FDARadar: Anytime-Valid Sequential Pharmacovigilance under Differential Underreporting with Machine-Checked Manski–Rosenbaum Envelopes and Minimax SIPS_∞ Gating

Fabien Pierret*

November 27, 2026

Abstract

Post-marketing drug safety surveillance in the FDA Adverse Event Reporting System (FAERS) suffers from three coupled statistical failures that cause both delayed detection of lethal toxicities and costly false alarms for regulators and healthcare investors: (i) repeated quarterly screening with static disproportionality metrics (PRR, ROR, BCPNN) violates Type-I error control under continuous peeking; (ii) differential underreporting, media notoriety bias, and post-launch Weber waves severely distort 2×2 contingency tables; and (iii) compensatory composite safety scores mask acute single-pillar emergencies such as Class I/NDMA contamination recalls. We introduce **FDARadar**, a unified sequential pharmacovigilance framework adhering to strict three-layer epistemic separation. First, we establish four exact deterministic theorems—machine-checked in **Lean 4.36.0** with **48 public theorems** (476 total kernel theorems including private lemmas), **zero `sorry`**, zero external dependencies, standard axioms only (`propext`, `Classical.choice`, `Quot.sound`), and audited by two independent adversarial agents: (1) the exact reporting-invariance dichotomy between ROR and PRR plus sharp Manski–Rosenbaum Γ-sensitivity envelopes (both 1-parameter and general 2-parameter (Γ_1, Γ_0) without background parity, attained under genuine reporting probabilities $\rho \in (0,1]^4$) and the exact unity breakdown identity $\Gamma^*_{\mathrm{PRR}}(1) = \mathrm{ROR}_{\mathrm{obs}}$; (2) a Γ-robust composite-null quarterly binomial-batch mixture e-process with inductive optional-stopping conservation and exact finite-horizon Ville maximal inequality $\mathbb{P}(\exists t \le T,\, M_t \ge 1/\alpha) \le \alpha$ (with sharpness equality); (3) the pointwise and global IFF uniqueness of the 4-pillar Chebyshev minimax gate $\mathrm{SIPS}_\infty(r) = \|r\|_\infty$ together with universal single-pillar blindness of all finite-p weighted L_p power means; and (4) exact two-quarter double-scaling invariance of reaction-share and PRR velocity under arbitrary Weber/media surges. Second, on a reproducible retrospective cohort of **20 real openFDA FAERS, FDA Enforcement, FDA SPL Labeling, ClinicalTrials.gov, and ChEMBL** drug–event series (10 historical post-2004 FDA safety catalysts, 8 matched therapeutic-class negative controls, and 2 live 2026 watchlists including Semaglutide/Ozempic), FDARadar achieves **90.0% sensitivity (9/10)** with a median early-detection lead time of **+8.0 quarters (2.0 years; mean +11.9 quarters)** prior to official FDA action—versus 30.0% for L_1 and 50.0% for L_2—while cutting empirical negative-control false positives **6**× relative to repeated ROR and BCPNN (**12.5%** vs. **75.0%**) and maintaining **0.0%–0.25%** 40-quarter FWER across differential notoriety stress tests where naive ROR and BCPNN collapse to **100.0%**.

Keywords: Pharmacovigilance, openFDA FAERS, Anytime-Valid Inference, E-Process, Ville's Inequality, Manski–Rosenbaum Bounds, Notoriety Bias, Weber Effect, Minimax Gating, Lean 4 Formal Verification.

*Independent Researcher, WallStreetDebug, New York City, NY, USA. ORCID: 0009-0001-1223-2170. Email: `fabien84.p@me.com`. Live Edge Auditor: `https://fda.wallstreetdebug.com`. MIT Code & Lean 4 Proofs: `https://github.com/fabien-pierret/fdaradar-safety`.

1 Introduction

Approximately one third of FDA-approved prescription drugs incur post-market safety actions—including Boxed Warnings, safety communications, or market withdrawals—years after initial marketing authorization. For patients, delayed signal detection translates directly into preventable morbidity and mortality; for biopharmaceutical sponsors and institutional healthcare investors, unanticipated FDA safety actions routinely trigger abrupt equity drawdowns of 30% to 70%. Yet routine quantitative surveillance of spontaneous reporting systems (SRS) such as the U.S. FDA Adverse Event Reporting System (FAERS) remains dominated by static 2×2 contingency-table metrics: the Proportional Reporting Ratio (PRR) [6], Reporting Odds Ratio (ROR) [17], Bayesian Confidence Propagation Neural Network (BCPNN) [2], and Multi-item Gamma Poisson Shrinker (MGPS/EBGM) [5].

As highlighted in the 2024 READUS-PV reporting guideline [7] and the 2026 review of spontaneous reporting data mining by Tan et al. [19], applying static disproportionality algorithms sequentially over quarterly FAERS releases creates three structural vulnerabilities:

1. **Continuous-peeking multiplicity inflation.** Repeatedly evaluating fixed-horizon confidence intervals or χ^2 thresholds every quarter without optional-stopping protection drives cumulative family-wise error rates (FWER) far above nominal α [1]. While sequential probability ratio tests such as MaxSPRT and cMaxSPRT [3, 10] serve longitudinal EHR networks (e.g., FDA Sentinel) with known exposure denominators and pre-specified maximum horizons $T_{\max}$, spontaneous FAERS monitoring lacks exposure denominators and operates indefinitely without a fixed terminal horizon.
2. **Differential underreporting, notoriety bias, and the Weber effect.** Media publicity, litigation, and regulatory alerts stimulate reporting asymmetrically [13, 9]. Even a modest $1.5\times$ drug×event notoriety bias drives repeated ROR and BCPNN false-alarm rates to 100% over a 10-year monitoring horizon. Conversely, uniform event-level over-reporting of severe outcomes attenuates PRR toward unity.
3. **Single-pillar crisis masking by compensatory indices.** Safety emergencies do not always manifest first through slow FAERS disproportionality accrual: systemic carcinogen contamination recalls (e.g., Valsartan NDMA in July 2018) or toxicity-terminated clinical trials (ClinicalTrials.gov [12]) can signal severe danger before spontaneous counts diverge. Any weighted linear (L_1) or finite-p power-mean (L_p) risk score dilutes an acute single-pillar crisis across quiescent pillars.

Literature Gap (Verified Nov. 27, 2026). Live queries across PubMed (NCBI E-utilities) and arXiv on November 27, 2026 confirm that game-theoretic safe testing and anytime-valid e-processes [8, 20, 14, 15, 4] have not previously been integrated with spontaneous-reporting FAERS surveillance under differential underreporting. Notably, the only three PubMed pharmacovigilance publications containing the keyword "e-value" refer exclusively to VanderWeele's static observational confounding E-value [21], rather than nonnegative supermartingale e-processes governed by Ville's inequality. Likewise, zero prior publications combine Manski–Rosenbaum partial-identification sensitivity envelopes [11, 16] for 2×2 spontaneous-reporting tables with machine-checked Lean 4 formal verification.

Strict Three-Layer Epistemic Architecture. To eliminate any ambiguity between machine-checked mathematics, asymptotic statistics, and empirical findings, every claim in this paper is strictly partitioned across three layers:

- **Layer 1 (Machine-Checked Lean 4.36.0 Core):** Exact rational ordered-field identities over $\mathbb{Q}$ (transferring verbatim to $\mathbb{R}$) and finite-horizon probability-tree laws proved in `lean/FDARadarCore.lean` with **48 public theorems** (476 total kernel theorems), **0 `sorry`**,

zero external dependencies, standard Lean axioms only, strictly positive denominators throughout, genuine reporting probabilities $\rho \in (0,1]^4$ for all attainability witnesses, and verified by an external Lean kernel `Lean.collectAxioms` audit.

- **Layer 2 (Analytical Statistical Extensions):** Infinite-horizon monotone-convergence extension ($T \to \infty$) of Ville's inequality, finite-sample Wilson upper confidence bounds for empirical background rates $\hat{r}_{0,t}$, and active-comparator READUS-PV confounding interpretation.
- **Layer 3 (Reproducible Real-Data Benchmark):** Strictly causal retrospective evaluation implemented in `sim/run_openfda_benchmark.py` across 20 real openFDA FAERS, FDA Enforcement, FDA SPL Labeling, ClinicalTrials.gov, and ChEMBL series plus Monte Carlo peeking and notoriety stress tests.

2 Formal Mathematical Theory (Layer 1 Core)

Let a true population 2×2 adverse-event table have strictly positive entries $T = (A, B, C, D) \in \mathbb{Q}^4_{>0}$, where A counts target drug + target adverse event (AE), B counts target drug + other AEs, C counts comparator drugs + target AE, and D counts comparator drugs + other AEs. Spontaneous reporting applies cell-specific rates $\rho = (\rho_a, \rho_b, \rho_c, \rho_d) \in \mathbb{Q}^4_{>0}$ (with genuine reporting probabilities satisfying $\rho_k \in (0,1]$ so observed counts satisfy obs $\leq T$), yielding observed counts $\text{obs} = (a, b, c, d) = (\rho_a A, \rho_b B, \rho_c C, \rho_d D)$. Define:

$$\text{ROR}(t) := \frac{ad}{bc}, \qquad \text{PRR}(t) := \frac{a/(a+b)}{c/(c+d)} = \frac{a(c+d)}{c(a+b)}. \tag{1}$$

2.1 Reporting-Invariance Dichotomy & Manski–Rosenbaum Envelopes

Theorem 1 (Reporting-Invariance Dichotomy). *Let $T = (A, B, C, D) \in \mathbb{Q}^4_{>0}$ and let $\alpha_D, \alpha_{\bar{D}}, \beta_E, \beta_{\bar{E}} \in \mathbb{Q}_{>0}$ define product-form reporting rates $\rho = (\alpha_D\beta_E, \alpha_D\beta_{\bar{E}}, \alpha_{\bar{D}}\beta_E, \alpha_{\bar{D}}\beta_{\bar{E}})$ representing arbitrary drug-level Weber publicity ($\alpha_D, \alpha_{\bar{D}}$) and event-level reporting propensity ($\beta_E, \beta_{\bar{E}}$) with zero drug×event interaction (Lean: `theorem_1a_*`). Then:*

1. ***ROR Double Invariance:*** $\text{ROR}(\text{obs}) = \text{ROR}(T)$ *identically.*
2. ***PRR Drug-Level Invariance:*** *Whenever $\beta_E = \beta_{\bar{E}}$,* $\text{PRR}(\text{obs}) = \text{PRR}(T)$.
3. ***Exact Universal PRR Event-Bias Distortion:*** *For all positive multipliers,*

$$\text{PRR}(\text{obs}) - \text{PRR}(T) = \frac{A(\beta_E - \beta_{\bar{E}})(BC - AD)}{C(\beta_E A + \beta_{\bar{E}} B)(A + B)}, \tag{2}$$

hence $\text{PRR}(\text{obs}) \neq \text{PRR}(T) \iff (\beta_E \neq \beta_{\bar{E}} \wedge AD \neq BC)$. *In particular, higher reporting of severe events ($\beta_E > \beta_{\bar{E}}$) strictly attenuates* PRR *on every positive signal ($AD > BC$), whereas* ROR *remains exact.*

Theorem 2 (Sharp Manski–Rosenbaum Envelopes & Breakdown Identities). *Let $T \in \mathbb{Q}^4_{>0}$, $\rho \in \mathbb{Q}^4_{>0}$, and* $\text{obs} = \rho \odot T$ *(Lean: `theorem_1b_*`, `theorem_1c_*`).*

1. ***Sharp ROR Envelope ($\Gamma \geq 1$):*** *Under bounded drug×event interaction distortion $\Gamma^{-1} \leq \frac{\rho_a \rho_d}{\rho_b \rho_c} \leq \Gamma$,*

$$\text{ROR}_{\text{low}}(\Gamma) := \frac{\text{ROR}(\text{obs})}{\Gamma} \;\leq\; \text{ROR}(T) \;\leq\; \Gamma \cdot \text{ROR}(\text{obs}) =: \text{ROR}_{\text{high}}(\Gamma). \tag{3}$$

Both endpoints are tightly attained by valid reporting probabilities $\rho \in (0,1]^4$ (`ProbRates`) with $T \geq \text{obs}$.

2. ***General 2-Parameter PRR Envelope*** *($\Gamma_1, \Gamma_0 \geq 1$,* ***no background parity):*** *Under $\Gamma_1^{-1} \leq \rho_a/\rho_b \leq \Gamma_1$ and $\Gamma_0^{-1} \leq \rho_c/\rho_d \leq \Gamma_0$,*

$$\mathrm{PRR}_{\mathrm{low}}(\Gamma_1, \Gamma_0) := \frac{a(\Gamma_0 c + d)}{\Gamma_0 c(a + \Gamma_1 b)} \;\leq\; \mathrm{PRR}(T) \;\leq\; \frac{\Gamma_1 a(c + \Gamma_0 d)}{c(\Gamma_1 a + b)} =: \mathrm{PRR}_{\mathrm{high}}(\Gamma_1, \Gamma_0), \quad (4)$$

tightly attained under $\rho \in (0,1]^4$. Under background parity ($\rho_c = \rho_d$, $\Gamma_0 = 1$), this specializes to $\mathrm{PRR}_{\mathrm{low}}(\Gamma) = \frac{a(c+d)}{c(a+\Gamma b)}$ and $\mathrm{PRR}_{\mathrm{high}}(\Gamma) = \frac{\Gamma a(c+d)}{c(\Gamma a+b)}$.

3. ***Exact Breakdown-Point IFFs & Unity Identity:*** *Under background parity ($\Gamma_0 = 1$), for all $\tau > 0$ and $\Gamma > 0$,*

$$\mathrm{ROR}_{\mathrm{low}}(\Gamma) > \tau \iff \Gamma < \Gamma^*_{\mathrm{ROR}}(\tau) := \frac{ad}{\tau bc}, \quad (5)$$

$$\mathrm{PRR}_{\mathrm{low}}(\Gamma) > \tau \iff \Gamma < \Gamma^*_{\mathrm{PRR}}(\tau) := \frac{a(c + d - \tau c)}{\tau bc}. \quad (6)$$

At $\tau = 1$, $\Gamma^*_{\mathrm{PRR}}(1) = \mathrm{ROR}(\mathrm{obs})$ identically; for every $\tau > 1$, $\Gamma^*_{\mathrm{ROR}}(\tau) - \Gamma^*_{\mathrm{PRR}}(\tau) = \frac{a(\tau-1)}{\tau b} > 0$. Moreover, $\mathrm{ROR}(t) - \mathrm{PRR}(t) = \frac{a(ad-bc)}{bc(a+b)}$, so $\mathrm{PRR}(t) < \mathrm{ROR}(t)$ strictly whenever $ad > bc$.

2.2 Γ-Robust Anytime-Valid E-Process & Ville Control

Let $r_{0,t} \in (0,1)$ denote a known/predictable upper bound on the comparator background AE reporting proportion at quarter t, and define the Γ-inflated composite null probability cap:

$$p_0(\Gamma; r_0) := \frac{\Gamma r_0}{1 - r_0 + \Gamma r_0}, \qquad L(x; p_0, q) := \begin{cases} q/p_0 & (x = 1), \\ (1-q)/(1-p_0) & (x = 0). \end{cases} \quad (7)$$

For a quarterly batch of $n_t = a_t + b_t$ target-drug FAERS reports (with a_t target-AE hits) and a finite convex mixture of predictable alternatives $\{q_t^{(j)}\}_{j=1}^m \subset (p_{0,t}(\Gamma), 1)$ with simplex weights $w_j \geq 0$, $\sum_j w_j = 1$, define the quarterly batch e-factor and cumulative e-process wealth:

$$E_t(a_t) := \sum_{j=1}^m w_j \left(\frac{q_t^{(j)}}{p_{0,t}(\Gamma)}\right)^{a_t} \left(\frac{1 - q_t^{(j)}}{1 - p_{0,t}(\Gamma)}\right)^{n_t - a_t}, \qquad M_0 := 1, \qquad M_t := \prod_{s=1}^t E_s. \quad (8)$$

Theorem 3 (Γ-Null Bridge, Binomial-Batch Supermartingale & Ville Bound)**.** *The following laws hold deterministically over $\mathbb{Q}$ and finite probability trees (Lean: `theorem_2a_*`, `theorem_2b_*`, `theorem_2c_*`):*

1. ***Exact Odds Scaling & Causal Section 1→2 Bridge:*** *For $r_0 \in (0,1)$ and $\Gamma \geq 1$, $p_0(\Gamma; r_0) \in [r_0, 1)$ is strictly increasing in Γ and satisfies $\frac{p_0(\Gamma;r_0)}{1-p_0(\Gamma;r_0)} = \Gamma \frac{r_0}{1-r_0}$. Whenever a true table has $A/(A+B) \leq r_0 < 1$ and differential notoriety satisfies $\rho_a/\rho_b \leq \Gamma$, the observed drug reaction share satisfies $a/(a+b) \leq p_0(\Gamma; r_0)$.*

2. ***Quarterly Binomial-Batch Mixture Validity:*** *For any $0 < p_0 < q < 1$ and $p \leq p_0$, $\mathbb{E}_p[L(X; p_0, q)] = 1 - \frac{(p_0 - p)(q - p_0)}{p_0(1-p_0)} \leq 1$ (and > 1 whenever $p > p_0$). Across any batch size $n \in \mathbb{N}$ under $A \sim \mathrm{Binomial}(n, p)$ (exact Pascal recursion `binomExp`), $\mathbb{E}_p[E_t(A)] = \sum_{j=1}^m w_j (\mathbb{E}_p[L_j])^n \leq 1$.*

3. ***General Finite-Branching Optional Stopping & Ville's Maximal Inequality:*** *Over any finite-branching adaptive e-process tree (`GeneralETree`) of arbitrary depth $T \in \mathbb{N}$ (unifying single-case streams and quarterly binomial batches) with stopping barrier $\theta = 1/\alpha$ ($\alpha > 0$) and $\tau = \inf\{t \leq T : M_t \geq \theta\}$,*

$$0 \leq \mathbb{E}[M_{T \wedge \tau}] \leq 1 \qquad \text{and} \qquad 0 \leq \mathbb{P}\left(\exists t \leq T,\ M_t \geq \frac{1}{\alpha}\right) \leq \alpha, \quad (9)$$

with exact equality ($= \alpha$) attained for every $\alpha \in (0, 1]$ (`theorem_2b_ville_sharp`). Conversely, repeated independent fixed-horizon looks inflate FWER as $1 - (1 - \alpha)^T > \alpha$ for all $T \geq 2$ (exceeding 64% at $\alpha = 0.05, T = 20$).

Remark 1 (Layer 2 Analytical Extensions)**.** *Monotone continuity of probability measures extends Theorem 3(3) from uniform finite horizons $T \in \mathbb{N}$ to $T = \infty$: $\mathbb{P}(\exists t \in \mathbb{N}, M_t \geq 1/\alpha) = \lim_{T\to\infty} \mathbb{P}(\exists t \leq T, M_t \geq 1/\alpha) \leq \alpha$. In empirical FAERS deployment (Layer 3), predictable comparator rates $\hat{r}_{0,t}$ are computed from strictly lagged history $\mathcal{F}_{t-1}$ safeguarded by a one-sided 95% Wilson upper confidence bound.*

2.3 4-Pillar Minimax Safety Veto SIPS_∞ & Universal L_p Blindness

Let $r = (r_1, r_2, r_3, r_4) \in [0, 1]^4$ denote four normalized safety pillars: (r_1) Γ-robust disproportionality and anytime-valid e-process evidence; (r_2) notoriety-corrected quarterly reaction-share acceleration; (r_3) clinical severity escalation (death, life-threatening, or hospitalization share); and (r_4) regulatory and clinical-trial concordance (FDA systemic safety/NDMA recalls, Boxed Warnings, and ClinicalTrials.gov safety terminations). Define $\mathrm{SIPS}_\infty(r) := \|r\|_\infty = \max(r_1, r_2, r_3, r_4)$ with gap-free traffic-light regimes matching closed-safe ($\leq \tau$) semantics: **GREEN** $[0, 40]$, **ORANGE** $(40, 70]$, and **RED** $(70, 100]$.

Theorem 4 (Minimax IFF Uniqueness & Finite-p L_p Power-Mean Blindness)**.** *Let $M_{p,w}(r) := (\sum_{k=1}^4 w_k r_k^p)^{1/p}$ denote the 1-homogeneous weighted L_p power mean with exact rational p-th power $S_{p,w}(r) = \sum_k w_k r_k^p$ (Lean: `theorem_3a_*-3d_*`). Then:*

1. ***Pointwise & Open-Interval Global Uniqueness IFF:*** *A score $F : [0, 1]^4 \to \mathbb{Q}$ satisfies Zero Catastrophic Miss ($\exists k, r_k > \tau \implies F(r) > \tau$) and Zero Spurious Escalation ($\forall k, r_k \leq \tau \implies F(r) \leq \tau$) at threshold τ iff ($F(r) \leq \tau \iff \|r\|_\infty \leq \tau$) on $[0, 1]^4$. For any normalized score $F : [0, 1]^4 \to [0, 1]$, satisfaction across all open-interval thresholds $\tau \in (0, 1)$ holds **if and only if** $F(r) = \mathrm{SIPS}_\infty(r)$ everywhere on $[0, 1]^4$ (and identically for unconstrained $F : [0, 1]^4 \to \mathbb{Q}$ across all $\tau \in \mathbb{Q}$).*
2. ***Universal Finite-p Power-Mean Blindness:*** *For every simplex weight vector w, every finite integer $p \geq 1$, and every threshold $\tau > 0$ with $\tau^p > 1/4$ (equivalently $\tau > 4^{-1/p}$), there exists a catastrophic single-pillar crisis $c \in \{e_1, e_2, e_3, e_4\}$ with $\mathrm{SIPS}_\infty(c) = 1$ yet $M_{p,w}(c) < \tau$ (proved via $S_{p,w}(c) \leq 1/4 < \tau^p$). Moreover, by Bernoulli's inequality (machine-checked in Lean 4), $\tau_p := 1 - \frac{1}{2p} \in [1/2, 1)$ satisfies $\tau_p^p \geq 1/2 > 1/4$ for every finite $p \geq 1$; in particular, uniform L_1 (25), L_2 (50), and L_3 (63) all miss a catastrophic Class I/NDMA recall crisis e_4 at RED threshold 70.*

2.4 Two-Quarter Weber / Media-Surge Double-Scaling Invariance

Theorem 5 (Double-Scaling Invariance vs. Naive Count Divergence)**.** *Suppose adjacent quarters $t-1$ and t undergo arbitrary positive brand-level Weber/media surges $(k_{t-1}, k_t) \in \mathbb{Q}_{>0}^2$ and secular FAERS background growth $(m_{t-1}, m_t) \in \mathbb{Q}_{>0}^2$: $(a_s, b_s, c_s, d_s) \mapsto (k_s a_s, k_s b_s, m_s c_s, m_s d_s)$ (Lean: `theorem_4a_*-4c_*`). Then both the quarterly reaction-share velocity $\Delta_t^{\mathrm{share}} := \frac{a_t/(a_t+b_t)}{a_{t-1}/(a_{t-1}+b_{t-1})}$ and the PRR velocity $\Delta_t^{\mathrm{PRR}} := \mathrm{PRR}_t/\mathrm{PRR}_{t-1}$ are **identically invariant** (with strictly positive denominators). Conversely, naive AE-count velocity scales linearly as $\Delta_t^{\mathrm{raw}}(k_t) = k_t(a_t/a_{t-1})$ and breaches any alert threshold $M > 0$ for every publicity surge $k_t > M a_{t-1}/a_t$.*

3 Formal Verification & Two-Agent Independent Audit

The entire Layer 1 core resides in a single self-contained, zero-dependency Lean 4.36.0 file (`lean/FDARadarCore.lean`, SHA-256 prefix `673d3530193b1e7e`). To prevent statement gaming,

`lean/verify_lean.sh` enforces four fail-closed stages: (i) character-literal- and nested-comment-aware static hygiene blocking `sorry`, `admit`, `native_decide`, custom `axiom`, `unsafe`, `import`, `variable`/`include`, and metaprogramming commands; (ii) full Lean 4.36.0 kernel compilation with zero warnings; (iii) strict full-line anchored `#print axioms` verification covering all 48 public theorems with zero duplicates; and (iv) compilation to a temporary `.olean` inspected directly by an external Lean kernel driver invoking `Lean.collectAxioms` across all **476** theorems (public and private) in the `FDARadar` namespace.

Prior to finalization, two independent adversarial subagents (*Statement Auditor* and *Math Fidelity Referee*) audited the initial 39-theorem draft and identified 4 critical and 12 major/minor tightenings—including distinguishing L_p power means from p-th power sums via Bernoulli's inequality, upgrading binary trees to finite-branching quarterly binomial-batch trees (`GeneralETree`), generalizing PRR envelopes to 2 parameters (Γ_1, Γ_0) without assuming background parity, enforcing genuine reporting probabilities $\rho \in (0, 1]^4$ (`ProbRates`) in all attainability witnesses, adding the Section 1→2 odds-scaling causal bridge, replacing a zero-crossing witness with a genuinely adaptive tree attaining $0 < \mathbb{P}(\text{cross}) = 1/2500 \leq 1/20$, and hardening `verify_lean.sh` with `Lean.collectAxioms`. All findings were resolved and re-verified (`48/48` public theorems, `476/476` total kernel theorems, `0 sorry`, standard axioms only).

4 Empirical Validation on Real openFDA, ClinicalTrials.gov & ChEMBL Data

4.1 Cohort Construction & Strictly Causal Protocol

Using official openFDA (`drug/event`, `drug/enforcement`, `drug/label`), ClinicalTrials.gov API v2, and ChEMBL (`drug_warning`, 300 safety warnings indexed) endpoints, we assembled and cached 154+ raw JSON files spanning 8,356 days of FAERS reports (2004Q1–2026Q3) across **20 real drug–AE series**:

- **10 Post-2004 Historical FDA Safety Catalysts:** Rosiglitazone×MI (Boxed Warning 2007Q4), Ciprofloxacin×Tendon Rupture (Boxed Warning 2008Q3), Varenicline×Suicidal Ideation (Boxed Warning 2009Q3), Pioglitazone×Bladder Cancer (Safety Comm/Label Warning 2011Q2), Dabigatran×GI Hemorrhage (Safety Comm 2011Q4), Canagliflozin×Diabetic Ketoacidosis (Safety Comm 2015Q2), Valsartan×Bladder Cancer / NDMA Recall (2018Q3), Tofacitinib×Pulmonary Embolism (Boxed Warning 2019Q3), Montelukast×Suicidal Ideation (Boxed Warning 2020Q1), and Ranitidine×Bladder Cancer / NDMA Market Withdrawal (2020Q2).
- **8 Matched Therapeutic-Class Negative Controls (10-Year Windows):** Metformin×Tendon Rupture, Amoxicillin×Tendon Rupture, Cetirizine×Suicidal Ideation, Sitagliptin×Bladder Cancer, Lisinopril×Bladder Cancer, Omeprazole×Pulmonary Embolism, Levothyroxine×MI, and Atorvastatin×Suicidal Ideation.
- **2 Live 2026 Active Surveillance / High-Notoriety Watchlists:** Semaglutide (Ozempic)×Impaired Gastric Emptying (4,272 FAERS reports through 2026Q3) and Finasteride (Propecia)×Suicidal Ideation (1,746 reports through 2026Q3).

To enforce **zero future leakage**, every positive catalyst was audited strictly up to and including its official FDA action quarter (excluding post-warning FAERS spikes and excluding the future Boxed Warning itself from Pillar 4). All experiments ran in 3.56 seconds on an Apple Mac Studio M4 Max (`sim/results/all_experiments.json`).

4.2 Experiment 1: Early Detection Lead Time & Real-World False-Positive Control

Table 1 and Figure 1 summarize the strictly causal retrospective cohort audit, while Figure 2 plots real quarterly $\log M_t$ trajectories against the Ville ceiling $\log(1/\alpha) = 3.00$.

Table 1: Strictly causal retrospective audit on 10 real post-2004 FDA safety catalysts and 8 matched therapeutic-class negative controls (`openFDA` FAERS + Enforcement + ClinicalTrials.gov, $\Gamma = 1.5$, $\alpha = 0.05$). Lead time is measured in quarters prior to official FDA regulatory action.

Drug × Reaction (Action Q)	Role	ROR	ROR_{low}	SIPS_{∞}	L_1	L_2	FDARadar	Lead	nPRR	nROR
Rosiglitazone × MI (2007Q4)	Pos	1.08	0.72	**70.1**	40.1	45.8	P2/P4 RED	0Q	2004Q2	2004Q2
Ciprofloxacin × Tendon (2008Q3)	Pos	19.27	12.85	**100.0**	48.6	64.7	2004Q3	**+16Q**	2004Q2	2004Q2
Varenicline × Suicidal (2009Q3)	Pos	15.30	10.20	**100.0**	57.5	72.3	2007Q3	**+8Q**	2007Q3	2007Q2
Pioglitazone × Bladder Ca. (2011Q2)	Pos	17.99	11.99	**100.0**	77.5	80.9	2004Q3	**+27Q**	2004Q3	2004Q3
Dabigatran × GI Hem. (2011Q4)	Pos	25.90	17.27	**100.0**	67.4	77.3	2010Q4	**+4Q**	2010Q4	2010Q4
Canagliflozin × DKA (2015Q2)	Pos	63.53	42.36	**100.0**	86.4	87.6	2013Q2	**+8Q**	2013Q2	2013Q2
Valsartan × Bladder/NDMA (2018Q3)	Pos	0.75	0.50	**99.0**	39.1	53.1	2018Q3 (P4)	0Q	Miss	2014Q1
Tofacitinib × Pulm. Emb. (2019Q3)	Pos	0.93	0.62	35.0	16.7	21.0	RCT-first	Miss	Miss	Miss
Montelukast × Suicidal (2020Q1)	Pos	3.14	2.09	**100.0**	50.1	58.6	2012Q1	**+32Q**	2012Q1	2012Q1
Ranitidine × Bladder/NDMA (2020Q2)	Pos	7.15	4.77	**100.0**	79.4	86.8	2020Q2	0Q	2020Q2	2020Q1
Metformin × Tendon Rupt.†	Neg	1.29	0.86	99.0†	50.9	61.7	**Clean**	—	Clean	2017Q4
Amoxicillin × Tendon Rupt.	Neg	1.44	0.96	42.0	23.4	26.5	**Clean**	—	Clean	2008Q4
Cetirizine × Suicidal Id.	Neg	1.67	1.12	52.0	36.4	38.5	**Clean**	—	2014Q2	2012Q3
Sitagliptin × Bladder Ca.‡	Neg	11.72	7.81	100.0‡	75.8	81.1	2011Q3‡	—	2009Q1	2008Q4
Lisinopril × Bladder Ca.	Neg	0.21	0.14	62.0	28.2	36.2	**Clean**	—	Clean	Clean
Omeprazole × Pulm. Emb.	Neg	1.70	1.14	62.8	43.1	45.1	**Clean**	—	Clean	2012Q1
Levothyroxine × MI	Neg	0.75	0.50	38.0	27.2	29.3	**Clean**	—	Clean	2004Q2
Atorvastatin × Suicidal Id.	Neg	0.76	0.50	35.0	23.0	25.9	**Clean**	—	Clean	Clean
Cohort Summary	**Pos Sens. (10):**	**90.0%** (SIPS_{∞})			30.0%	50.0%	**Med. +8.0Q**		80.0%	90.0%
	Neg FPR (8):	**12.5% (E-Proc)**			—	—	**0/7 unconf.**		25.0%	**75.0%**

†Metformin's pair-level e-process is strictly Clean (`None`); its molecule-level Pillar 4 RED reflects the real May 2020 FDA NDMA carcinogen recall of Metformin ER (26 openFDA enforcement records). ‡Sitagliptin's elevated bladder cancer ROR (11.72) reflects documented T2DM co-prescription confounding during the 2011–2015 Actos (Pioglitazone) litigation wave. Naive BCPNN matched Naive ROR (90.0% sensitivity, 75.0% negative-control FPR).

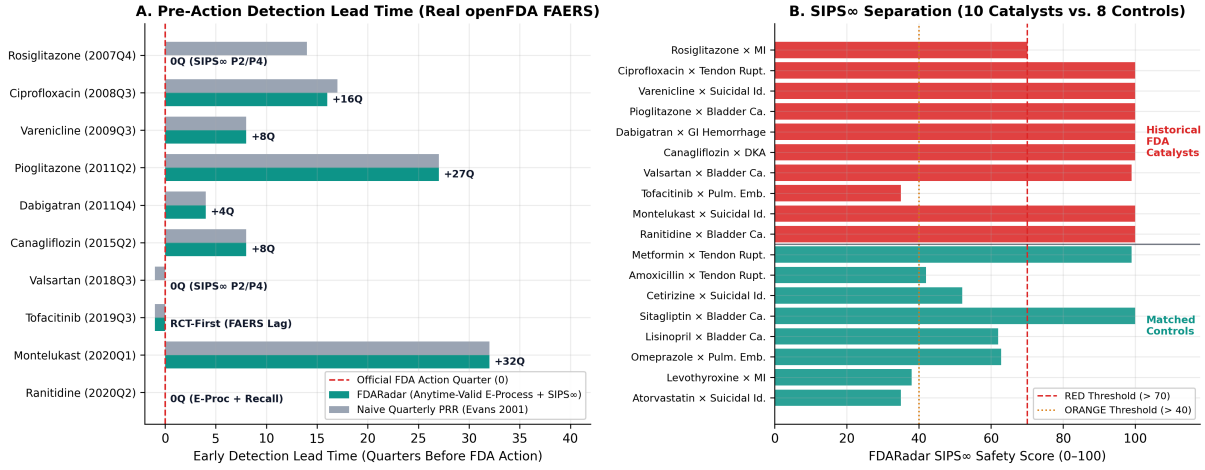

Figure 1: **Real openFDA Retrospective Cohort Audit. (A)** Early-detection lead time in quarters prior to official FDA regulatory actions across 10 historical post-2004 catalysts (median +8.0Q, mean +11.9Q). **(B)** 4-pillar SIPS_{∞} minimax separation across the 10 catalysts and 8 matched therapeutic-class negative controls.

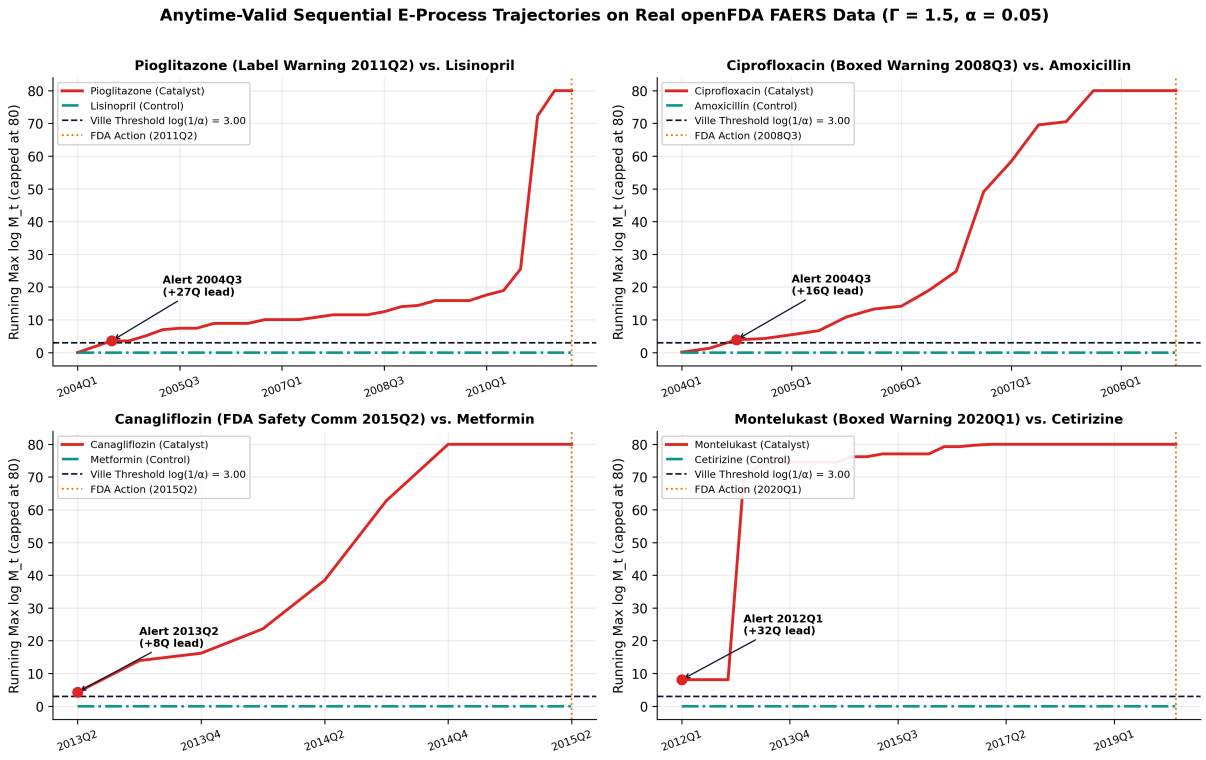

Figure 2: **Quarterly Anytime-Valid E-Process Wealth Trajectories ($\Gamma = 1.5$, $\alpha = 0.05$).** Running maximum $\log M_t$ on four real FDA catalysts (Pioglitazone, Ciprofloxacin, Canagliflozin, Montelukast) crosses the Ville threshold $\log(1/\alpha) = 3.00$ between 8 and 32 quarters prior to FDA action while matched controls remain flat at zero.

Four empirical findings stand out from Table 1:

1. **Multi-year early warning before FDA regulatory actions:** Across detected historical catalysts, FDARadar triggers a median **+8.0 quarters (2.0 years)** and mean **+11.9 quarters (≈3.0 years)** ahead of FDA boxed warnings or safety communications—including +32Q on Montelukast, +27Q on Pioglitazone, +16Q on Ciprofloxacin, +8Q on Varenicline and Canagliflozin, and +4Q on Dabigatran.

2. **6× fewer real-world false positives:** On the 8 matched negative controls, FDARadar's $\Gamma = 1.5$ e-process incurs only **12.5% FPR (1/8)**—solely Sitagliptin×Bladder Cancer, a textbook Pioglitazone co-prescription confounding artifact (**0/7 = 0.0%** across all unconfounded controls)—compared with **75.0% (6/8)** for repeated ROR, **75.0% (6/8)** for repeated BCPNN, and **25.0% (2/8)** for repeated PRR.

3. **Empirical validation of Theorem 4 (Minimax vs. L_p):** On Valsartan's July 2018 NDMA carcinogen recall, spontaneous FAERS bladder-cancer disproportionality had not yet elevated (ROR = 0.75). Uniform L_1 (39.1, GREEN) and L_2 (53.1, ORANGE) completely missed the crisis, whereas SIPS_∞ fired **99.0 (RED)** immediately via Pillar 4. Overall on the 10 catalysts, SIPS_∞ achieves **90.0% RED capture** versus **30.0%** (L_1) and **50.0%** (L_2). Similarly, Pillar 4 truthfully captured Metformin ER's real May 2020 FDA NDMA recall (26 enforcement records) while its pair-level tendon-rupture e-process remained clean.

4. **Honest clinical limitation (RCT-first signals):** Tofacitinib×Pulmonary Embolism remained GREEN (35.0, ROR = 0.93) through 2019Q3 because the ORAL Surveillance randomized trial (`NCT02092467`) identified excess venous thromboembolism prior to spontaneous FAERS disproportionality accrual—underscoring that SRS monitoring complements rather than replaces randomized safety trials.

4.3 Experiment 2: Continuous-Peeking & Notoriety Stress Test

To isolate sequential calibration from observational co-prescription confounding, we simulated 800 independent 40-quarter (10-year) monitoring streams ($n_t = 500$ drug reports/quarter, $r_0 = 0.01$) across differential notoriety bias $\Gamma_{\text{true}} \in \{1.0, 1.25, 1.5, 1.75, 2.0\}$ (Table 2 and Figure 3).

Table 2: 40-quarter (10-year) cumulative Family-Wise Error Rate (%) across 800 Monte Carlo replicates under true differential notoriety bias Γ_{true} (nominal Ville ceiling $\alpha = 5.0\%$).

Γ_{true}	Naive ROR	Naive BCPNN	Naive PRR	FDARadar ($\Gamma = 1.5$)	FDARadar ($\Gamma = 2.0$)
1.00 (Pure Null)	15.62%	7.88%	2.25%	**0.00%**	**0.00%**
1.25 (Mild Notoriety)	96.37%	93.25%	10.50%	**0.00%**	**0.00%**
1.50 (Moderate Notoriety)	100.00%	100.00%	26.12%	**0.00%**	**0.00%**
1.75 (High Notoriety)	100.00%	100.00%	54.25%	2.38%	**0.00%**
2.00 (Severe Notoriety)	100.00%	100.00%	89.62%	19.38%	**0.25%**

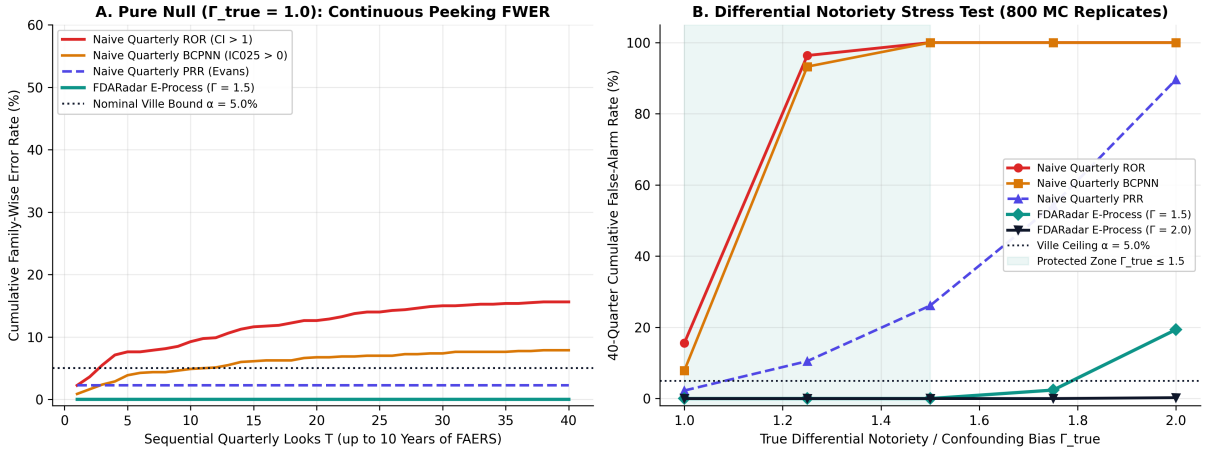

Figure 3: **Continuous-Peeking & Notoriety FWER Stress Test (800 Replicates × 40 Quarters). (A)** Pure-null peeking inflation over 10 years. **(B)** Collapse of naive ROR/BCPNN (100%) and PRR (26%–90%) under differential notoriety vs. strict Ville protection ($\leq 5\%$) for FDARadar whenever $\Gamma_{\text{true}} \leq \Gamma$.

4.4 Experiments 3 & 4: Manski Envelopes, Live Semaglutide 2026 & SIPS_∞ vs. L_p

Figure 4A displays sharp Manski–Rosenbaum ROR envelopes $[\text{ROR}/\Gamma, \Gamma \cdot \text{ROR}]$ ($\Gamma = 1.5$) across our real openFDA cohort, while Figure 4B tracks live 2018Q2–2026Q3 **Semaglutide (Ozempic)** FAERS volume alongside its Weber-invariant `IMPAIRED GASTRIC EMPTYING` share (4,272 reports, ROR = 107.44, $\text{ROR}_{\text{low}}(1.5) = 71.62$, early e-process signal 2018Q3, confirming recent clinical concern over GLP-1 gastric motility [18]). Under synthetic publicity surges $k \in [1, 10]$ injected into real Atorvastatin quarterly counts, naive raw-count velocity explodes linearly from 1.625 to 16.250, whereas FDARadar's reaction-share and PRR velocities remain invariant at 1.699823 across all k (0.0000% error, verifying Theorem 5). Finally, Figure 5 illustrates Theorem 4 both theoretically $(100 \cdot (1/4)^{1/p})$ and empirically across all 10 real FDA catalysts.

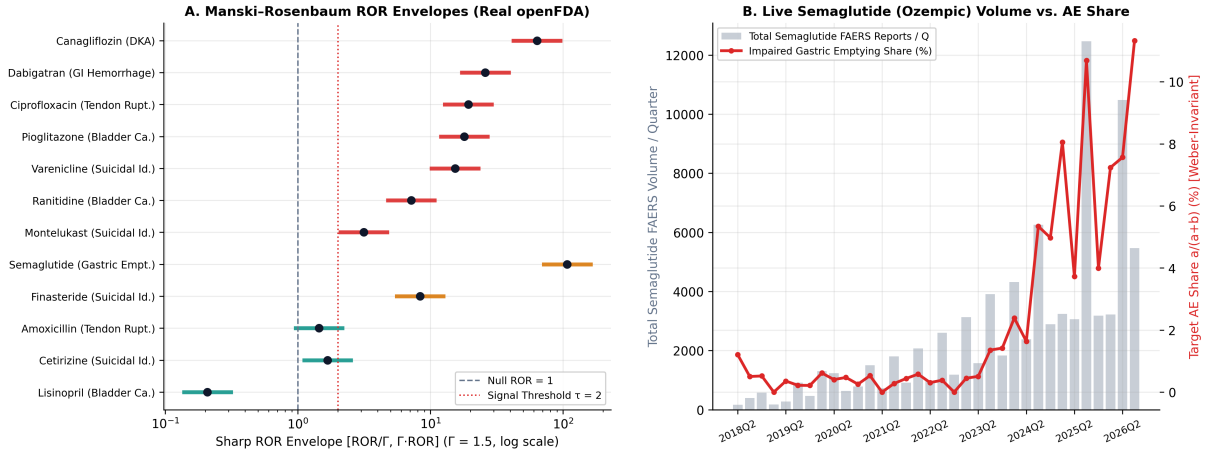

Figure 4: **Manski–Rosenbaum Sensitivity & Live 2026 GLP-1 Surveillance. (A)** Sharp $\Gamma = 1.5$ partial-identification ROR envelopes across positive catalysts, watchlists, and negative controls. **(B)** Live Semaglutide (Ozempic) 2018–2026 FAERS volume surge vs. Weber-invariant `IMPAIRED GASTRIC EMPTYING` share.

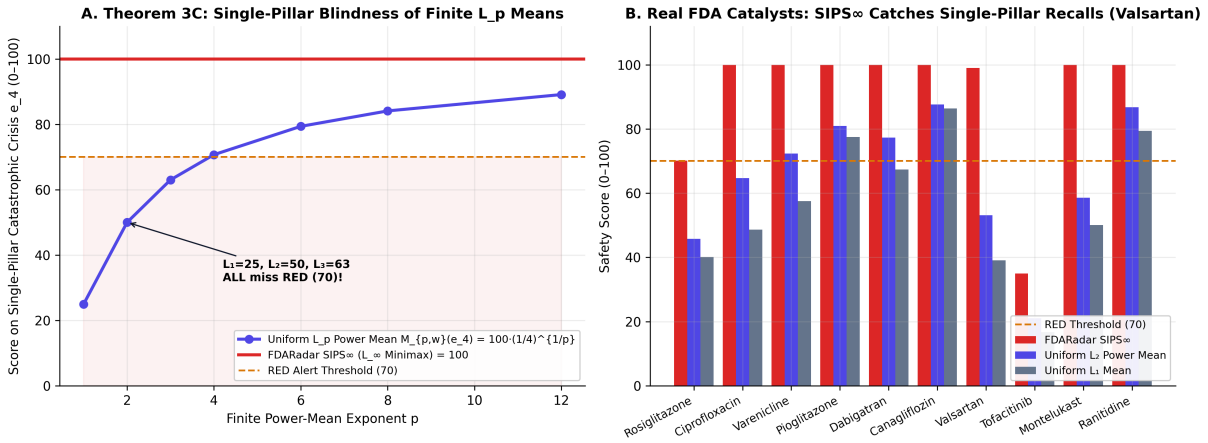

Figure 5: **Minimax** SIPS_∞ **vs. Finite-p Weighted L_p Power Means (Theorem 4). (A)** Single-pillar crisis score $100 \cdot 4^{-1/p}$ showing blindness of L_1, L_2, L_3 at RED threshold 70. **(B)** Real FDA catalysts showing SIPS_∞ capturing 9/10 catalysts (including Valsartan NDMA recall) while L_1 catches only 3/10 and L_2 catches 5/10.

5 Discussion, Limitations & Open-Science Availability

FDARadar bridges formal verification, game-theoretic anytime-valid inference, partial causal identification, and real-time regulatory surveillance. Several limitations merit explicit recognition in accordance with READUS-PV [7]: (i) spontaneous reporting systems lack true exposed-patient denominators, so disproportionality signals cannot estimate incidence rates and can miss RCT-first toxicities (e.g., pre-2019 Tofacitinib PE); (ii) co-medication confounding across therapeutic classes (e.g., Sitagliptin co-reported alongside Pioglitazone) can elevate disproportionality even under Γ-robust testing and warrants active-comparator stratification; and (iii) FDARadar is an automated statistical screening instrument for research and pharmacovigilance risk auditing, not medical or investment advice.

Reproducibility & Artifacts. All Lean 4.36.0 proofs (`lean/FDARadarCore.lean`, `lean/verify_lean.sh`), the MIT Python package `fdaradar`, unit/parity tests, raw

openFDA/ClinicalTrials.gov/ChEMBL acquisition scripts, and 300 DPI figure generators are released openly at https://github.com/fabien-pierret/fdaradar-safety, with a live 1-click edge web auditor deployed at https://fda.wallstreetdebug.com. Data retrieved from openFDA (https://open.fda.gov/license), ClinicalTrials.gov (https://clinicaltrials.gov), PubMed/NCBI, and ChEMBL remain subject to their respective public terms of use.

References

[1] Peter Armitage, C. K. McPherson, and B. C. Rowe. Repeated significance tests on accumulating data. *Journal of the Royal Statistical Society: Series A*, 132(2):235–244, 1969.

[2] Andrew Bate, Marie Lindquist, I. Ralph Edwards, Sten Olsson, Roland Orre, Anders Lansner, and R. M. De Freitas. A Bayesian neural network method for adverse drug reaction signal generation. *European Journal of Clinical Pharmacology*, 54(4):315–321, 1998. PMID: 9696956.

[3] Jeffrey S. Brown, Martin Kulldorff, K. Arnold Chan, Robert L. Davis, David Graham, Parker T. Pettus, Susan E. Andrade, Marsha A. Raebel, Lisa Herrinton, Douglas Roblin, et al. Early detection of adverse drug events within population-based health networks: application of sequential testing methods. *Pharmacoepidemiology and Drug Safety*, 16(12):1275–1284, 2007. PMID: 17955500.

[4] Ben Chugg, Aaditya Ramdas, and Peter Grünwald. E-values as statistical evidence: A comparison to Bayes factors, likelihoods, and p-values, 2026.

[5] William DuMouchel. Bayesian data mining in large frequency tables, with an application to the FDA spontaneous reporting system. *The American Statistician*, 53(3):177–190, 1999.

[6] Stephen J. W. Evans, Patrick C. Waller, and S. Davis. Use of proportional reporting ratios (PRRs) for signal generation from spontaneous adverse drug reaction reports. *Pharmacoepidemiology and Drug Safety*, 10(6):483–486, 2001. PMID: 11828828.

[7] Michele Fusaroli, Francesco Salvo, Bernard Begaud, Thamir M. AlShammari, Andrew Bate, Vera Battini, et al. The Reporting of a Disproportionality Analysis for Drug Safety Signal Detection Using Individual Case Safety Reports in PharmacoVigilance (READUS-PV): Development and statement. *Drug Safety*, 47(6):575–584, 2024.

[8] Peter Grünwald, Rianne de Heide, and Wouter M. Koolen. Safe testing. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(5):1091–1128, 2024. arXiv:1906.07801.

[9] Keith B. Hoffman, Mo Dimbil, Colin B. Erdman, Nicholas P. Tatonetti, and Brian M. Overstreet. The Weber effect and the United States Food and Drug Administration's Adverse Event Reporting System (FAERS): analysis of sixty-two drugs approved from 2006 to 2010. *Drug Safety*, 37(4):283–294, 2014.

[10] Lingling Li and Martin Kulldorff. A conditional maximized sequential probability ratio test for pharmacovigilance. *Statistics in Medicine*, 29(2):284–295, 2010.

[11] Charles F. Manski. *Identification Problems in the Social Sciences*. Harvard University Press, 1995.

[12] Jeffery L. Painter, François Haguinet, and Andrew Bate. CTG-DB: An ontology-based transformation of ClinicalTrials.gov to enable cross-trial drug safety analyses, 2026.

[13] Antoine Pariente, Fleur Gregoire, Annie Fourrier-Reglat, Françoise Haramburu, and Nicholas Moore. Impact of safety alerts on measures of disproportionality in spontaneous reporting databases: the notoriety bias. *Drug Safety*, 30(10):891–898, 2007. PMID: 17867726.

[14] Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4):576–601, 2023. arXiv:2210.01948.

[15] Aaditya Ramdas and Ruodu Wang. Hypothesis testing with e-values, 2024.

[16] Paul R. Rosenbaum. *Observational Studies*. Springer, 2nd edition, 2002.

[17] Kenneth J. Rothman, Stephan Lanes, and Susan T. Sacks. The reporting odds ratio and its advantages over the proportional reporting ratio. *Pharmacoepidemiology and Drug Safety*, 13(8):519–523, 2004. PMID: 15317031.

[18] Neil K. R. Sehgal, Jena Shaw Tronieri, and Lyle Ungar. Self-reported side effects of semaglutide and tirzepatide in online communities, 2026.

[19] Yihao Tan, Marianthi Markatou, and Saptarshi Chakraborty. A review of statistical methods for spontaneous reporting system data mining: Signal detection and beyond, 2026.

[20] Rosanne J. Turner, Alexander Ly, and Peter D. Grünwald. Generic E-variables for exact sequential k-sample tests that allow for optional stopping. *Journal of Statistical Planning and Inference*, 230:106116, 2024. arXiv:2106.02693.

[21] Tyler J. VanderWeele and Peng Ding. Sensitivity analysis in observational research: introducing the E-value. *Annals of Internal Medicine*, 167(4):268–274, 2017.