Injection Risk Triage (IRT) Field Manual

A causal-inference pipeline for induced-seismicity management in the Permian Basin

Lewis Matthews

2026-06-02

Table of contents

Welcome

This is the field manual for **Injection Risk Triage (IRT)** — a causal-inference pipeline that answers, with regulatory-grade evidence standards, the question every saltwater-disposal permit decision in the Permian Basin implicitly turns on:

> *If this well injects less, does the earthquake risk actually go down — and by how much?*

The manual is the translation layer between the technical work (a targeted-learning analysis of 7,581 TexNet earthquakes and 918,720 well-days of Railroad Commission injection records) and the three audiences who can act on it:

- **Regulators** — who today curtail injection *after* felt events, using spatial-proximity rules, and who could instead condition permits on causal dose-response evidence *before* events occur.
- **Operators** — who need to know which of their wells carry attributable risk, what a volume reduction buys, and how to demonstrate proactive management to regulators and insurers.
- **Researchers** — who need the methods to be auditable, the failure modes documented, and the estimates reproducible.

How to read this book

If you have five minutes, read the Executive Summary. If you advise a regulator or legislator, the Regulator One-Pager is written to be lifted directly into a briefing packet. If you operate disposal wells, start with the Decision Tables. If you intend to attack the methods — please do — start with the Evidence Scoreboard and the Methods Defense appendix, which document not only the estimates but the artifacts we found in our own pipeline and how we fixed them.

The live system

This manual describes a running system, not a one-time study:

- **Interactive dashboard** — per-event well attribution, dose-response curves, and per-well volume-threshold curves, refreshed against live TexNet and RRC data: https://alphanet.tail098a15.ts.net/ (access on request: lewis@projectgeminae.net)
- **Source code** — the full pipeline, paper draft, and change log: https://github.com/Project-Geminae/induced-seismicity

1. About the Author

Who wrote this — and why the numbers can be trusted

1.1. Lewis Matthews (Corresponding Author)

Lewis Matthews is the analyst of record for the causal-inference pipeline this manual is built on: the confounder model, the targeted-learning estimators, the GPU solver that made full-panel analysis feasible, and the negative-control and sensitivity suites. Why those numbers should be trusted is a question of track record — and his runs through the three fields this project deliberately fuses:

- **Permian Basin operations, from the inside.** Five years (2016–2021) as Senior Data Scientist at CrownQuest Operating, a Midland-based operator — building statistical and machine-learning systems on production and operational data and translating them for non-technical decision-makers. The wells in this dataset are not abstractions to the author: he has sat in the rooms where injection, permitting, and curtailment decisions get made, which informs every policy translation in this book.
- **Institutional finance.** A decade (2003–2013) leading a family-office advisory practice in institutional capital markets — the discipline of fiduciary-grade evidence, and the language of the operators, insurers, and capital partners this manual is written for.
- **Causal methods.** Founder of Project Geminae (applying causal inference to industrial optimization problems) and corresponding author on SPE-228051, the targeted-maximum-likelihood analysis this pipeline extends. The estimators in this book are drawn from the current targeted-learning literature (van der Laan group, 2024–2025) and were implemented and validated against reference implementations before any result was reported.

Contact: lewis@projectgeminae.net.

1.2. A disclosure, made deliberately

The author spent five years inside a Permian Basin operator and maintains long-standing industry relationships — and is publishing evidence that supports *limits on injection.* That dual position is disclosed here because it is the right frame for reading this work: the analysis was built by someone with the operational context to get the domain details right and no incentive to overstate the seismicity risk. Where the evidence is weak or unstable, this manual says so explicitly — see the Evidence Scoreboard, which documents estimate revisions across data vintages, including an artifact we found in our own cross-validation machinery and the fix.

No funding was received from any operator, regulator, or advocacy organization for this analysis.

1.3. Project Geminae

Project Geminae is an independent research effort applying modern causal inference — targeted learning, the highly adaptive lasso, and double-machine-learning methods — to high-stakes industrial decision problems. The induced-seismicity program produces:

- **This field manual** — the policy and operations translation.
- **The IRT dashboard** — per-event, per-well causal attribution on live data.
- **The technical paper** — the SPE-facing methods contribution (regHAL-TMLE Delta-method inference and a GPU active-set solver for full-panel highly-adaptive-lasso estimation).
- **The codebase** — open at https://github.com/Project-Geminae/induced-seismicity.

1.4. A note on attribution

This manual is one part of a research program in active development. Where future editions include findings produced with external collaborators, methods reviewers, or partner institutions, the author block will reflect that. For the current edition, the authorship is Matthews.

2. Executive Summary

2.1. The one-sentence version

Cutting saltwater-disposal injection volumes measurably reduces earthquake *frequency* — how often quakes happen — within 7–19 km of disposal wells in the Permian Basin, and the evidence for this meets a causal standard that the current regulatory process does not require but should.

2.2. The five-minute version

What we did. We linked every TexNet-catalogued earthquake in the Permian Basin (7,581 events, 2017–2026) to every saltwater-disposal well within 20 km (795 wells, 918,720 well-days of daily Railroad Commission injection records) and asked a counterfactual question: *what would seismicity have been if each well had injected 10% less?* Answering that question credibly requires controlling for the things that make high-injection wells different from low-injection wells — fault proximity, formation depth, well age, neighborhood injection — and we use targeted learning, the same family of methods used in FDA-grade epidemiology, to do it.

What we found.

1. **The effect is real and positive at pressure-diffusion distances.** Pooled across the 7–19 km band — the distances where pore-pressure diffusion operates on a one-year timescale — a 10% volume increase causes a statistically significant increase in expected seismicity (basin-scale combined test: $z = 4.25$, p 2×10 under the current data vintage; significant in every vintage tested once a documented cross-validation artifact was repaired).
2. **The mechanism is frequency, not magnitude.** Essentially all of the causal effect operates through *how often* events occur, not *how large* a given event is. A volume cap is therefore an event-rate-reduction tool. This is the single most policy-relevant finding: it means volume management reduces the number of felt-event triggers that today's reactive curtailment responds to.
3. **Per-well attribution is operational.** For any catalogued event, the pipeline ranks every well within radius by its causal contribution, with honest confidence intervals — and for any well, it produces a volume-threshold curve: the injection level above which that specific well's marginal contribution to expected seismicity exceeds a chosen risk tolerance.
4. **Near-field estimates (1–6 km) are inconclusive** and we say so. At very short distances the positive-event counts are too small for the basis-function machinery to behave stably; we document the sensitivity rather than reporting a fragile number.

What it means for policy. The Railroad Commission's current Seismic Response Area protocol is reactive: events happen, then injection is curtailed nearby. The evidence in this manual supports a complementary *proactive* instrument: permit conditions tied to causal dose-response thresholds, applied before events occur, targeted at the wells whose attribution and threshold curves identify them as marginal-risk carriers. Model permit-condition language is provided in the appendix.

What it means for operators. The same pipeline that supports regulation also supports defense: an operator whose wells sit below their threshold curves has auditable, causal-standard evidence of proactive risk management — for regulators, insurers, and litigation.

2.3. What this evidence is not

It is not a claim that any specific earthquake was caused by any specific well; per-event attributions are model-based estimates with stated uncertainty. It is not a near-field (under 7 km) claim. And it does not yet model the operator-feedback loop (operators reducing volumes in response to nearby events), which is the program's next methodological step — documented openly in the Methods Defense.

Part I.

The Evidence at a Glance

3. Regulator One-Pager

The case for causal permit conditions, on one page

Written to be lifted directly into a briefing packet. Every claim links to the underlying evidence.

3.1. The problem with the current standard

Texas regulates injection-induced seismicity reactively. Under the Railroad Commission's Seismic Response Area (SRA) protocol, a felt earthquake triggers curtailment of disposal wells within a radius. This works — but it pays for risk reduction *after* the seismic event, with a blunt spatial rule that treats every well within the circle identically, regardless of whether that well's injection actually contributes to seismic hazard.

3.2. What the evidence now supports

A causal analysis of nine years of TexNet earthquake data and daily RRC injection records (7,581 events; 795 disposal wells; 918,720 well-days) establishes:

Finding	Evidence grade
Injection volume **causes** increased seismicity at 7–19 km	Basin-scale combined test, z = 4.25, p 2 × 10 ; direction stable across every data vintage and estimator tested
The mechanism is **event frequency**, not magnitude	~100% of the pooled causal effect operates through P(event), not E[magnitude] — stable across vintages
Per-well causal attribution is **feasible today**	Honest-forest per-well contributions with confidence intervals, live on the dashboard for every catalogued event
Per-well **volume thresholds** are computable	Dose-response curves per well identify the injection level at which marginal risk exceeds tolerance

In plain terms: **a 10% volume reduction at pressure-diffusion distances reduces how often earthquakes happen**, and we can now say *which wells* carry the marginal risk, instead of curtailing every well in a circle.

3.3. The instrument this enables

A **causal permit condition**: disposal permits in seismically active areas carry a volume schedule derived from the well's position on its dose-response curve, reviewed annually against refreshed data. Compared to SRA curtailment it is:

- **Proactive** — applied before felt events, not after.
- **Targeted** — calibrated per well, not per circle. Wells with no attributable risk are not penalized.
- **Auditable** — the dose-response machinery, its assumptions, and its failure modes are public and reproducible (github.com/Project-Geminae/induced-seismicity).
- **Defensible for operators** — compliance generates causal-standard evidence of proactive management.

Draft regulatory text is provided in Appendix A — Model Permit Language.

3.4. What we are *not* claiming

- No specific earthquake is attributed to a specific well with certainty; attributions carry stated uncertainty.
- Near-field effects (under 7 km) are statistically inconclusive in this data and are not a basis for the proposed instrument.
- The headline effect size varies across data vintages (the *direction* and *mechanism* do not); the instrument is therefore designed around rankings and thresholds, which are stable, rather than point estimates, which move.

3.5. Contact

Lewis Matthews · Project Geminae · lewis@projectgeminae.net Live dashboard access on request.

4. Evidence Scoreboard

Every estimator, every vintage, including the one we had to fix

Trust in this program should rest on the audit trail, not on any single number. This scoreboard reports every basin-scale estimate produced to date — including the revision history and the artifact we found in our own pipeline.

4.1. The estimators

Two independent estimators target the basin-scale effect of a 10% volume shift, pooled by inverse variance across the 13 pressure-band radii (7–19 km):

- **Estimator A — regHAL-TMLE Delta-method** (Li, Qiu, Wang & van der Laan 2025, arXiv:2506.17214). Targets expected maximum local magnitude directly. Runs on a ~50,000-row cluster-aware subsample (HAL basis enumeration is the constraint).
- **Estimator B — full-panel hurdle HAL-TMLE.** Decomposes the effect into frequency (does an event occur?) and magnitude (how large, given occurrence) channels. Runs on the full panel (459,105 spatially deduplicated rows) using a GPU active-set solver built for this project.

4.2. The scoreboard

Date	Estimator	(pooled, 7–19 km)	z	p	Status
2026-04	A (regHAL-TMLE, n=50k)	$+7.65 \times 10^{3}$	+3.38	7.2×10	Published headline
2026-04	B (full-n hurdle CV)	$+4.81 \times 10$	+8.35	$< 10^{1}$	Superseded — see artifact note
2026-05-06	B (May data, pre-patch)	-6.6×10	−1.44	0.15	**Artifact** — documented below
2026-05-06	A (May data rerun)	$+1.03 \times 10^{3}$	+0.64	0.52	Diagnostic only

Date	Estimator	(pooled, 7–19 km)	z	p	Status
2026-05-08	**B (May data, patched CV)**	**+8.83 × 10**	**+4.25**	**2.15 × 10**	**Current best estimate**

4.3. What survived every vintage — and what didn't

Stable across all vintages and estimators:

- The **direction**: positive. More injection → more seismicity at pressure-diffusion distances.
- The **mechanism**: frequency-dominant. The effect operates through event *rates*, not event *magnitudes*.
- The **spatial structure**: effects concentrate at 7–19 km, consistent with pore-pressure diffusion on a one-year timescale; near-field (1–6 km) inconclusive.

Not stable: the headline point estimate, which moved with data vintage and with a cross-validation repair. Policy artifacts in this manual are therefore built on rankings, thresholds, and mechanism — not on any single value.

4.4. The artifact we found, and the fix

In the May 2026 data refresh, Estimator B's pooled estimate collapsed to near-null with a sign flip in the 11–15 km band. Diagnosis: the Stage-1 (frequency-channel) cross-validation was selecting a regularization penalty large enough to **prune every basis function** — leaving the frequency channel with literally no model, which forces its contribution to zero by construction and leaves only a noisy magnitude residual.

The fix (public in the repository, commit `ef51baf`): a denser, wider penalty grid, plus an *active-floor* selection rule — the cross-validation may only select penalties that retain a working model (median active basis count 5 across folds). Under the patched selection, all 13 radii produce stable models (active sets of 1,365–1,407 basis functions, versus 0–37 before) and uniformly positive estimates.

We report this openly for a reason: **a pipeline that can detect and repair its own artifacts is the thing a regulator should demand** before conditioning permits on model output. The repair is itself now a unit test (`gpu_hal/tests/test_active_floor.py`) that prevents regression.

4.5. Why Estimator A weakened in May — resolved

The May panel doubled in size (451k → 918k well-days) not because of one new month of data, but because the refreshed RRC export **backfilled historical records to 2016**. Estimator A's fixed-size subsample now covers ~5% of the panel instead of ~11%, with a different confounder distribution. A rolling-window analysis (26 windows, 2020–2026) shows Estimator A has *never* had single-radius, single-window detection power — its April significance was always an emergent property of pooling 13 radii. Estimator B, which uses the full panel and has stable per-radius behavior, is the more reproducible instrument and is the program's primary going forward.

4.6. Reproduce it

Every number in this table regenerates from public code against public data (TexNet catalog; RRC H-10 filings):

```
git clone https://github.com/Project-Geminae/induced-seismicity
# See CHANGES.md for the full revision history with per-date entries.
```

5. Decision Tables

What an operator, a regulator, or an insurer does with this — by situation

This chapter translates the pipeline's outputs into actions. Each table is keyed to a situation, names the dashboard artifact that informs it, and states what the evidence does and does not support.

5.1. For operators

Situation	Dashboard artifact	Supported action	Not supported
Planning a new disposal well	Dose-response curve at candidate location's radius band	Site/size the well below the volume where marginal seismicity contribution becomes material	Treating the curve as a guarantee — it is a population estimate with CIs
Felt event near your well	Per-event attribution panel (click the event)	Check your well's ranked contribution and CI before assuming liability; wells with CI crossing zero have no resolvable contribution	Using a low ranking as proof of no contribution
Approaching a threshold	Per-well volume-threshold curve (click your well)	Trim volumes to re-enter the supported region; document the action	Extrapolating the curve beyond the red-shaded support region
Insurance / ESG reporting	Threshold-curve compliance over time	Causal-standard evidence of proactive risk management	Claims about magnitude reduction (the effect is frequency-channel)
SRA curtailment dispute	Attribution + threshold curve for the curtailed well	Evidence that a specific well's contribution is/isn't resolvable from data	Re-litigating the SRA itself — that's the regulator's instrument

5.2. For regulators

Situation	Artifact	Supported action	Not supported
Permit application in active area	Dose-response curve, radius band of application	Volume schedule conditioned on the curve's supported region	Conditions tied to a single point estimate (use percentile bands)

Situation	Artifact	Supported action	Not supported
Post-event response	Per-event attribution ranking	Targeted review of high-attribution wells instead of (or before) blanket-radius curtailment	Enforcement against a single well on attribution alone
Annual SRA review	Refreshed combined test + per-radius profile	Evidence-based widening/narrowing of response areas	Near-field (<7 km) rule-making from this data — inconclusive
Rule-making record	Evidence Scoreboard + Methods Defense	A documented, reproducible causal standard with stated failure modes	Presenting any single as settled magnitude

5.3. For insurers and lenders

Situation	Artifact	Supported use
Underwriting an operator	Portfolio of per-well threshold positions	Risk-differentiate premiums on causal exposure, not just well count and location
Post-event claim	Attribution panel with CIs	Independent, model-based estimate of contribution, with uncertainty, from public data
Portfolio stress test	Dose-response curves by radius band	What a basin-wide volume-growth scenario implies for event frequency

5.4. The one rule that governs every table

Rankings and thresholds are robust; point magnitudes are not. The Evidence Scoreboard documents how the pooled effect size moved across data vintages while the direction, mechanism, and spatial structure held. Every action in these tables is therefore designed to depend on *which wells*, *which direction*, and *which side of a threshold* — never on the third decimal of .

Part II.

Using This Work

6. Reading the Numbers

A plain-language guide to the statistics in this book

This chapter is for readers who want to follow the evidence without a statistics background. It covers the five concepts that appear throughout the manual and the dashboard.

6.1. “Causal” versus “correlated”

Wells that inject more sit in different places than wells that inject less — closer to faults, in different formations, with different neighbors. A correlation between injection and earthquakes could just be geography. A **causal** estimate answers a different question: *for the same well, in the same place, what changes if only the volume changes?* The methods in this pipeline (targeted learning) adjust for the measurable differences between high- and low-injection wells — fault distance, depth, well age, neighborhood injection — so that the comparison is volume-against-volume, not place-against-place.

6.2. (“psi”) — the effect size

is the answer to a specific what-if: *if every well in the band had injected 10% more, how much would expected seismicity change?* It is measured in units of expected local magnitude per well-day, which makes the raw number small (e.g. $+8.8 \times 10$) — but small-per-well-day across 800 wells and nine years is a material amount of seismicity. What matters for decisions is the **sign** (positive = more injection, more quakes), the **confidence interval** (does it exclude zero?), and the **comparison across wells** (which wells carry the effect).

6.3. z and p — how sure are we?

The z-score counts how many standard errors the estimate sits away from zero; the p-value translates that into the probability of seeing such an estimate if the true effect were zero. The current basin-scale result ($z = 4.25$, p 2×10) means: if injection truly had no effect, data like ours would occur about twice in a hundred thousand tries. Individual radii are weaker (z around 1); the strength comes from 13 radii agreeing in direction — which is itself the expected signature of a real physical mechanism operating across a distance band.

6.4. Frequency versus magnitude — the mechanism

Total seismicity can change two ways: more events, or bigger events. The pipeline separates these. The finding — stable in every data vintage — is that volume operates almost entirely on **frequency**: more injection means *more* earthquakes, not *larger* ones. For policy this is good news: event frequency is exactly what reactive curtailment programs trigger on, so volume management reduces the thing the existing system is built to respond to.

6.5. Confidence intervals on the dashboard

Every per-well attribution carries a 95% confidence interval. Reading them honestly:

- **CI entirely above zero** — the data resolve a positive contribution from this well.
- **CI crossing zero** — the data cannot distinguish this well's contribution from nothing. This is not evidence of innocence; it is absence of resolvable evidence.
- **Red-shaded regions** on dose-response curves mark volumes beyond the 99th percentile of observed data. Estimates there are extrapolation and are displayed only to show where the supported region ends.

6.6. What would change our minds

Good evidence states its falsifiers. The program's standing list:

1. A data vintage in which the patched pipeline produces a *negative* pooled pressure-band estimate with stable active sets.
2. A demonstration that the operator-feedback loop (operators cutting volume after nearby events) reverses the sign when modeled longitudinally — this is the next planned methods step.
3. A negative-control failure: detectable "effects" of injection on earthquakes too distant or too early for any physical mechanism.

None of these has occurred. The day one does, the Evidence Scoreboard reports it.

7. How to Cite

7.1. This manual

Matthews, L. (2026). *Injection Risk Triage (IRT) Field Manual: A causal-inference pipeline for induced-seismicity management in the Permian Basin.* Project Geminae. https://github.com/Project-Geminae/induced-seismicity

7.2. The technical paper

The methods contribution (regHAL-TMLE Delta-method inference for shift interventions on injection volume, and the GPU active-set solver for full-panel highly-adaptive-lasso estimation) is described in the paper draft in the repository (`PAPER_DRAFT.md`), which extends:

Matthews, L. et al. SPE-228051: *Targeted Maximum Likelihood Estimation for Induced Seismicity Attribution.*

7.3. Key methods literature

- Li, Z., Qiu, H., Wang, G., & van der Laan, M. (2025). Delta-method inference for regularized HAL. arXiv:2506.17214.
- van der Laan, M. & Rose, S. (2011). *Targeted Learning.* Springer.
- Nugent, J. et al. (2024). Two-Stage TMLE. arXiv:2208.09508.
- Xiao, et al. (2025). MSM-IPTW for induced seismicity under time-varying confounding. arXiv:2510.16360. (The complementary longitudinal approach; see the Methods Defense appendix for positioning.)

7.4. Data sources

- **TexNet Earthquake Catalog**, Bureau of Economic Geology, UT Austin.
- **Railroad Commission of Texas**, H-10 injection filings (daily volumes).

Both are public. The processed panel and all derived artifacts regenerate from the repository pipeline.

A. Appendix A — Model Permit-Condition Language

Draft regulatory text for a causal volume condition on saltwater-disposal permits. Written for adaptation by Commission staff or legislative counsel; bracketed terms are policy choices, not analytical outputs.

A.1. A.1 Definitions

Dose-response curve. For a disposal well, the estimated relationship between 365-day cumulative injection volume and expected seismicity within [7–19] km, produced by a targeted-learning estimator meeting the standards of §A.4, with 95% confidence bands, evaluated only within the supported data region (volumes at or below the 99th percentile of observed well-volumes).

Volume threshold. The lowest 365-day cumulative volume at which the lower bound of the well's dose-response confidence band exceeds [the Commission-adopted risk tolerance].

Attributable-risk ranking. For a catalogued seismic event, the ranking of wells within [19] km by estimated causal contribution, with confidence intervals.

A.2. A.2 Permit condition (new and amended permits)

1. For any disposal permit within [a Seismic Response Area or a Commission-designated monitoring area], the permitted 365-day cumulative volume shall not exceed the well's volume threshold as most recently computed under §A.4.
2. Where the well's dose-response curve does not resolve a threshold within the supported data region, the permit shall instead carry [the area's standard volumetric condition], and the well shall be re-evaluated at each annual refresh.
3. A permittee may at any time submit volumes below the threshold as evidence of proactive seismic-risk management in any Commission proceeding.

A.3. A.3 Event response (complement to existing SRA protocol)

1. Upon a catalogued event of [M3.0]+ within a monitored area, Commission staff shall obtain the attributable-risk ranking for the event.

2. Wells whose contribution interval excludes zero shall receive priority review. Curtailment of wells whose interval includes zero may proceed under existing protocol but shall be flagged for threshold re-evaluation rather than presumptive contribution.

A.4. A.4 Evidence standards

A dose-response curve or attribution is cognizable under this section only if produced by a method that:

1. **Adjusts for confounding** between injection volume and location (minimally: fault distance, formation depth, well age, and neighboring injection) using a doubly-robust estimator with published convergence guarantees;
2. **Publishes its uncertainty** — confidence intervals constructed with cluster-robust variance at the well level;
3. **Declares its supported region** and renders no operative estimate by extrapolation;
4. **Is reproducible** — code and processed data sufficient to regenerate every operative number are available to the Commission and, [save for trade-secret well identifiers], to the public;
5. **Maintains a revision record** documenting every material change in estimates across data refreshes, including artifacts detected and repaired.

Note: standard 5 is deliberate. The program maintaining this manual has itself detected and repaired a cross-validation artifact (see the Evidence Scoreboard); a method without a revision record has not been tested hard enough to be load-bearing in a permit.

A.5. A.5 Annual refresh

Thresholds and rankings shall be recomputed against refreshed TexNet and H-10 data at least annually. A permittee whose threshold rises may apply the increase immediately; a threshold reduction takes effect at the next permit anniversary, with [180 days] notice.

B. Appendix B — Methods Defense

The hard questions, answered before they're asked

This appendix collects the strongest objections to the analysis and the program's answers. The full technical treatment is in the repository's `PAPER_DRAFT.md` (§6, *Diagnostics, sensitivity, and limitations*); this is the briefing-room version.

B.1. "Correlation isn't causation — wells sit where the faults are."

That is precisely the problem the estimator class is built for. Targeted learning adjusts for fault distance, fault-segment density, formation depth, well age, and neighborhood injection before comparing volumes. The residual question — unmeasured confounding — is addressed by negative controls and by the spatial signature: the effect concentrates at pore-pressure-diffusion distances (7–19 km) and vanishes in the far field, which is the physical fingerprint of a real mechanism, not of geographic confounding.

B.2. "Your headline number changed. Why should anyone rely on it?"

Because we publish the revision record and design the policy instruments around what *didn't* change. Across every data vintage and estimator: the direction (positive), the mechanism (frequency-channel), and the spatial band (7–19 km) held. The pooled magnitude moved — and the Evidence Scoreboard documents each movement, including a cross-validation artifact we found and repaired ourselves. The model permit language (§A.4.5) makes a revision record an evidence *requirement* for exactly this reason.

B.3. "The per-radius estimates aren't individually significant."

Correct, and disclosed. Individual radii carry z   1; the basin-scale result is the inverse-variance pool of 13 correlated radii agreeing in direction. This is the appropriate inference for a mechanism that operates across a distance band, and it is why the policy artifacts use rankings and thresholds rather than single-radius point estimates.

B.4. "You don't model operators reacting to earthquakes."

True, and it is the program's most important open limitation. If operators cut volumes after nearby events, a cross-sectional estimator can *understate* the true effect (the feedback loop biases toward null). Xiao et al. (2025, arXiv:2510.16360) demonstrate this longitudinal bias in the Fort Worth Basin with marginal structural models. Note the direction: correcting for the feedback loop would most plausibly *strengthen* the causal case for volume management, not weaken it. A longitudinal (LTMLE) extension is the program's next methods milestone.

B.5. "Why should a 50,000-row subsample or a GPU solver be trusted?"

The two estimators cross-check each other. Estimator A (subsample, reference R implementation of HAL) and Estimator B (full panel, our GPU solver) agree on direction, mechanism, and spatial structure. The GPU solver was validated to 7-digit agreement against the reference implementation on synthetic problems before any production use, and its cross-validation selection rule carries a regression test born from the one artifact it produced.

B.6. "Near-field is where the public worries. Why no claim under 7 km?"

Because the data don't support one. At 1–6 km, positive-event counts are small and the basis machinery is demonstrably unstable (sign flips under basis-degree changes, documented in `PAPER_DRAFT.md` §6.4). Reporting a fragile near-field number would be the fastest way to discredit the stable pressure-band result. The program reports what holds and labels what doesn't.

B.7. "Who paid for this?"

No operator, regulator, or advocacy organization. See About the Author for the dual-role disclosure: an analyst who spent five years inside a Permian Basin operator, publishing evidence that supports injection limits, with the incentive structure that implies.