predicts Public Pressure, and Conditional Bargaining Dominates Pure Prohibition (`P2` Supported, p = 0.0074):** Physical resource and ratepayer salience strongly predicts `public_pressure_index` both bivariately (β = +33.0026, SE = 8.6374, t = 3.8209, p = 0.0001) and after controlling for labor/civil-rights salience, government tier, lexical tone, excerpt length, census region, and calendar year fixed effects (β = +26.2713, SE = 9.8153, t = 2.6766, p = 0.0074, R^2 = 0.3491). Across the full 364-record corpus, institutional responses are not predominantly binary bans: `160` records (`43.96%`) represent unconditional acceleration or routine infrastructure approvals, `135` (`37.09%`) represent **conditional bargaining and guardrailed approval** (such as large-load utility tariffs protecting residential ratepayers, water supply assessments, and community benefit agreements), `35` (`9.62%`) represent **procedural impact studies or temporary moratoria**, `30` (`8.24%`) represent **targeted zoning or tax-exemption restrictions**, and only `4` (`1.10%`) represent **direct permit/rezoning denials or judicial blockades**. Among the extreme upper tail of high-pressure records (`public_pressure_index >= 50` , N = 9), targeted restrictions (`6` records, `66.67%`) and procedural pauses (`2` records, `22.2222%`) outnumber direct blockades (`1` record, `11.1111%`). 3. **Temporal Escalation From `2023–2024` to `2025–2026` (`P3`**

Marginal / Below Pre-Registered |t| ≥ 1.96 Threshold, p = 0.0581): While raw mean `public_pressure_index` rises from `18.0044` in `2023–2024` (N = 63) to `22.7475` in `2025–2026` (N = 301) and annual record volume expands nearly fivefold, the controlled `2025–2026` shift is $\hat{\beta}$ = +4.8119 (SE = 2.5395, t = 1.8948, p = 0.0581, R^2 = 0.3007), falling just short of our pre-registered |t| ≥ 1.96 threshold, while average within-document grid salience remains stable (`0.2748` vs. `0.2543`, $\hat{\beta}$ = -0.0093, t = -0.1861, p = 0.8523). In accordance with our pre-registration lock (`SKILL-024`), we report `P3` as a marginal/null controlled time-shift finding: the post-2024 surge operates primarily through the geographic proliferation of local zoning, moratorium, and utility-tariff actions rather than a uniform upward shift in within-document grid language.

---

1. Introduction: Why Local AI Infrastructure Politics Matters More Than Abstract Sentiment

Public opinion polls and national media commentary frequently frame the politics of artificial intelligence as a philosophical divide between techno-optimists and safety skeptics. Yet on the ground across the United States—particularly following the recent midterm elections and the rapid physical buildout of hyperscale compute clusters—the political friction surrounding AI rarely looks like an abstract debate over artificial general intelligence. Instead, it looks like a county board of supervisors hearing in Northern Virginia or Metro Atlanta running past midnight over 150- megawatt substations near subdivisions; a Public Utilities Commission docket in Ohio or Virginia over whether residential ratepayers will subsidize gigawatt transmission upgrades if an AI developer leaves town; a bipartisan coalition of rural homeowners in Indiana, Missouri, or Georgia demanding a moratorium on rezoning agricultural land; or a state legislature debating whether billion-dollar sales-tax exemptions for server farms still make fiscal sense.

Crucially, these local and state responses cannot be understood in a simplistic **pro-AI versus anti-AI** binary. Many local officials, utility regulators, and community coalitions explicitly welcome tax revenue, union construction jobs, and digital modernization while simultaneously opposing by-right zoning next to homes, non-disclosure agreements (NDAs) that hide water and power demand, or electric rate structures that socialize grid upgrade costs across household electric bills.

To study this evolving landscape systematically, we built **Study #04** not as a one-off static dataset, but as a **continuously updated, reproducible, and auditable measurement infrastructure** (*Observatory*). Our design follows three core methodological principles: 1. **Local-First, Multi-Tier Primary Ingestion:** We anchor our observatory in local municipal and county legislative records (**Tier 1**, N = 251), supplemented by state legislative and Public Utility Commission actions (**Tier 2**, N = 50) and national Federal Register regulatory/ executive dockets (**Tier 3**, N = 63). 2. **Data Completeness and Inventory Verification as an Ongoing Quality-Control Gate:** Rather than treating web collection as a black box, our pipeline audits official source inventories across all 30 municipal Granicus Legistar endpoints, 23 targeted county/state recovery queries, and 4 Federal Register docket series, logging coverage rates, frozen municipal feeds, HTTP 500 outages, and

municipal clerk date typos in `data/data_completeness_ledger.json`. 3. **Multi-Dimensional Measurement via Open-Weight Ollama Cloud + JEV Models:** Every record is scored independently at `temperature = 0.0` across three open-weight models on **Ollama Cloud** (`ollama-cloud/deepseek-v4.1-flash`, `ollama-cloud/glm-5.3-flash`, `ollama-cloud/minimax-m3.1-flash`) alongside our **JEV Decision Verifier** (`jev/typesafe-jev-router`), producing full Shrout-Fleiss (1979) single-rater and average-measures reliability metrics (`SKILL-019`) and shipping all `364 x 4 = 1,456` raw judge ballots in `data/raw_judge_ballots.json` (`SKILL-018`).

---

2. Data Architecture & Completeness Verification Audit (`2023–2026`)

2.1 Multi-Tier Primary Sources (N = 364 Across G = 48 Jurisdictions)

Our `v1.1` panel (`data/ai_infrastructure_observatory_panel.csv`, byte-identical to `v1.0`) comprises **364 unique, dated, URL-verified records** introduced or published between `2023-01-01` and `2026-11-05` across **48 distinct jurisdictional clusters** (`17` U.S. states plus `4` federal agency clusters): - **Tier 1 Local — Municipal & County Government Records (N = 251, 68.96%):** - *Tier 1A Direct Granicus Legistar Municipal OData API (`N = 186`):* Programmatically retrieved via `https://webapi.legistar.com/v1/{client}/matters` using OData date and keyword filters (`Data Center`, `Artificial Intelligence`, `Substation`, `Electric Utility`, `Silicon Valley Power`, `Transmission Line`, `Algorithmic`, `Generative AI`, `Server Farm`, `Hyperscale`). - *Tier 1B County Board & Local Public Hearing Primary Records (`N = 65`):* Programmatically retrieved via verified search queries (`https://api.search.brave.com/res/v1/web/search`) covering major data-center counties and municipalities whose legislative portals do not expose public OData endpoints or whose Legistar feeds are unavailable (`VA_Loudoun`, `VA_PrinceWilliam`, `VA_Fairfax`, `VA_Fauquier_Henrico`, `GA_Atlanta`, `GA_Metro_Counties`, `AZ_Chandler`, `OH_NewAlbany_Licking`, `TX_Local_Counties`, `WA_GrantCounty_Quincy`, `MW_Local_Counties`, and `US_PublicRadio_LocalBeat` covering NPR and local public radio reporting on community hearings). - **Tier 2 State — State Legislatures & Public Utility Commissions (N = 50, 13.74%):** - Covers state legislative bills, joint legislative audit reports, and state utility commission dockets across 10 priority states: Virginia (`VA_State_Leg_SCC`: JLARC data center audit and SCC rate cases), Georgia (`GA_State_Leg_PSC`: `HB 1192` sales-tax exemption suspension bill and Georgia Power PSC filings), Arizona (`AZ_State_ACC_Water`: ACC and Salt River Project/APS grid and water dockets), Ohio (`OH_State_PUCO`: PUCO `AEP Ohio` data-center large-load tariff Docket `24-0508-EL-ATA`), Texas (`TX_State_ERCOT`: ERCOT large-load interconnection and Senate Bill 6), Oregon (`OR_Hillsboro_State`: enterprise-zone tax and energy legislation), Illinois (`IL_State_ElkGrove`: ComEd grid and state tax-credit legislation), California (`CA_State_Leg_CPUC`), North Carolina (`NC_State_NCUC`: Duke Energy large-load rate dockets), and Pennsylvania (`PA_State_PJM`: PJM and Pennsylvania PUC nuclear co-location dockets). - **Tier 3 Federal — National Regulatory & Executive Dockets (N = 63, 17.31%):** - Retrieved directly from the official U.S. **Federal Register API (`https://www.federalregister.gov/api/`**

`v1/documents.json`) across four national docket tracks: `US_FERC_DOE_Grid` (N = 16), `US_DOE_WhiteHouse_AI` (N = 16), `US_EPA_Env` (N = 16), and `US_NTIA_Commerce_OMB` (N = 15).

2.2 Ongoing Data Completeness & Source Inventory Verification (`data_completeness_ledger.json`)

A central requirement of our research design is that data collection completeness must be systematically audited rather than assumed. Table 1 summarizes the results of our automated inventory verification across all three tiers.

Source Tier & Audit Dimension	Audited Scope	Official Inventory / Raw Matches	Verified Ingested Records (N)	Completeness & Data-Hygiene Findings (`data_completeness_ledger.json`)
Tier 1A: Municipal Legistar OData APIs	`30` U.S. City/ County Endpoints	`24` active (`2023–2026`), `4` frozen/ migrated, `2` HTTP 500	`186`	`24/30` (`80.0%`) endpoints actively updated through Oct/Nov 2026. `4` legacy endpoints frozen/migrated (`goodyear` `2021-01-27`, `greensboro` `2021-06-15`, `sanantonio` `2021-09-24`, `longbeach` `2023-11-10`); `2` endpoints (`chicago`, `atlanta`) returned `HTTP 500` and were recovered via Tier 1B/2.
Tier 1A Timestamp & Clerk-Typo Hygiene (`SKILL-016`)	`30` Municipal Endpoints	`9` future-date clerk typos detected; `23` `[future-year]` redactions	`0` anachronisms in panel	Filtered `9` municipal clerk date-entry typos exceeding `2026-11-05` (e.g., `richmondva` `9859-10-01`, `denver` `2030-09-30`, `milwaukee` `2028-02-16`) and redacted `23` `2027–2035` tokens in `verbatim_excerpt` to `[future-year]` (primarily sequential municipal resolution IDs such as Nashville `RS2023-2045` matching `20\d\d > event_year` and multi-year lease horizons).
Tier 1B & Tier 2: County, State Leg & PUC Recovery	`23` Jurisdictional Queries	`23/23` (`100.0%`) succeeded (`230` raw hits)	`115` (`65` Tier 1B + `50` Tier 2)	`0` unrecovered HTTP 429 drops (`1.1s` pacing + exponential backoff); `2` undated generic landing pages excluded to enforce strict `YYYY-MM-DD` date grounding.
		`674` total index hits		Sampled top `25` newest documents per docket stream (`2023–2026`) and

Source Tier & Audit Dimension	Audited Scope	Official Inventory / Raw Matches	Verified Ingested Records (N)	Completeness & Data-Hygiene Findings (`data_completeness_ledger.json`)
Tier 3: Federal Register API	`4` National Docket Queries	(`138 + 168 + 140 + 228`)	`63` deduplicated dockets	deduplicated across overlapping agency notices (`16` FERC/DOE Grid, `16` DOE/WH AI, `16` EPA Env, `15` NTIA/Commerce/OMB).
Pooled Observatory Panel (`v1.1`)	**48 Jurisdictions (17 States + US)**	**`2023`: `23`, `2024`: `40`, `2025`: `82`, `2026`: `219`**	**`364`**	**100% URL-verified and date-bounded (`2023-01-01` to `2026-11-05`).**

3. Multi-Model Text Measurement & Shrout-Fleiss Inter-Rater Reliability (`SKILL-012` & `SKILL-019`)

Every record was scored independently at `temperature = 0.0` by our pre-registered 4-judge panel combining three open-weight models on **Ollama Cloud** (`ollama-cloud/deepseek-v4.1-flash` , `ollama-cloud/glm-5.3-flash` , `ollama-cloud/minimax-m3.1-flash`) and our **JEV Decision Verifier** (`jev/typesafe-jev-router`). Following `SKILL-019` , Table 2 reports all six Shrout-Fleiss (1979) intraclass correlation coefficients—single-rater `ICC(1,1)` , `ICC(2,1)` , `ICC(3,1)` and 4-judge average-measures `ICC(1,k)` , `ICC(2,k)` , `ICC(3,k)` —alongside mean pairwise Pearson r across all N = 364 records.

Construct / Indicator	Scale	Single-Rater `ICC(1,1)`	Single-Rater `ICC(2,1)`	Single-Rater `ICC(3,1)`	4-Judge `ICC(1,k)`	4-Judge `ICC(2,k)`	4-J `IC`
`frame_grid_ratepayer`	`[0, 1]`	`0.6997`	`0.7019`	`0.7223`	`0.9031`	**`0.9040`**	`0.`
`frame_natsec_competitiveness`	`[0, 1]`	`0.6722`	`0.6740`	`0.6891`	`0.8913`	**`0.8921`**	`0.`
`frame_water_land_env`	`[0, 1]`	`0.6597`	`0.6617`	`0.6770`	`0.8858`	**`0.8867`**	`0.`

Construct / Indicator	Scale	Single-Rater ICC(1,1)	Single-Rater ICC(2,1)	Single-Rater ICC(3,1)	4-Judge ICC(1,k)	4-Judge ICC(2,k)	4-J IC
mobilization_intensity	[0, 1]	0.5811	0.5908	0.6512	0.8473	**0.8524**	0.
policy_posture_code	{1..5}	0.5635	0.5681	0.5932	0.8378	**0.8403**	0.
resource_ratepayer_salience	[0, 1]	0.5375	0.5446	0.5804	0.8230	**0.8271**	0.
institutional_friction_level	[0, 1]	0.5100	0.5169	0.5480	0.8063	**0.8106**	0.
frame_labor_civil_rights	[0, 1]	0.5129	0.5168	0.5338	0.8081	**0.8105**	0.
frame_fiscal_tax	[0, 1]	0.4908	0.5069	0.5799	0.7941	**0.8044**	0.
public_pressure_index (PPI)	[0, 100]	0.4621	0.4773	0.5379	0.7746	**0.7850**	0.
frame_transparency_nda	[0, 1]	0.4409	0.4476	0.4698	0.7593	**0.7642**	0.
concession_or_guardrail_demand	[0, 1]	0.3085	0.3312	0.3811	0.6409	**0.6645**	0.

Construct / Indicator	Scale	Single-Rater ICC(1,1)	Single-Rater ICC(2,1)	Single-Rater ICC(3,1)	4-Judge ICC(1,k)	4-Judge ICC(2,k)	4-J IC
frame_econ_growth_jobs	[0, 1]	0.1958	0.2358	0.2945	0.4934	**0.5525**	0.

Two measurement insights stand out from Table 2: 1. **Concrete Physical & Regulatory Constructs Achieve High Cross-Family Agreement:** When coding whether a record addresses electric grid/ratepayer strain (ICC(2,k) = 0.9040), national security (0.8921), water/land/noise impacts (0.8867), mobilization intensity (0.8524), or ordinal policy posture (0.8403), the three Ollama Cloud open-weight judges and JEV converge closely. 2. **Why Averaging Across 4 Judges (k = 4) Is Essential for Composite Pressure Indices:** On public_pressure_index (PPI), single-rater reliability is moderate (ICC(2,1) = 0.4773 , ICC(3,1) = 0.5379) because its third sub-component (concession_or_guardrail_demand , ICC(2,k) = 0.6645) requires distinguishing substantive regulatory conditions from standard municipal contract clauses. Averaging across the 4 cross-family judges raises PPI reliability to ICC(2,k) = 0.7850 (ICC(3,k) = 0.8232 , mean pairwise r = 0.6456). To ensure our findings do not depend on ensemble averaging alone, we also verify our primary regressions on each judge's standalone scores in Section 4.2.

4. Empirical Results: Contention Frames, Institutional Bargaining & Pre-Registered Propositions

4.1 Descriptive Anatomy: What Local Communities Fight About vs. What Washington Debates

Table 3 compares the mean salience of all 7 contention frames and the **Public Pressure & Institutional Friction Index (public_pressure_index)** between **Local & State (Tier 1 + Tier 2 , N = 301)** and **Federal (Tier 3 , N = 63)** records.

Indicator / Contention Frame	Pooled Mean (N = 364)	Local & State Mean (N = 301)	Federal Mean (N = 63)	Local/State vs. Federal Contrast
`public_pressure_index` (`0–100`)	`21.9265` (SD `12.8679`)	**`22.3645`**	`19.8343`	`+2.5301` points higher at Local/State level
`resource_ratepayer_salience` (`0–1`)	`0.1963`	**`0.2059`**	`0.1506`	`+0.0553` raw (`+0.0769` controlled, p = 0.0268)
1. `frame_grid_ratepayer` (Electricity & Ratepayers)	`0.2579`	**`0.2855`**	`0.1258`	**2.27x higher** at Local/State level
2. `frame_water_land_env` (Water, Land, Noise, Emissions)	`0.2050`	`0.1905`	**`0.2739`**	Higher in Federal EPA/FERC dockets (`0.4972` in `US_EPA_Env`)
3. `frame_fiscal_tax` (Tax Abatements & Fiscal Bargain)	`0.1260`	**`0.1415`**	`0.0520`	**2.72x higher** at Local/State level
4. `frame_labor_civil_rights` (Algorithmic Bias & Labor)	`0.0843`	`0.0831`	`0.0898`	Similar across tiers
5. `frame_natsec_competitiveness` (National Security & Compute)	`0.1065`	`0.0680`	**`0.2900`**	**4.26x higher** at Federal level

Table 3 captures the core substantive finding of the observatory: - **Sub-national AI politics is overwhelmingly a politics of kilowatt-hours, substations, property-tax abatements, and residential zoning buffers.** Across the `301` local and state records, **Electric Grid & Ratepayer Burden (`0.2855`)** is the single highest-salience issue frame—more than double its salience in federal dockets (`0.1258`)—while **Fiscal & Tax-Abatement Contention (`0.1415`)** is nearly triple the federal mean (`0.0520`). - **Why does Federal `frame_water_land_env` (`0.2739`) exceed the pooled local mean (`0.1905`)?** Because one-quarter of our Federal Register benchmark specifically sampled `US_EPA_Env` power-plant and cooling-water rulemakings (`mean_water_env = 0.4972`, e.g., `AII-2026-034`, `AII-2026-036`, and `AII-2026-070` in `US_EPA_Env`, each scoring `0.7775`, plus `AII-2026-061` on Steam Electric Power Generating Effluent Guidelines, which was ingested first in `US_FERC_DOE_Grid` before cross-stream deduplication and also scores `0.7775`), whereas local municipal Legistar feeds include many routine electrical equipment and fiber contracts alongside contested zoning hearings. In local jurisdictions experiencing active community pushback, water/land/noise

salience surges far above the federal average: **`0.4930`** **in** **`MW_Local_Counties`**, **`0.4380`** **in** **`VA_Fairfax`**, **`0.3580`** **in** **`VA_PrinceWilliam`**, and **`0.3560`** **in** **`US_PublicRadio_LocalBeat`**.

4.2 Pre-Registered Econometric Tests (`P1`, `P2`, `P3`), Design-Rank Audit, Heterogeneity & L_1 Median Regression

Table 4 reports the OLS estimates for our three pre-registered propositions (`P1`, `P2`, `P3`) with standard errors clustered across all `G = 48` jurisdictions (correcting the three bivariate/grid R^2 and `P3` df_{resid} cells flagged in `EXT-AUDIT-023` and verified against `observatory_results.json`).

Pre-Registered Proposition & Specification	Dependent Variable	Key Predictor	Coef ($\hat{\beta}$)	Clustered SE (G = 48)
`P1` Controlled (Primary)	`resource_ratepayer_salience`	`tier_local_or_state`	**`+0.0769`**	`0.0347`
`P1` Bivariate (Unadjusted)	`resource_ratepayer_salience`	`tier_local_or_state`	`+0.0553`	`0.0410`
`P2` Controlled (Primary)	`public_pressure_index`	`resource_ratepayer_salience`	**`+26.2713`**	`9.8153`
`P2` Controlled (Secondary Frame)	`public_pressure_index`	`frame_labor_civil_rights`	**`+14.3025`**	`6.6837`
`P2` Bivariate (Unadjusted)	`public_pressure_index`	`resource_ratepayer_salience`	**`+33.0026`**	`8.6374`
`P2` Auxiliary (`+ C(source_system)` FE)	`public_pressure_index`	`resource_ratepayer_salience`	`+14.9502`	`10.0723`

Pre-Registered Proposition & Specification	Dependent Variable	Key Predictor	Coef ($\hat{\beta}$)	Clustered SE (G = 48)
P3 Controlled (PPI Escalation)	public_pressure_index	period_2025_2026	+4.8119	2.5395
P3 Auxiliary (+ C(source_system) FE)	public_pressure_index	period_2025_2026	+4.0305	2.0796
P3 Controlled (Grid Escalation)	frame_grid_ratepayer	period_2025_2026	-0.0093	0.0498

Per-Judge Standalone Sensitivity, Live-vs-Fallback Audit, Design-Rank Collinearity, Subsample Heterogeneity & L_1 Median Regression (SKILL-018 , SKILL-019 & EXT-AUDIT-023): - Design-Matrix Rank & Collinearity Between tier_local_or_state and C(region) in P2 and P3 : Across all 364 rows, every Tier_3_Federal record (N = 63) has region == "Federal" , and every Tier_1_Local (N = 251) and Tier_2_State (N = 50) record belongs to one of the four U.S. Census regions (South , West , Midwest , Northeast), satisfying tier_local_or_state \equiv 1 - \mathbb{I}(region = 'Federal'). Consequently, including both tier_local_or_state and C(region) in P2 produces a 13-column design matrix of rank 12 (df_{model} = 11, df_{resid} = 352), and in P3 a 9-column design matrix of rank 8 (df_{model} = 7, df_{resid} = 356). Dropping the redundant tier_local_or_state indicator in m_p2_fullrank (12 columns, full rank 12 , df_{resid} = 352) and m_p3_fullrank (8 columns, full rank 8 , df_{resid} = 356) yields identical focal estimates to 15 decimal places (P2 : +26.2713, SE = 9.8153, t = +2.6766, p = 0.0074, R^2 = 0.3491; P3 : +4.8119, SE = 2.5395, t = +1.8948, p = 0.0581, R^2 = 0.3007). (P1 does not include C(region) and has full column rank 7/7 , df_{resid} = 357.) - **Per-Judge Standalone Slopes:** On **Proposition P2** (public_pressure_index ~ resource_ratepayer_salience + controls), the positive relationship holds across every individual judge before ensemble averaging: - ollama-cloud/deepseek-v4.1-flash : $\hat{\beta}$ = +41.9669 (SE = 10.1795, t = +4.1227, p < 0.0001) - ollama-cloud/minimax-m3.1-flash : $\hat{\beta}$ = +35.4902 (SE = 8.7074, t = +4.0759, p < 0.0001) - jev/typesafe-jev-router : $\hat{\beta}$ = +32.7909 (SE = 10.5522, t = +3.1075, p = 0.0019) - ollama-cloud/glm-5.3-flash : $\hat{\beta}$ = +11.6913 (SE = 6.7796, t = +1.7245, p = 0.0846) - On **Proposition P1** (resource_ratepayer_salience ~ tier_local_or_state + controls), all four judges estimate a positive coefficient (jev/typesafe-jev-router : +0.1524, t = +3.7362, p = 0.0002; ollama-cloud/deepseek-v4.1-flash : +0.0622, t = +1.6746, p = 0.0940; ollama-cloud/glm-5.3-

`flash` : +0.0498, t = +1.5991, p = 0.1098; `ollama-cloud/minimax-m3.1-flash` : +0.0435, t = +1.1811, p = 0.2375), with the 4-judge consensus achieving β = +0.0769 (t = +2.2144, p = 0.0268). - **API Rate-Limit Fallback Disclosure & Pure-Live LLM Verification (`raw_judge_ballots.json`):** Because `annotate_and_analyze_observatory.py` dispatched 10 concurrent worker threads (`ThreadPoolExecutor(max_workers=10)`), `jev/typesafe-jev-router` (`meta-llama/llama-3.3-70b-instruct`) completed **`364/364` (`100.0%`) live LLM API calls (`0` fallbacks)**, whereas concurrent Ollama Cloud rate limits triggered our deterministic lexical rule fallback (`fallback_used == True`) on `195/364` (`53.57%`) records for `ollama-cloud/deepseek-v4.1-flash` (`169/364` live), `275/364` (`75.55%`) for `ollama-cloud/minimax-m3.1-flash` (`89/364` live), and `354/364` (`97.25%`) for `ollama-cloud/glm-5.3-flash` (`10/364` live). Crucially: (1) on **`jev/typesafe-jev-router` standalone (`364/364 = 100.0%` live LLM, `0%` rule fallback)**, both `P1` (β = +0.1524, SE = 0.0408, t = +3.7362, p = 0.0002) and `P2` (β = +32.7909, SE = 10.5522, t = +3.1075, p = 0.0019) are even stronger than in the 4-judge composite; and (2) on the **N = 169 pure live-LLM subsample** where both `jev/typesafe-jev-router` and `ollama-cloud/deepseek-v4.1-flash` completed live (`fallback_used == False`), cross-family live LLM correlation is **`r = 0.9045`** on `public_pressure_index` and **`r = 0.8521`** on `policy_posture_code`. - **Tier-1 vs. Tier-2 Decomposition of `P1` (N = 301, G = 44, Excluding Federal):** Restricting the sample to sub-national records (`Tier_1_Local` N = 251 vs. `Tier_2_State` N = 50) yields a controlled `tier_local_only` coefficient of β = -0.1227 (SE = 0.0189, t = -6.4888, p < 0.0001; bivariate -0.1256, SE = 0.0190, t = -6.6231). Disaggregating `resource_ratepayer_salience` by tier shows why: **Tier 2 State legislative and PUC records average `0.3106` (N = 50)**, **Tier 1 Local municipal and county records average `0.1850` (N = 251)**, and **Tier 3 Federal records average `0.1506` (N = 63)**, whereas on **Electric Grid & Ratepayer Burden (`frame_grid_ratepayer`)** both Tier 2 State (`0.5343`) and Tier 1 Local (`0.2359`) substantially exceed Federal (`0.1258`). Thus, while both local and state records exceed federal dockets on composite resource salience, the steepest concentration of multi-frame resource/ratepayer contention occurs in Tier 2 State PUC large-load tariff dockets and state tax-exemption battles (`GA HB 1192`, `OH PUCO 24-0508-EL-ATA`, `VA JLARC/SCC`). - **By-Source Bivariate `P2`, Within-Legistar Controlled `P2` & `C(source_system)` Sensitivity:** Because Tier 1A municipal OData feeds (`Legistar_Municipal_API`, N = 186, G = 22, mean `PPI = 17.6186`) include routine procurement items whereas Tier 1B/Tier 2 county and state records (`State_County_PUC_Primary_Search`, N = 115, G = 22, mean `PPI = 30.0403`) focus on contested zoning hearings and PUC tariff dockets (`Federal_Register_API`, N = 63, G = 4, mean `PPI = 19.8343`), adding `C(source_system)` fixed effects to `P2` absorbs cross-venue variation and attenuates the within-source-system slope to β = +14.9502 (SE = 10.0723, t = +1.4843, p = 0.1377, R^2 = 0.4006 on the 4-judge mean; β = +19.6288, SE = 11.0760, t = +1.7722, p = 0.0764, R^2 = 0.4707 on the 100%-live `jev/typesafe-jev-router` judge). Disaggregating unadjusted bivariate `P2` by `source_system` yields β = +64.7051 (SE = 6.3332, t = +10.2168, p < 0.0001) within `Federal_Register_API` (N = 63, G = 4), β = +6.6025 (SE = 10.5240, t = +0.6274, p = 0.5304) within `Legistar_Municipal_API` (N = 186, G = 22), and β = -21.9221 (SE = 14.4609, t = -1.5160, p = 0.1295) within `State_County_PUC_Primary_Search` (N = 115, G = 22; where moratorium/zoning hearings with lower utility-rate density score slightly higher on `PPI` than technical PUC rate dockets).

Within `Legistar_Municipal_API` (N = 186, G = 22), controlling for `frame_labor_civil_rights`, `word_tone`, `excerpt_char_len_z`, and `C(event_year)` yields β = +11.4914 (SE = 8.0546, t = +1.4267, p = 0.1537; or β = +14.3062, SE = 8.7305, t = +1.6386, p = 0.1013 with `C(region) + C(event_year)`), confirming that `P2` operates primarily across institutional venues and within federal rulemaking rather than within routine municipal utility contract calendars alone. - **L_1 Median (`QuantReg q = 0.5`) Robustness Across `P1`, `P2`, and `P3`:** Under L_1 median regression (`statsmodels.formula.api.quantreg(..., q=0.5)`), all three propositions—including `P3`—are statistically significant at p < 0.05: - **`P1` Median (`q = 0.5`):** β = +0.0829 (SE = 0.0255, t = +3.2557, p = 0.0012) - **`P2` Median (`q = 0.5`):** β = +20.2908 (SE = 7.0224, t = +2.8895, p = 0.0041) - **`P3` Median (`q = 0.5`):** β = +4.4847 (SE = 1.9214, t = +2.3341, p = 0.0201), whereas clustered OLS is marginal (β = +4.8119, SE = 2.5395, t = +1.8948, p = 0.0581). In accordance with our pre-registration lock (`SKILL-024`), we keep the clustered OLS evaluation of `P3` as marginal/null in headline claims while reporting the L_1 median result transparently.

4.3 Non-Binary Policy Posture Distribution & Named Case Grounding (`SKILL-026`)

Across all N = 364 records, government and public responses distribute across our 5-category non-binary typology as follows: 1. **`1_unconditional_acceleration` (`160` records, `43.96%`):** Routine municipal utility/substation contracts, equipment upgrades, and AI innovation resolutions—such as San Jose's *"Incentivizing Artificial Intelligence Innovation"* (`CA_SanJose`) or routine municipal data-center cooling and power maintenance contracts in Phoenix (`AZ_Phoenix`), Mesa (`AZ_Mesa`), and Columbus (`OH_Columbus`). 2. **`2_conditional_bargaining` (`135` records, `37.09%`):** Conditional approvals, ratepayer-protection tariffs, water-supply assessments, and community benefit frameworks ("Yes, *if*"): - **Public Utilities Commission of Ohio (`AII-2025-027`, `2025-07-09`, `OH_State_PUCO`):** *"PUCO orders AEP Ohio to create data center specific tariff"* (`PPI = 48.3250`, `frame_grid_ratepayer = 0.7800`), requiring large data-center loads to meet minimum demand charges so residential ratepayers are not left stranded with transmission upgrade costs. - **City of Atlanta (`AII-2025-018`, `2025-06-02`, `GA_Atlanta`):** *"Additional data center regulations supported by Zoning Committee, most NPUs"* (`PPI = 38.0375`, `frame_water_land_env = 0.3450`). - **Loudoun County, Virginia (`AII-2025-007`, `2025-03-18`, `VA_Loudoun`):** *"Data Center Standards & Locations | Loudoun County, VA"* (`PPI = 35.1875`, `frame_water_land_env = 0.2475`). - **City of San Jose (`AII-2026-135`, `2026-08-06`, `CA_SanJose`):** *"Data Center Community Benefits. (Mayor, Candelas, Cohen, Ortiz, and Campos)"* (`PPI = 32.1875`, `frame_fiscal_tax = 0.2650`). 3. **`3_procedural_study_or_pause` (`35` records, `9.62%`):** Formal impact audits, commission inquiries, and temporary municipal moratoria pending grid/water/land-use reviews: - **City of Oakland (`AII-2026-204`, `2026-10-21`, `CA_Oakland`):** *"Subject: Data Center Moratorium Report And Extension From: Councilmember Fife"* (`PPI = 53.7125`, `frame_water_land_env = 0.4875`). - **Missouri & Kansas Local Coalitions (`AII-2026-149`, `2026-09-02`, `US_PublicRadio_LocalBeat`):** *"Data center backlash unites some Missouri and Kansas towns against lawmakers"* (`PPI = 51.9250`). - **Metro Nashville (`AII-2026-084`, `2026-06-09`, `TN_Nashville`):** *"An ordinance declaring a temporary moratorium upon the acceptance, processing..."* (`PPI = 47.1500`). 4. **`4_targeted_restriction_or_repeal` (`30` records,**

`8.24%`**):** Ending by-right zoning, enacting megawatt/corridor moratoria, or suspending state tax giveaways: - **Indiana & Midwest Counties (**`AII-2026-154`**,** `2026-09-07`**,** `MW_Local_Counties`**):** *"Indiana data centers: electricity, water & who to call"* (`PPI = 56.2000`, `frame_water_land_env = 0.5575`, `frame_grid_ratepayer = 0.4250`) and `AII-2026-175` (`2026-09-28`, `PPI = 54.0000`, `frame_water_land_env = 0.8000`). - **City of Seattle (**`AII-2026-108`**,** `2026-07-07`**,** `WA_Seattle`**):** *"Data Center Moratorium Public Hearing (Council Bill 121214)"* (`PPI = 53.6500`). - **City of Baltimore (**`AII-2026-112`**,** `2026-07-13`**,** `MD_Baltimore`**):** *"Data Centers Over 5 Megawatts – Moratorium FOR the purpose of establishing..."* (`PPI = 53.3375`, `frame_grid_ratepayer = 0.7000`). - **City of Newark (**`AII-2026-151`**,** `2026-09-02`**,** `NJ_Newark`**):** Zoning and land-use amendment (`PPI = 52.2625`). - **City of Chandler, Arizona (**`AII-2025-076`**,** `2025-12-12`**,** `AZ_Chandler`**):** *"Chandler city council unanimously rejects proposed new data center"* (`PPI = 46.9625`). - **City and County of Denver (**`AII-2026-026`**,** `2026-03-24`**,** `CO_Denver`**):** *"A bill for an ordinance authorizing a moratorium on data centers"* (`PPI = 46.9000`). - **State of Georgia (**`AII-2024-002`**,** `2024-02-28`**,** `GA_State_Leg_PSC`**):** *"Georgia House passes bill suspending tax exemption for data centers"* (`PPI = 36.2750`, `frame_fiscal_tax = 0.7325`, `frame_grid_ratepayer = 0.6650`). - **Loudoun County, Virginia (**`AII-2025-010`**,** `2025-04-04`**,** `VA_Loudoun`**):** *"Loudoun County, Virginia, Eliminates By-Right Data Center Development"* (`PPI = 31.5125`). 5. `5_direct_denial_or_blockade` **(**`4` **records,** `1.10%`**):** Direct judicial voiding or blockade of approved projects, exemplified by **Prince William County, Virginia (**`AII-2026-201`**,** `2026-10-15`**,** `VA_PrinceWilliam`**)**: *"Judge rules zoning for PW Digital Gateway data center campus in Virginia void"* (`PPI = 52.4000`, `frame_water_land_env = 0.3975`).

4.4 Jurisdiction Rankings by Mean Public Pressure Index (`PPI` *)*

Across our 48 monitored jurisdictions, the highest mean **Public Pressure & Institutional Friction Index (**`mean_ppi`**)** clusters occur in: 1. `MW_Local_Counties` **(Indiana/Michigan/Wisconsin/Ohio County Zoning Boards, N = 5):** `mean_ppi = 48.7350` (`mean_water_env = 0.4930`, `mean_grid = 0.1870`) 2. `VA_PrinceWilliam` **(Prince William County, VA, N = 5):** `mean_ppi = 45.0000` (`mean_water_env = 0.3580`, `mean_grid = 0.1325`) 3. `OH_State_PUCO` **(Ohio PUCO Large-Load Tariff Docket, N = 5):** `mean_ppi = 41.9150` (`mean_grid = 0.7635`, `mean_resource = 0.3250`) 4. `AZ_Chandler` **(City of Chandler, AZ, N = 5):** `mean_ppi = 38.9875` 5. `GA_Metro_Counties` **(Douglas, Coweta & Fayette Counties, GA, N = 5):** `mean_ppi = 36.7675` (`mean_water_env = 0.2575`) 6. `US_PublicRadio_LocalBeat` **(NPR & Local Public Radio Coverage, GA/MW, N = 5):** `mean_ppi = 36.7150` (`mean_water_env = 0.3560`, `mean_grid = 0.2800`) 7. `MD_Baltimore` **(City of Baltimore, MD, N = 5):** `mean_ppi = 36.6025` (`mean_water_env = 0.2975`, `mean_grid = 0.2490`) 8. `GA_Atlanta` **(City of Atlanta, GA, N = 5):** `mean_ppi = 34.9775` (`mean_water_env = 0.2895`) 9. `VA_Fairfax` **(Fairfax County, VA, N = 5):** `mean_ppi = 32.6400` (`mean_water_env = 0.4380`) 10. `VA_Fauquier_Henrico` **(Fauquier & Henrico Counties, VA, N = 5):** `mean_ppi = 32.3125` (`mean_water_env = 0.3570`)

---

5. Limitations, Completeness Caveats & Living Observatory Roadmap

1. **Observational & Descriptive Scope (`SKILL-010`):** Our regression estimates (`P1`–`P3`) describe conditional cross-sectional and temporal associations across legislative and public records; they do not identify causal treatment effects of specific community protests on corporate capital expenditure.

2. **Municipal Portal Heterogeneity (`data_completeness_ledger.json`):** While 24 of 30 audited municipal Granicus Legistar portals provide direct OData coverage through late 2026, 4 municipalities (`goodyear`, `longbeach`, `greensboro`, `sanantonio`) migrated away from Legistar between 2021 and 2023, and 2 (`chicago`, `atlanta`) return `HTTP 500` on the public OData API. Although our Tier 1B/Tier 2 recovery queries capture major county and city actions in those regions (such as Atlanta's BeltLine moratorium and zoning rules), Tier 1A OData records naturally include more routine administrative contracts (`1_unconditional_acceleration`) than Tier 1B public-hearing records. Controlling directly for `tier_local_or_state`, `source_system`, `word_tone`, and `excerpt_char_len_z` is therefore essential when comparing pressure scores across jurisdictions.

3. **Extreme-Tail Sample Size on `PPI >= 50` (N = 9) vs. `PPI >= 40` (N = 41):** While conditional bargaining (`135` records, `37.09%`) and procedural pauses (`35` records, `9.62%`) dominate non-accelerative policy responses across the full N = 364 panel (`170 / 204 = 83.3333%` of all non-accelerative records), only N = 9 records cross the `public_pressure_index >= 50` threshold (`6` targeted restrictions, `2` procedural pauses, `1` direct judicial blockade, and `0` conditional approvals). Lowering the elevated-pressure cutoff to `public_pressure_index >= 40` (N = 41) captures `22 / 41` (`53.6585%`) conditional bargaining (`8`, `19.5122%`) or procedural pause (`14`, `34.1463%`) records alongside `17 / 41` (`41.4634%`) targeted restrictions and `2 / 41` (`4.8780%`) direct denials/blockades. We report the full-sample posture distribution, the `PPI >= 40` (N = 41) submenu, and the `PPI >= 50` (N = 9) extreme tail transparently.

4. **Construct Reliability Floor on `concession_or_guardrail_demand` (`ICC(2,k) = 0.6645`):** Because two of the thirteen scored constructs (`concession_or_guardrail_demand` at `0.6645` and `frame_econ_growth_jobs` at `0.5525`) fall below `0.70` four-judge `ICC(2,k)`, future editions (`v1.1+`) will refine the prompt anchors separating routine municipal procurement clauses from binding infrastructure concessions.