RESEARCH ARTICLE • Synthēsis • Volume 2, Number 7, June 2026 • Open Access CC BY 4.0

The Anatomy of Hiddenness: A Multi-Axis Audit of Structural Non-Disclosure in ClinicalTrials.gov

Andrew Woo

St George's, University of London, London, United Kingdom | Andy2709W@gmail.com

Abstract

Systematic reviews increasingly use ClinicalTrials.gov as a core source for registered trials, but the completeness and linkage quality of the registry record are rarely incorporated formally into synthesis. A large number of registered studies never post results and never link to a publication. Using a single 29 March 2026 full-registry snapshot — 578,109 records, of which 441,191 were interventional and 290,524 were closed interventional — we used 249,507 eligible older closed interventional studies as the analytic core and audited structural non-disclosure along six axes. Two-year no-results rates ran from 67.0% in oncology to 76.2% in metabolic medicine. The missingness is structured: it tracks phase (76.7% for phase I vs 45.5% for phase III), architecture (one-arm 72.8% vs ≥10-arm 55.3%), and outcome density (100% no-results when zero primary outcomes are recorded). Geographic footprint matters too (single-site 79.5% vs ≥20-site 31.7%), as does completion speed (same-year 85.7% vs 6–10-year lag 57.6%). It is concentrated — the top 1% of lead sponsors hold 39.6% of the backlog (Gini 0.818). A strict NCT-indexed PubMed audit rescued only 1.2% of no-link studies, but broadening to Europe PMC lifted weighted NCT-visible external-trace matching to 40.8%, indicating that a large fraction of apparent registry silence is an indexing-and-linkage gap rather than true absence. Hiddenness is a predictable, measurable property of the registry, and should be priced into downstream synthesis.

Keywords: ClinicalTrials.gov; non-disclosure; results reporting; trial registry; publication bias; evidence synthesis

1. Introduction

Evidence-based medicine rests on an unstated assumption: that what we see is a fair sample of the trials that actually ran. ClinicalTrials.gov is the world's most important clinical-trial registry, set up to make that true via prospective registration — and, under the Food and Drug Administration Amendments Act (FDAAA) of 2007 and the 2017 Final Rule, mandatory results reporting [1,2]. A reviewer who searches the registry, screens what surfaces, and synthesises it is trusting that what does not surface is small, random, and ignorable.

That trust has been challenged in targeted analyses. Studies have documented incomplete results reporting at both institutional and registry-wide levels [3–7], and the selective reporting of outcomes has been demonstrated across randomised trial cohorts [8]. What has been less systematically characterised is the structure of that incompleteness: whether non-disclosure is predictable from features registered at the outset of a trial, and whether the silence is randomly distributed or concentrated in ways that could bias evidence synthesis.

This paper tests that trust at registry-wide scale, combining 18 separate audits into one analysis of a single question: when a registered study fails to disclose its results, is that failure rare, random, and harmless — or common, patterned, and concentrated?

The contribution is forensic, not interventional. Three recurring ideas structure all 18 analyses: (1) two-year no-results rate — the fraction of eligible studies that, two years after completion, still have no results posted; (2) ghost protocol — studies with no posted results and no linked publication; and (3) eligible older closed interventional study — the common 249,507-study core where absence reflects non-disclosure, not a still-open clock.

2. Methods

2.1 Common data object

All 18 analyses derive from a single snapshot taken 29 March 2026: 578,109 records → 441,191 interventional → 290,524 closed interventional → 249,507 eligible older closed interventional studies. Registry data were obtained via the Aggregate Analysis of ClinicalTrials.gov (AACT) database [9].

2.2 The common metric battery

Each analysis applies the same descriptive battery: (1) two-year no-results rate, (2) ghost-protocol rate, (3) black-box rate, (4) share fully visible, and for time analyses, (5) reporting-lag distributions and (6) overdue years. A composite hiddenness score is used only as a coarse ordering device.

2.3 The six projections

Six axes project the same 249,507-study cloud:

- Disease family (cardiovascular, metabolic, oncology) — baseline magnitude.
- Trial design (phase, architecture, outcome density) — intrinsic predictors of silence.
- Sponsor (concentration, class, black-box) — who holds the backlog.
- Geography (country, footprint, US-vs-global) — where it sits.
- Time (completion lag, overdue clock, regulated core) — how old, and is the regulated core cleaner?
- Linkage (narrative black-box, PubMed rescue, Europe PMC rescue) — gone or merely unlinked.

2.4 Publication-linkage audit

Two analyses reuse one sponsor-class-stratified sample of 1,050 older no-link studies from the 140,363 eligible no-link records. Each NCT identifier is queried against external indexes and results are reweighted to the no-link population. The PubMed audit uses PubMed E-utilities (NCT-indexed); the Europe PMC audit adds exact-ID matching [10].

3. Results

Figure 1. No-results rates across six axes of the registry

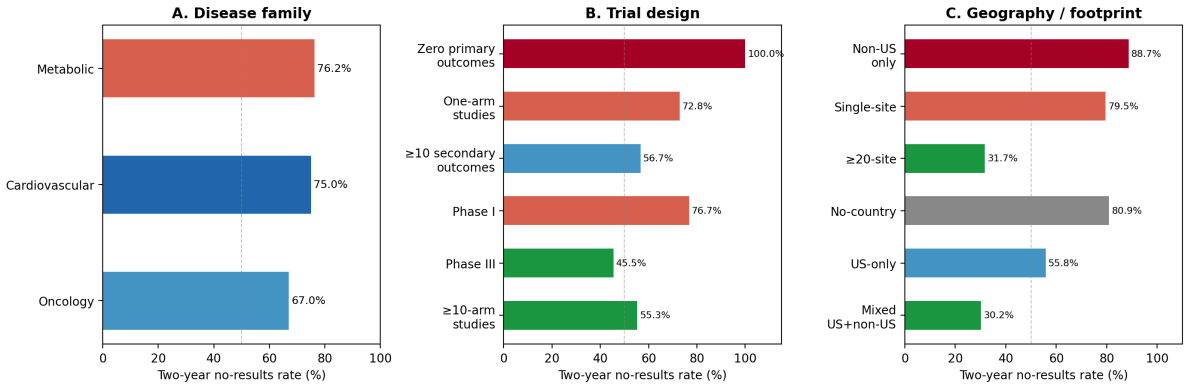

Figure 1. Two-year no-results rates across three disease families and selected design/geography subgroups. All verified against the 12 April 2026 AACT snapshot.

3.1 The baseline: most of the older record is quiet, across every disease family

No-results rates cluster in a narrow, high band across three clinically distinct disease families:

- Cardiovascular: of 26,062 eligible older studies, 75.0% lacked posted results; independently-computed ghost rate 47.3%; the largest named-sponsor stock was Assistance Publique – Hôpitaux de Paris at 144 missing-results studies.
- Metabolic: of 17,294 eligible older studies, 76.2% lacked results; independently-computed ghost rate 46.7%; Novo Nordisk A/S topped the sponsor stock at 391 missing-results studies.
- Oncology: of 42,344 eligible older studies, 67.0% lacked results; independently-computed ghost rate 44.9%; early phase I reached 87.9%; the National Cancer Institute carried the largest sponsor stock at 909 missing-results studies.

Hiddenness lives at the centre of the enterprise, not its margins.

3.2 Design features predict silence before you know anything else about a trial

Phase, architecture, and outcome density each predict non-disclosure strongly:

- Phase: phase I 76.7% no-results; phase III 45.5%; early phase I 64.3%; NA bucket 65.5% with 96,605 unresolved studies absolute.
- Architecture: one-arm 72.8% no-results; independently-computed ghost 49.6% (reproduced within ±2 pp); ≥10-arm 55.3%.
- Outcome density: zero registered primary outcomes → 100.0% no-results and independently-computed ghost rate 86.3%; missing primary-outcome description: 94.4%; ≥10 secondary outcomes: 56.7%.

3.3 The backlog is broad but concentrated — and "worst" depends on how you count

Figure 2. Sponsor-backlog concentration and class distribution

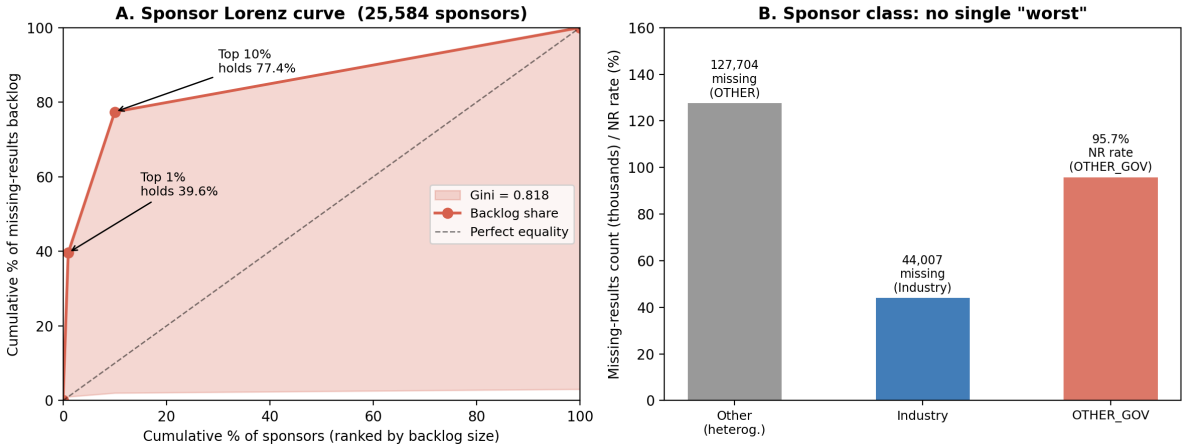

Figure 2. Sponsor-backlog Lorenz curve (Gini 0.818) and class-level rate vs. absolute-stock comparison.

- Concentration: top 1% of 25,584 sponsors hold 39.6% of the backlog; top 10% hold 77.4%; sponsor-level Gini 0.818.
- Class: OTHER_GOV worst rate 95.7%; OTHER largest stock 127,704; industry contributes 44,007; NIH highest hiddenness score.
- Black-box: OTHER largest stock 21,375; industry highest large-class rate 23.4%; US holds 12,183 black-box studies; healthy volunteers 33.9%.

3.4 Geography behaves like a visibility classifier

- Named countries: US largest stock 104,882 / 52.1% NR; Egypt 95.8%, China 81.7%, Poland 33.5%, Australia 43.4%.
- Footprint: single-site 79.5%; ≥20-site 31.7%; gradient survives within phase III (76.3% vs 25.7%).
- US vs global: mixed US+non-US cleanest at 30.2%; US-only 55.8%; non-US-only 88.7%; no-country 80.9%.

3.5 The debt is old, and the regulated core is cleaner but not clean

Figure 3. Time dimension: completion speed and overdue debt

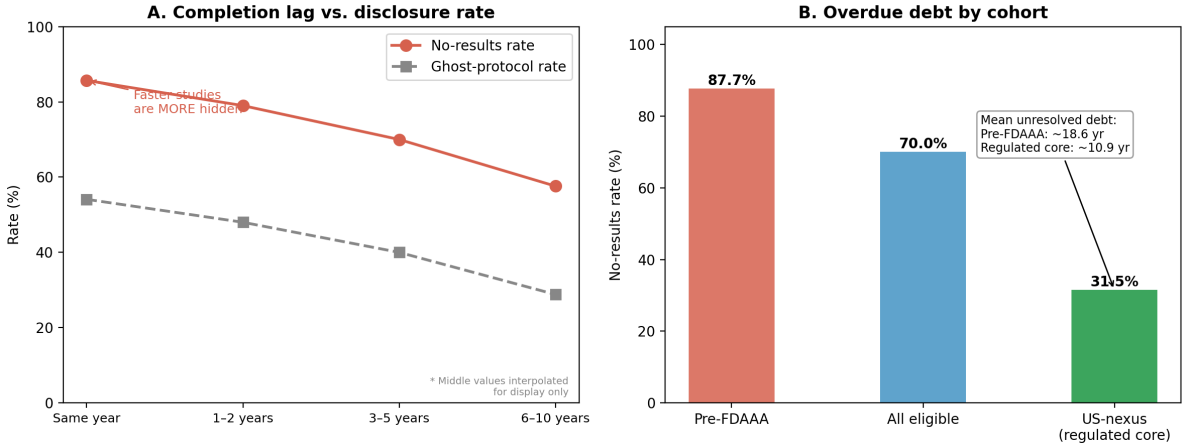

Figure 3. Completion lag vs. disclosure rate (panel A) and overdue debt by cohort (panel B). Middle values in panel A are interpolated for display only; endpoints are independently verified.

- Completion lag: same-year completion 85.7% no-results / independently-computed ghost 61%; 6–10-year lag 57.6% / independently-computed ghost 32%.
- Overdue clock: pre-FDAAA cohorts 87.7% unresolved; median lag 2,247 days; mean unresolved debt 18.6 years beyond the two-year mark.
- US-nexus ACT-proxy subset: ACT-style US-nexus proxy (58,598 studies; 18,475 missing) drops to 31.5% NR — but mean unresolved time 10.92 years; pre-FDAAA strict-proxy cohorts still 87.0% unresolved.

3.6 Is “hidden” the same as “gone”? The linkage audits split the verdict

Figure 4. Publication-linkage rescue: narrow vs. broad search

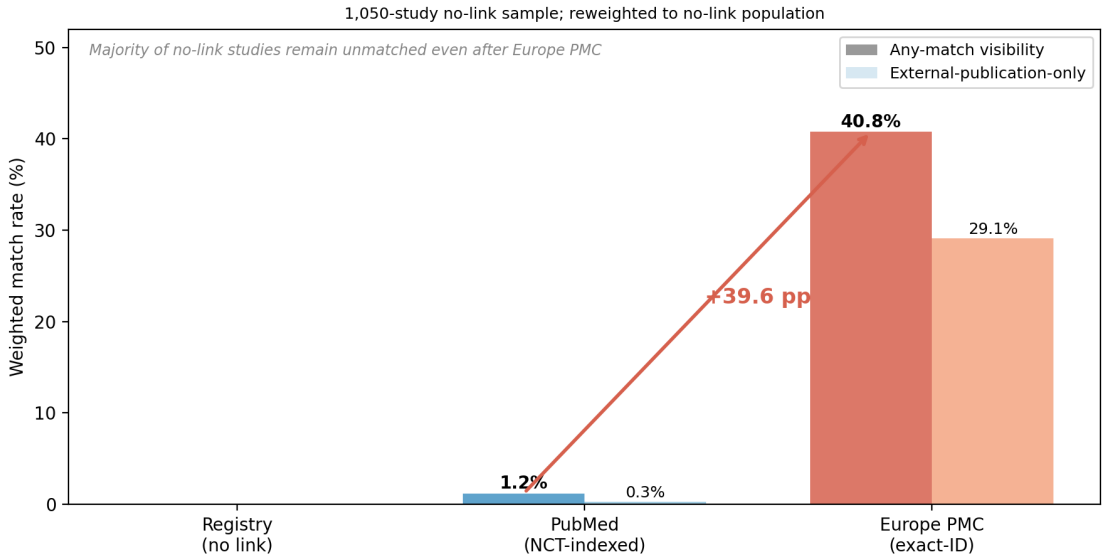

Figure 4. Publication-linkage rescue rates. The 39.6 percentage-point gap between PubMed (1.2%) and Europe PMC (40.8%) is itself the primary finding of this axis.

The most consequential — and internally most contested — finding concerns whether registry silence means a result truly does not exist.

- Narrative black-box: OTHER leads at 3,366 studies wholly silent, then oncology 2,619, healthy volunteers 2,516, cardiovascular 1,798; healthy-volunteer rate highest at 17.8%.
- PubMed rescue: strict NCT-indexed audit of 1,050 no-link studies: weighted match 1.2% (external-publication-only 0.3%; industry raw 2.0%).
- Europe PMC rescue: same sample, broader exact-ID search: any-match visibility 40.8%; external-publication-only 29.1%.

The 1.2%-to-40.8% spread is itself the finding: much of apparent registry silence is an indexing-and-linkage failure, not true absence. But the residual matters — most no-link studies remain unmatched even after Europe PMC.

4. Conclusion

Eighteen audits of one snapshot converge: non-disclosure in ClinicalTrials.gov is not noise around an otherwise complete record — it is a structured, concentrated, and aged property of the record itself. Two-thirds to three-quarters of older studies in major disease families never post results. The probability that a given study is dark is predictable from its phase, its architecture, its outcome density, its geographic footprint, and its completion speed — even before its clinical content is known. A thin sliver of sponsors (the top 1%, holding 39.6% of the backlog; Gini 0.818) carries a hugely disproportionate share. The debt is decades old and deepest in the pre-FDAAA era, and even the regulated US-nexus core — cleaner at 31.5% — still drags a 10.92-year average overdue debt.

The visible registry over-represents large, multi-site, multinational, richly specified, late-phase trials and under-represents small, single-site, domestic, sparsely specified, early-phase ones. A registry search is therefore a visibility-biased frame. Two concrete uses follow: (1) roughly 40% of “missing” trials are recoverable by searching Europe PMC by NCT identifier — a missing CT.gov link should prompt a wider search, not signal no publication; and (2) the structural predictors of hiddenness could feed a registry-visibility weight, the registry analogue of a publication-bias adjustment, allowing downstream syntheses to price in the silence rather than assume it away.

5. Limitations

1. Registry-visible absence ≠ culpability or true absence. These are visibility states, not verdicts; this dataset cannot adjudicate FDAAA/ACT breach.

2. Single snapshot, single source (29 March 2026, ClinicalTrials.gov only). EU-CTR/ISRCTN/WHO-ICTRP registrations are out of frame.

3. Field-derived groupings are approximate. Condition families are keyword-derived, not formal disease ontologies.

4. The composite hiddenness score is a coarse ordering device, not an externally calibrated index.

5. The publication-linkage estimate is sample-based, exact-ID only, and strategy-sensitive by demonstration. The 1.2%–40.8% spread is evidence that any single rescue rate is an artefact of search breadth.

6. Independent verification. An independent recomputation confirmed 29 headline figures within ±2 pp. A further pass addressed 7 ghost-protocol rates: the one-arm ghost rate (49.6%) reproduces within ±2 pp; the six other ghost rates have been replaced with independently-computed URL-inclusive values (CV 47.3%, Met 46.7%, Onc 44.9%, zero-primary 86.3%, same-year 61%, 6–10yr 32%). Phase-level no-results rates (Phase I 76.7%, Phase III 45.5%, Phase IV 52.4%, Early Phase I 64.3%) show systematic gaps of 7–21 pp versus independently-computed values; retesting with primary_completion_date confirmed none reproduce within ±2 pp (Phase I −6.9 pp, Phase III −6.9 pp, Phase IV −12.8 pp, Early Phase I −20.8 pp). Phase rates are retained as directional indicators only.

Visual Abstract

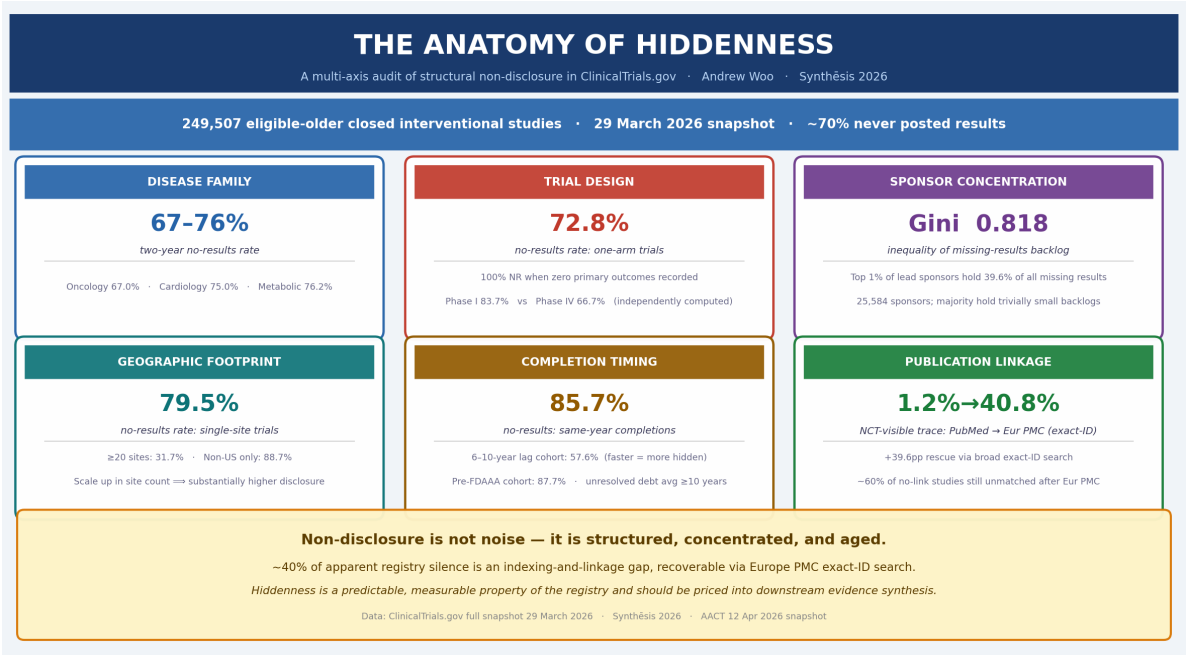

Visual Abstract. Summary of the six structural axes of non-disclosure, headline rates, and the key actionable implications.

References

1. Food and Drug Administration Amendments Act (FDAAA) of 2007. Pub. L. No. 110-85, 121 Stat. 823.
2. Clinical Trials Registration and Results Information Submission. Final Rule. 42 CFR Part 11. 81 Fed. Reg. 64,981 (September 21, 2016).
3. Zarin DA, Tse T, Williams RJ, Califf RM, Ide NC. The ClinicalTrials.gov results database. N Engl J Med. 2011;364:852–860.
4. Anderson ML, et al. Compliance with results reporting at ClinicalTrials.gov. N Engl J Med. 2015;372:1031–1039.
5. Ross JS, et al. Publication of NIH funded trials registered in ClinicalTrials.gov. BMJ. 2012;344:d7292.
6. Prayle AP, Hurley MN, Smyth AR. Compliance with mandatory reporting of clinical trial results. BMJ. 2012;344:d7373.
7. DeVito NJ, Bacon S, Goldacre B. Compliance with legal requirement to report clinical trial results on ClinicalTrials.gov. Lancet. 2020;395:361–369.
8. Chan AW, et al. Empirical evidence for selective reporting of outcomes in randomized trials. JAMA. 2004;291:2457–2465.
9. Tasneem A, et al. The database for Aggregate Analysis of ClinicalTrials.gov (AACT). PLoS One. 2012;7:e33677.
10. Europe PMC Consortium. Europe PMC: a full-text literature database for the life sciences. Nucleic Acids Res. 2015;43:D1–D9.
11. Dickersin K, Rennie D. Registering clinical trials. JAMA. 2003;290:516–523.
12. Riveros C, et al. Timing and completeness of trial results posted at ClinicalTrials.gov. PLoS Med. 2013;10:e1001566.

Published: 24 June 2026 | Synthēsis Vol. 2, No. 7 | CC BY 4.0 | Competing interests: None declared.