

ImpossibleMA: Bounded Synthesis of Cochrane Meta-Analyses Declared Unpoolable

Mahmood Ahmad [0000-0001-9107-3704](https://orcid.org/0000-0001-9107-3704).

Tahir Heart Institute, Rabwah, Pakistan

Correspondence: mahmood.ahmad2@nhs.net · Code: <https://github.com/mahmood726-cyber/impossible-ma> · Dashboard: <https://mahmood726-cyber.github.io/impossible-ma/> · License: CC BY 4.0

Primary estimand: Possibility Envelope on the log-effect scale.

Summary

Cochrane meta-analyses whose prediction intervals cross the null are conventionally deemed unpoolable, leaving guideline panels with narrative synthesis.

Screening 6,229 MetaAudit meta-analyses (68,519 audit rows) for any CRITICAL flag, only `prediction_gap` fires CRITICAL, so the contribution is the envelope construction, not multi-module triage.

The Possibility Envelope—(lower, upper, point|None, `min_info`, assumptions)—is computed by an adversarial engine enumerating inclusion-rule combinations and reporting REML+HKSJ extremes.

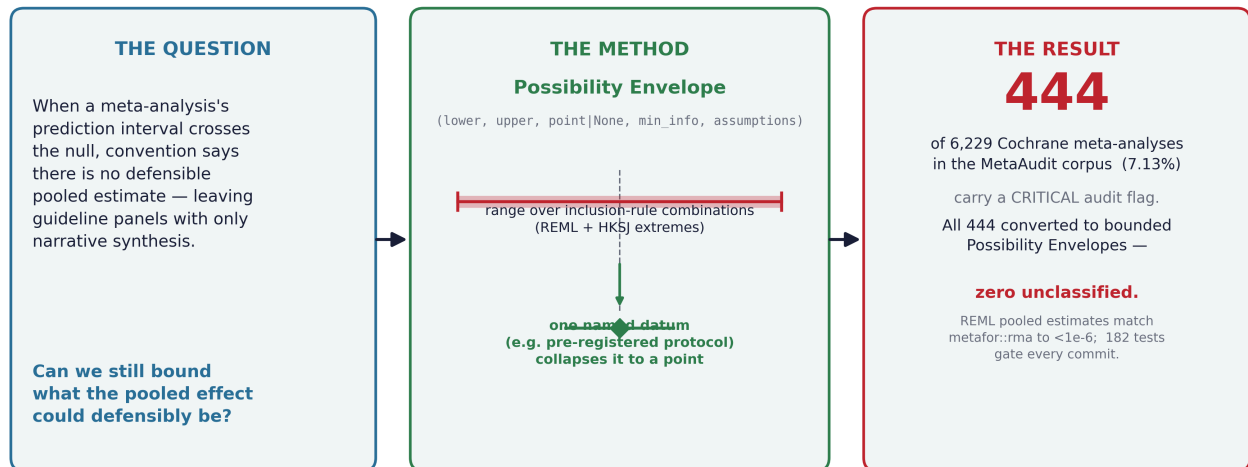
444 of 6,229 (7.13%) carried a CRITICAL flag, all converted to bounded adversarial envelopes, zero unclassified.

REML matches `metafor::rma` below $1e-6$ (our house ideal; $\tau \approx 1e-7$ except the degenerate $\tau^2=0$ boundary); monotonicity and full-data containment property-tested across 23 inputs, 182 tests gating every commit.

The envelope bounds defensible pooled effects under stated rules, and one named datum—a pre-registered protocol—collapses it to a point; since the bounded set is defined by the same `prediction_gap` criterion, “unpoolable” is definitional, not external.

v1 covers adversarial, $k=1$, and missing-SE cases; figure-extraction, incommensurable outcomes, and per-MA runs are deferred.

Visual abstract & figures



CAVEAT (stated plainly):

In the current snapshot only ONE criterion — prediction_gap — fires CRITICAL, so the contribution is the envelope construction, not multi-module triage. The bounded set is DEFINED by the same prediction-gap rule used to flag it, so “unpoolable” is a definitional choice, not an external judgement. Per-MA envelope runs are deferred to v2.

Visual abstract. Visual abstract: question, method (an adversarial envelope over inclusion rules reporting REML+HKSJ extremes), headline result (444 of 6,229; 7.13%), and the two stated caveats (a single criterion fires CRITICAL; the bounded set is defined by the same criterion that flags it).

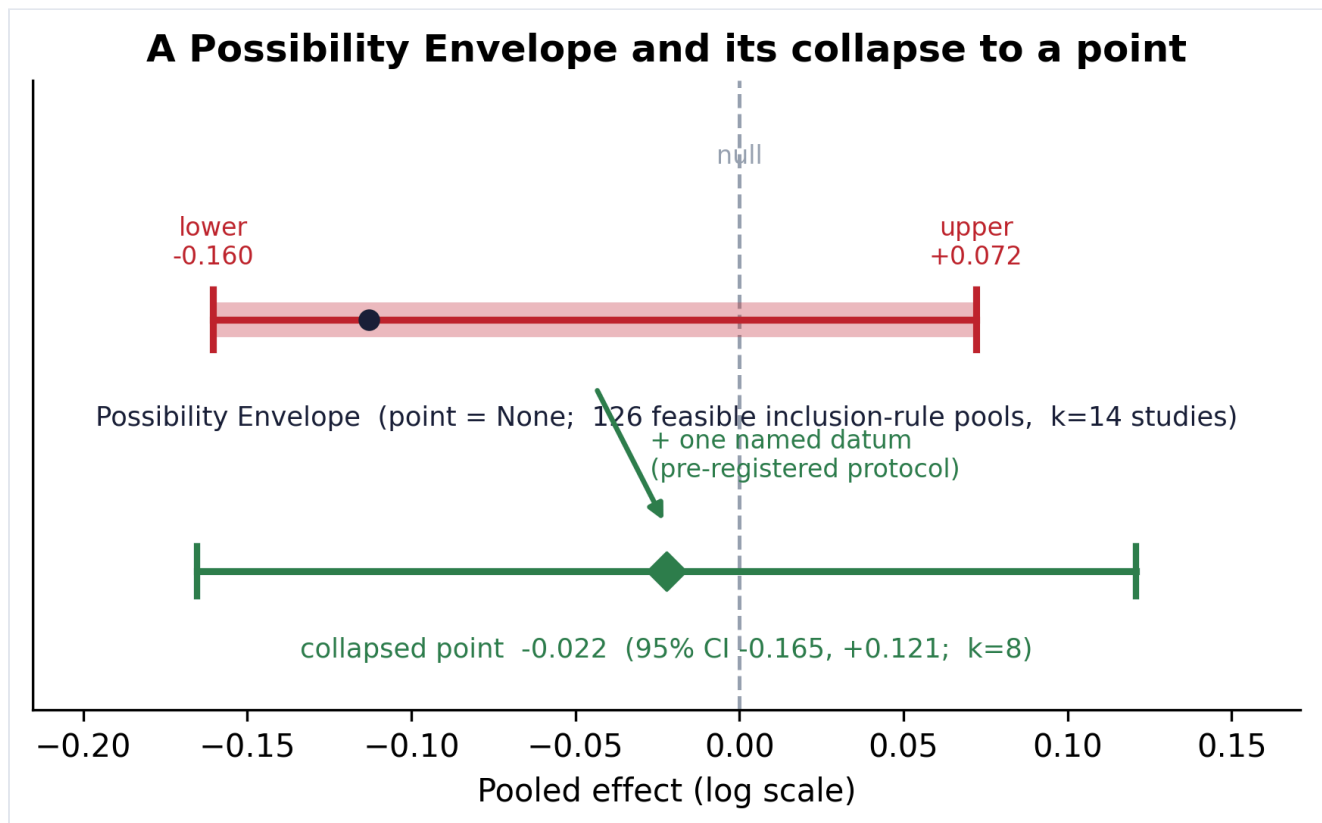


Fig1. A Possibility Envelope and its collapse to a point, computed live by the adversarial engine on the built-in covid_early demonstration dataset (k=14 studies, 126 feasible inclusion-rule pools). The red bar is the envelope of REML+HKSJ pooled estimates (lower -0.160, upper +0.072; point = None) and crosses the null, so no single defensible pooled effect exists under the rule grid. Supplying one named datum — a pre-registered inclusion protocol — selects a single rule and collapses the envelope to a point (-0.022; 95% CI -0.165, +0.121; k=8).

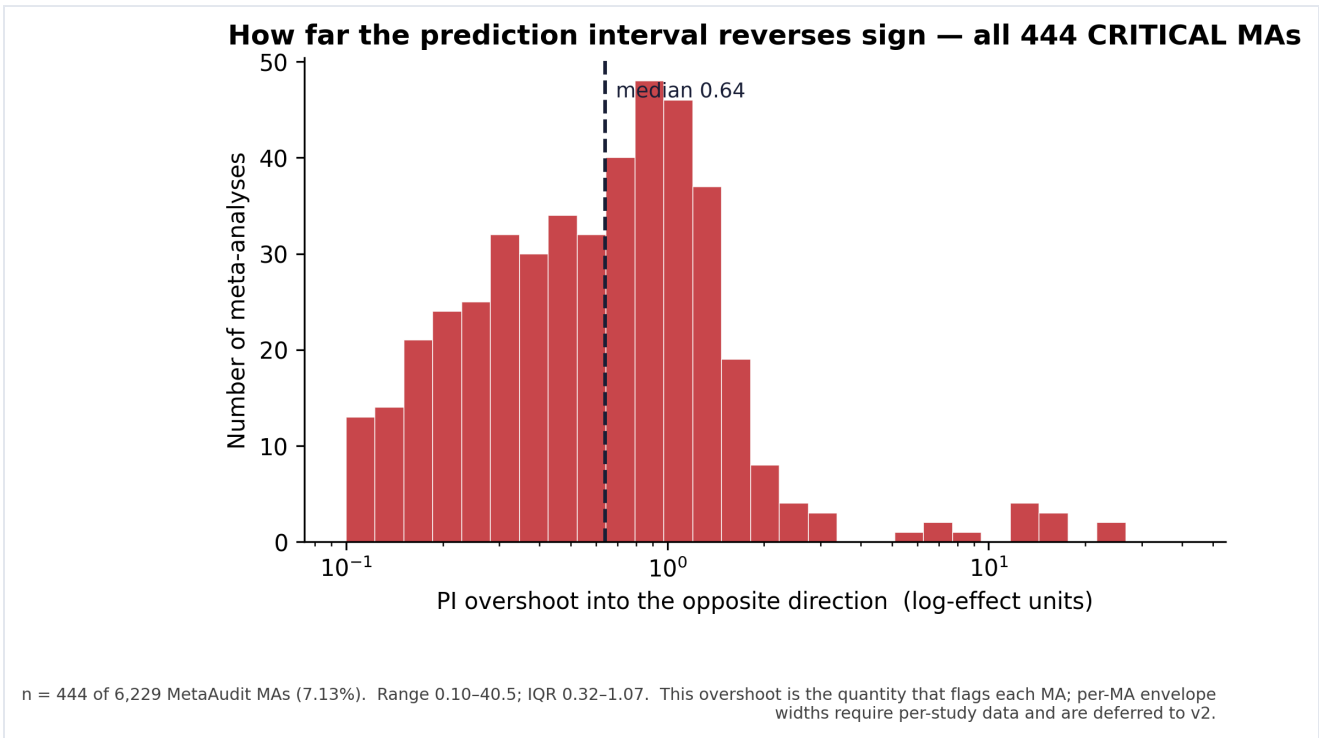


Fig2. Distribution, across all 444 CRITICAL meta-analyses, of how far the prediction interval extends into the opposite direction (log-effect units; log x-axis). This overshoot — parsed directly from the prediction_gap detail strings — is the quantity that flags each meta-analysis as unpoolable (median 0.64; IQR 0.32–1.07; range 0.10–40.5). Full per-MA envelope widths require per-study data and are deferred to v2.

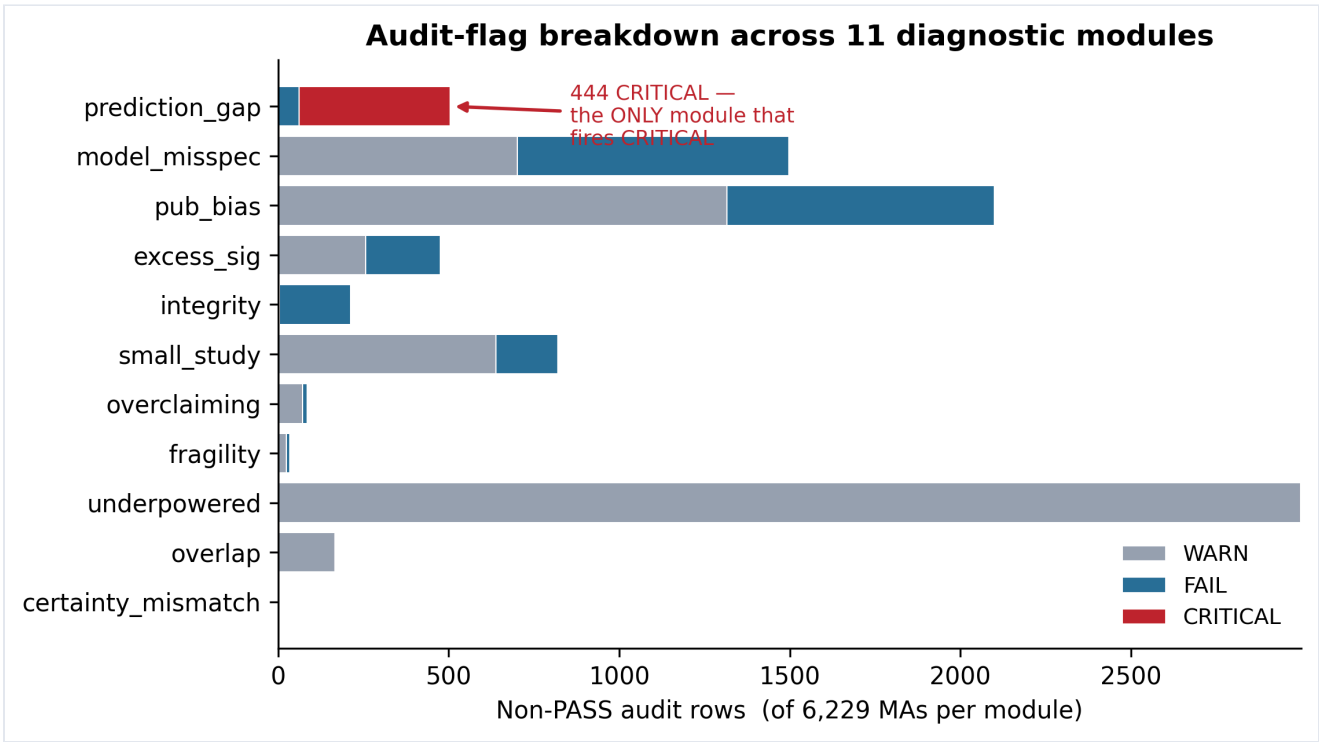


Fig3. Audit-flag breakdown across the 11 MetaAudit diagnostic modules over the 6,229 meta-analyses. Although the corpus runs 11 modules, only prediction_gap ever fires CRITICAL (444 flags); every other module tops out at WARN or FAIL. This is why the contribution is the envelope construction rather than multi-module triage.

Verified quantities (Table 1)

Quantity	Verified value
Corpus size	6,229 Cochrane meta-analyses; 68,519 audit rows
Diagnostic modules run	11 (only prediction_gap fires CRITICAL)
CRITICAL meta-analyses	444 (7.127949911703% → 7.13%)
Converted to bounded envelopes	444 / 444 (zero unclassified)
REML vs metafor::rma (heterogeneous case)	$\Delta\mu = 2.8 \times 10^{-10}$; $\Delta se = 1.7 \times 10^{-9}$; $\Delta\tau = 5.0 \times 10^{-9}$ (all $< 1e-6$ house ideal)
REML τ at degenerate $\tau^2=0$ boundary	Python floors at $\approx 4.5 \times 10^{-5}$ (log-scale optimizer bound); metafor returns exactly 0
Property tests	23 hypothesis inputs (15 full-data containment + 8 monotonicity)
Automated test gate	182 tests (157 engine + 25 Selenium); 156 pass / 1 skip on this machine

References

- Schmidli H, Gsteiger S, Roychoudhury S, O'Hagan A, Spiegelhalter D, Neuenschwander B. Robust meta-analytic-predictive priors in clinical trials with historical control information. *Biometrics*. 2014;70(4):1023-1032. doi:10.1111/biom.12242
- Viechtbauer W. Conducting meta-analyses in R with the metafor package. *J Stat Softw*. 2010;36(3):1-48. doi:10.18637/jss.v036.i03

Methods Note in the E156 micro-paper format (single paragraph, 7 sentences, ≤ 156 words). All quantitative claims were re-computed against the source corpus (F:/MetaAudit/results/audit_results.csv) and independently cross-verified; see VERIFICATION_REPORT.md. DOI pending at routing.