PPN WORLD RESEARCH PAPER NO. 3 · EDITION 1 OPEN ACCESS

ORIGINAL INVESTIGATION

Reporting of the Registered Primary Outcome in Drug-Trial Press Releases

Two Preregistered Cross-sectional Studies

Émile Bolduc[1*]

[1] PPN World Research Center, Bollé Communications inc., Montréal, Québec, Canada

* ppnw@ppnsource.com · ORCID https://orcid.org/0009-0005-7882-4741

Citation: Bolduc, É. (2026). Reporting of the registered primary outcome in drug-trial press releases: Two preregistered cross-sectional studies (PPN World Research Paper No. 3, Edition 1). PPN World. https://www.ppnworld.-com/research/clinical-trial-press-releases

Registered: November 11, 2026; November 12, 2026

Published: November 12, 2026

Copyright: © 2026 PPN World. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International licence (CC BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

Keywords: press release; clinical trial; primary outcome; trial registration; ClinicalTrials.gov; outcome reporting; spin; SEC filings; content analysis; pre-registration

Key Points

Question Do drug-trial press releases say whether the registered primary outcome was met, and does that depend on when the release is issued?

Findings In 2 preregistered cross-sectional studies of 262 and 2,141 release–trial pairs, releases said whether a randomized trial's primary outcome was met for 58% and 55% of pairs; headlines stated its result for 16% and 13%. First announcements gave a verdict for 82% of pairs, conference or publication releases for 39%; failures were 31% of verdicts at first announcement and 12% later.

Meaning A failed trial is announced mainly once; later releases mostly concern trials that succeeded.

Abstract

Importance. For many drug trials, the sponsor's press release is the first public account of the result, reaching investors and journalists before peer review. Whether releases report the registered primary outcome has not been studied.

Objective. To measure how drug-trial press releases report the registered primary outcome and whether this depends on the kind of announcement.

Design, Setting, and Participants. Two preregistered cross-sectional studies of English-language releases naming a ClinicalTrials.gov identifier: 770 published July to October 2026 in a press-release archive (study 1) and 6,856 filed with the US Securities and Exchange Commission from 2017 to mid-2026 (study 2), testing 3 hypotheses from study 1's exploratory analyses, registered before its releases were read. Pairs were compared with the registry record in force on the release date. Language models coded releases with a fixed codebook; a second model recoded a random subsample.

Main Outcomes and Measures. Whether the release said if the registered primary outcome was met (verdict), whether the headline stated that outcome's result, and the share of verdicts that were "not met," by kind of announcement.

Results. The main analysis sets comprised 262 pairs (196 trials) and 2,141 pairs (813 trials; 335 companies). Releases stated a verdict for 58.4% (95% CI, 50.2%-65.3%) and 54.7% (95% CI, 47.6%-60.9%) of pairs, and headlines stated the primary result for 15.6% and 12.6%. In study 2, first announcements stated a verdict more often than conference or publication releases (82.3% vs 39.3%; Holm-adjusted $P < .001$), and more of their verdicts were "not met" (30.8% vs 12.4%; Holm-adjusted $P < .001$). At first announcement, headlines stated the result for 42.6% of pairs whose trial met the primary outcome vs 28.6% of those whose trial did not (Holm-adjusted $P = .02$), a difference not significant among domestic filers or phase 3 trials. In exploratory analyses, 62% of trials first announced as successes and 31% of failures had a later release.

Conclusions and Relevance. Drug-trial press releases report failure mainly at the first announcement; later releases concern trials that succeeded or omit the ver-

dict. Readers of later releases should look to the first announcement or the registry.

Introduction

For many industry-sponsored drug trials, the sponsor's press release is the first public account of the result, issued before any journal article.[1] Releases from listed companies move share prices before any peer review,[2] and what a release says carries into news coverage: exaggeration in health press releases is associated with exaggeration in the news stories that follow them.[3,4]

Spin, reporting that makes a treatment look more beneficial than its results justify, is well documented in articles reporting trials with nonsignificant primary outcomes[5,6] and in press releases on randomized trials compared with the articles they announce.[7] Changes between registered and published primary outcomes are common in journals.[8,9] Prospective registration exists so that readers can check a reported outcome against the one named before the trial began.[10] We found no study that compared press releases with the registry record, and none that followed what releases say across a trial's life, from the first announcement of top-line results to later conference and journal releases.

We conducted 2 preregistered studies. Study 1 measured how releases published over 3 months in 2026 report a trial's registered primary outcome: whether they name it, say whether it was met, state its result in the headline, or call another outcome primary. Its exploratory analyses suggested that the answer depends on when in a trial's life the release is issued. Study 2 tested 3 hypotheses drawn from those analyses, registered before any of its releases were read, on 9.5 years of releases filed with the US Securities and Exchange Commission (SEC).

Methods

Registration and Reporting

Each study's plan, codebook, analysis code and pilot results were committed to a version-controlled repository, with a manifest of their SHA-256 hashes, before any release in the study's frame was read or coded (study 1, November 11, 2026; study 2, November 12, 2026); the manifests were timestamped with OpenTimestamps. Each study has 1 numbered deviation, neither of which affects a main estimate or hypothesis (eAppendix 1 in Supplement 1). The registrations, deviations, data, codings, prompts and code are deposited openly. The report follows the STROBE guideline for cross-sectional studies.[11] The studies used public documents and involved no human participants; ethics review was not sought.

Study 1: A Press-Release Archive

The frame was every English-language release published July 1 to October 6, 2026, in PPN World's press-release archive, which collects releases from distribution services, companies and public bodies, whose page had been saved in a census of that archive's releases in October 2026[12] and named at least 1 ClinicalTrials.gov identifier ("NCT" followed by 8 digits): 770 releases. Release text was extracted with Mozilla Readability. A pilot of 199 releases published after the window shaped the design and is excluded.

Study 2: Releases Filed With the SEC

Companies listed in the United States file material news with the SEC, usually by attaching the press release to a Form 8-K (domestic issuers) or Form 6-K (foreign issuers).[13] EDGAR full-text search for "NCT*" in Forms 8-K, 6-K and their amendments, run month by month on November 12, 2026, returned 10,708 documents; 9,820 were press-release exhibits (EX-99) or Form 6-K documents filed January 1, 2017, to June 30, 2026, outside a 140-document pilot, and 9,174 of these named a ClinicalTrials.gov identifier. A language model (Claude Haiku 5.5) read the first 700 words of each and returned the kind of document and, for a press release, its headline copied verbatim. After duplicate texts (earliest filing kept) and copies of pilot documents were removed, 6,856 press releases from 700 companies remained. Text was extracted from each document's HTML, block by block. The filing date stands in for the release date.

Registry Record

For each release–trial pair, the registry record in force on the release date was taken from AACT, a relational copy of ClinicalTrials.gov:[14] the latest snapshot on or before the release date (study 1, monthly snapshots from July to November 2026; study 2, the first snapshot of each quarter from January 2017 to July 2026). A pair was flagged when the next snapshot listed different primary outcomes.

Coding

A codebook (eAppendix 2 in Supplement 1) defined, for each release and each trial it names, whether the release was issued by the trial's side (sponsor, partner or investigators' institution); whether it reports results of that trial; the kind of announcement (first announcement of top-line results, interim, conference presentation, journal publication, regulatory, final or full results, other); whether it names the registered primary outcome; its verdict on that outcome (met, not met, mixed, a number without a verdict, no verdict); whether it calls another outcome primary; whether the headline states the registered primary outcome's result; 3 kinds of spin; and whether a "met" verdict comes with a number. Whether each registered primary outcome measures efficacy, or only safety, pharmacokinetics or dose, was coded from the registry record.

Language models can match human coders on text annotation;[15] ours were Anthropic's Claude models, called through the Message Batches API with one request per release, the codebook as system prompt, the release text and registry records as input (not the publisher's name or website), answers constrained to a JSON schema with quoted evidence, reasoning at medium effort and default sampling. In study 1, Claude Sonnet 5.5 coded every release and Claude Opus 5.5 a random 150 (November 11, 2026). In study 2, Claude Haiku 5.5 coded every release and Opus a random 200 (November 12, 2026); Haiku was chosen before registration because, on study 1's double-coded releases, it agreed with Opus at least

as well as Sonnet had on every measure the hypotheses use. Agreement was measured with Krippendorff α;[16] a hypothesis using a measure with α below .667 was to be reported as undecided. Because both coders come from one model family, their agreement is an upper bound on reliability.

Measures and Hypotheses

The main analysis set comprised pairs in which a release issued by the trial's side reported results of a randomized trial (registry allocation) with at least 1 efficacy primary outcome. Measures were whether the release named the registered primary outcome; stated a verdict (met, not met or mixed); called another outcome primary; contained spin, among pairs without a "met" verdict; stated the primary outcome's result in the headline; and gave a number with a "met" verdict; and the share of pairs "not met."

Study 1 registered 2 hypotheses: headlines state the primary result less often when it was not met than when it was met (H1), and releases from sponsors with fewer than 50 registered studies state a verdict less often than releases from larger sponsors (H2). Study 2 registered 3: among first announcements, headlines state the primary result less often when it was not met (H1); conference and publication releases state a verdict less often than first announcements (H2); and among pairs with a verdict, "not met" is a smaller share in releases that are not first announcements (H3).

Statistical Analysis

Shares are reported with 95% CIs from a cluster bootstrap with 10 000 draws, resampling lead sponsors (study 1) or filing companies (study 2). Each hypothesis was tested as a difference in shares with a one-sided bootstrap P value and Holm adjustment across the study's hypotheses, and was supported if the 90% CI excluded 0 in the predicted direction and the adjusted P was below .05. Prespecified robustness analyses, none of which decides a hypothesis, used the double-coded subsample with each coder, removed pairs whose registry record changed around the release, kept 1 pair per trial, restricted to phase 3 trials, and, in study 2, clustered by lead sponsor and split 8-K from 6-K filers. Exploratory analyses are labeled as such: study 1's results by kind of announcement, which motivated study 2; registry results for study 1's releases without a verdict; and, in study 2, filing time and what followed a first announcement. Analyses were run in Node.js 24.

Results

Releases and Pairs

Study 1 comprised 770 releases, 729 of them issued by the trial's side, naming trials in 1,222 release–trial pairs; 470 pairs reported results, and the main set comprised 262 pairs on 196 trials from 129 lead sponsors (Table 1). Study 2 comprised 6,856 releases; 6,834 were coded (22 long filings exceeded the model's output limit) and 6,815 were issued by the trial's side, in 13,906 pairs; 5,126 reported results, and the main set comprised 2,141 pairs on 813 trials from 335 filing companies (44% of pairs from foreign issuers filing Form 6-K).

Table 1. Releases, Pairs and Main Analysis Sets.

	Study 1: press-release archive	Study 2: SEC filings
Period	July 1–October 6, 2026	January 1, 2017–June 30, 2026
Releases naming a ClinicalTrials.gov identifier	770	6,856
Coded	770	6,834
Issued by the trial's side	729	6,815
Release–trial pairs	1,222	13,906
Pairs without a registry record on the release date	23	450
Pairs reporting results	470	5,126
Early-phase pairs (safety, pharmacokinetic or dose primary outcomes)	127	1,706
Main set: randomized, efficacy primary outcome	262	2,141
Trials in the main set	196	813
Clusters (lead sponsors; filing companies)	129	335
Registry record changed around the release	9	124
Measurement coder; check coder (releases)	Sonnet 5.5; Opus 5.5 (150)	Haiku 5.5; Opus 5.5 (200)

What the Releases Said

The 2 studies gave similar estimates (Table 2). Releases named the registered primary outcome for 82.4% and 74.9% of pairs and stated a verdict for 58.4% (95% CI, 50.2%-65.3%) and 54.7% (95% CI, 47.6%-60.9%). Headlines stated the primary outcome's result for 15.6% (95% CI, 10.6%-21.5%) and 12.6% (95% CI, 10.0%-15.3%). Releases rarely called another outcome primary (2.7% and 4.3%). The primary outcome was reported as not met for 8.8% and 10.7% of pairs. Among pairs without a "met" verdict, spin appeared in 79.9% and 73.6%; study 1's coders disagreed on spin (α = .29), so its estimate is undecided.

Table 2. Estimates and Hypotheses, Main Analysis Sets.

Measure	Study 1, % (95% CI)	Study 2, % (95% CI)
Names the registered primary outcome	82.4% (95% CI, 75.7%-88.0%)	74.9% (95% CI, 70.0%-79.0%)
States a verdict (met, not met, mixed)	58.4% (95% CI, 50.2%-65.3%)	54.7% (95% CI, 47.6%-60.9%)
Calls another outcome primary	2.7% (95% CI, 0.4%-5.3%)	4.3% (95% CI, 3.3%-5.4%)
Spin, pairs without a "met" verdict[a]	79.9% (95% CI, 71.8%-86.9%)	73.6% (95% CI, 68.9%-77.8%)
Headline states the primary outcome's result	15.6% (95% CI, 10.6%-21.5%)	12.6% (95% CI, 10.0%-15.3%)
Gives a number with a "met" verdict	69.5% (95% CI, 59.1%-80.6%)	57.8% (95% CI, 51.1%-66.9%)
Not met	8.8% (95% CI, 5.0%-13.6%)	10.7% (95% CI, 8.5%-13.4%)
Study 2 hypotheses	**Shares**	**Difference, percentage points (90% CI); Holm P**
H1 Headline states the result at first announcement: met vs not met	42.6% vs 28.6%	14.0 (2.5 to 23.7); .02
H2 Verdict: first announcement vs conference or publication	82.3% vs 39.3%	42.9 (35.8 to 50.8); <.001
H3 "Not met" among verdicts: first announcement vs other releases	30.8% vs 12.4%	18.4 (13.6 to 22.4); <.001

[a] Study 1: undecided, coders' α = .29.

Kind of Announcement

In study 1's exploratory analysis, first announcements stated a verdict for 95.8% (46 of 48) of pairs and conference or publication releases for 33.7% (28 of 83) (Figure). Study 2 showed the same gradient: first announcements, 82.3%; final or full results, 77.3%; regulatory, 72.3%; publication, 51.7%; other, 38.5%; conference, 37.0%; interim, 24.5%.

Hypotheses

In study 1, H1 was not supported: headlines stated the primary result for 25.8% (33 of 128) of "met" pairs and 17.4% (4 of 23) of "not met" pairs (difference, 8.4 percentage points [90% CI, −7.3 to 24.1]; Holm-adjusted P = .18). H2 was supported: small sponsors' releases stated a verdict for 51.0% (80 of 157) of pairs, large sponsors' for 69.5% (73 of 105) (18.6 percentage points [90% CI, 7.3 to 30.5]; Holm-adjusted P = .008), but not among phase 3 trials (0.3 percentage points).

In study 2, all 3 hypotheses were supported (Table 2), and every measure they use had α of at least .71 (eTable 1 in Supplement 1). First announcements stated a verdict for 82.3% (95% CI, 77.2%-86.6%) of pairs and conference or publication releases for 39.3% (95% CI, 30.7%-47.2%) (H2: 42.9 percentage points [90% CI, 35.8 to 50.8]; Holm-adjusted P < .001). Among pairs with a verdict, "not met" was 30.8% (95% CI, 25.5%-36.9%) at first announcement and 12.4% (95% CI, 8.6%-18.3%) in other releases (H3: 18.4 percentage points [90% CI, 13.6 to 22.4]; Holm-adjusted P < .001). At first announcement, headlines stated the primary result for 42.6% of the 291 pairs whose trial met it and 28.6% of the 140 whose trial did not (H1: 14.0 percentage points [90% CI, 2.5 to 23.7]; Holm-adjusted P = .02).

H2 held in every robustness analysis, including phase 3 trials only (33.8 points). H3 held in every one but the double-coded subsample of 200 releases (62 and 65 main-set pairs as each coder coded them), where its difference was 6.7 and 8.2 points, with 90% CIs including 0; among phase 3 trials it was 13.3 points. H1 did not: its difference was 31.2 points among foreign issuers and 2.5 points (90% CI, −7.6 to 13.1) among domestic ones, and 10.5 points (90% CI, −5.8 to 23.7) among phase 3 trials (eTable 2 in Supplement 1).

Exploratory Analyses

Of 201 trials whose first announcement, filed before July 2025, said the primary outcome was met, 61.7% (124 of 201; 95% CI, 52.7%-70.0%) had at least 1 later release reporting their results, and 35.3% a conference or publication release. Of 103 trials announced as not meeting it, 31.1% (32 of 103; 95% CI, 22.4%-39.8%) had a later release and 17.5% a conference or publication release. In study 1, 22 main-set pairs without a verdict concerned trials with results posted on ClinicalTrials.gov; for 16 of the 17 with a P value, a primary outcome analysis had P < .05.

In study 2, first announcements of failures were not filed late on Fridays (none of 140 after 4 PM Eastern time on a Friday, vs 6 of 291 successes) but were more often filed after the market closed at 4 PM (27.1% vs 17.5%) and less often before it opened at 9:30 AM (61.4% vs 73.9%; eTable 5 in Supplement 1).

Discussion

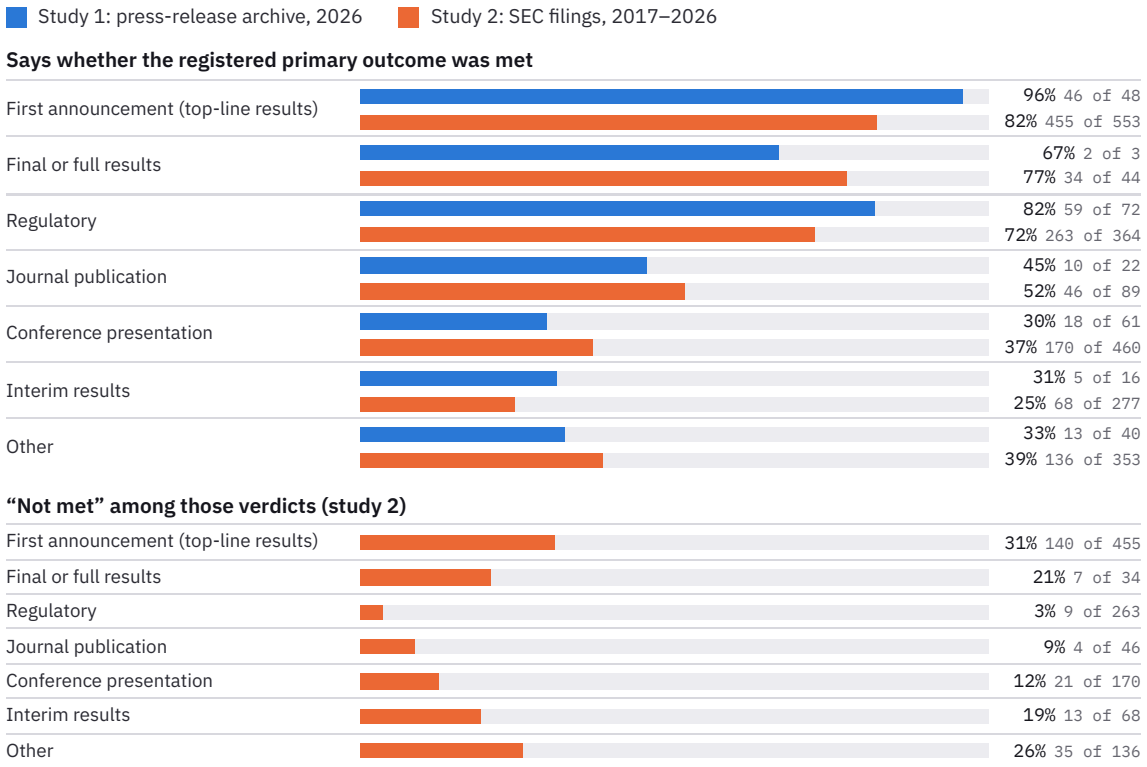


Figure. Verdict and Failure by Kind of Announcement. Share of main-set pairs whose release states whether the registered primary outcome was met (top) and, among those, the share "not met" (bottom), by kind of announcement, in study 1 (exploratory) and study 2 (descriptive; first announcements vs conference or publication tested as H2 and H3).

Kind of announcement	Study 1: verdict stated	Study 2: verdict stated	Study 2: "not met" among verdicts
First announcement (top-line)	95.8% (46 of 48)	82.3% (455 of 553)	30.8% (140 of 455)
Final or full results	66.7% (2 of 3)	77.3% (34 of 44)	20.6% (7 of 34)
Regulatory	81.9% (59 of 72)	72.3% (263 of 364)	3.4% (9 of 263)
Journal publication	45.5% (10 of 22)	51.7% (46 of 89)	8.7% (4 of 46)
Conference presentation	29.5% (18 of 61)	37.0% (170 of 460)	12.4% (21 of 170)
Interim	31.3% (5 of 16)	24.5% (68 of 277)	19.1% (13 of 68)
Other	32.5% (13 of 40)	38.5% (136 of 353)	25.7% (35 of 136)

In 2 preregistered studies, about half of the releases reporting a randomized drug trial's results said whether its registered primary outcome was met, and about 1 headline in 7 stated that outcome's result. The answer depended on when the release came. First announcements of top-line results almost always gave a verdict and carried most of the failures; releases tied to conference presentations and journal publications usually gave none and, when they did, rarely reported failure. The exploratory follow-up of individual trials points to the reason: a trial announced as a success was about twice as likely as a failed one to be the subject of another release.

The problem these releases show is omission rather than substitution. Releases rarely called an unregistered outcome primary, far less often than the outcome changes reported between registries and journal articles;[8,9] what they left out was the verdict. That omission is concentrated in releases tied to conference presentations and journal publications, which come after the first announcement. A reader who meets a trial through such a release, as journalists covering a congress or a journal article often do, sees mostly trials that worked, described without saying whether they met their primary outcome. In study 1, releases without a verdict on trials with posted results mostly concerned primary outcomes with $P < .05$, so silence did not hide failure there. The bias lies less in what later releases say than in which trials get them.

The evidence that headlines hide failure is weak. In study 1, the difference was not significant; in study 2, it was significant overall, came mostly from foreign issuers, and was not significant among domestic filers or phase 3 trials. Study 1's finding that small sponsors give a verdict less often did not hold within phase 3 trials, consistent with confounding by trial phase or kind of announcement.

Spin was common in pairs without a "met" verdict (80% and 74%), but these estimates are not comparable with earlier ones, which compared releases or abstracts with the full article and defined spin differently,[5,7] and study 1's coders disagreed on it.

Limitations

The frames cover releases that name a ClinicalTrials.gov identifier, not trials named only by acronym or in another registry, and English-language releases only. Study 1 covers 3 months of 1 archive. Study 2 covers companies registered with the SEC and the releases they chose to file; large companies file

only news they consider material, so their releases are underrepresented. The filing date may differ from the release date. Registry records were taken from monthly (study 1) and quarterly (study 2) snapshots, so a record changed and changed back within a quarter was missed. A release's verdict was coded as the release stated it; we did not check it against the trial's data, except in the exploratory analysis of posted results. Coding was done by language models from one family, so their agreement is an upper bound on reliability, and study 2's measurement coder is the smallest of them. 22 long filings could not be coded. The analyses of what followed a first announcement, of posted results and of filing time were not prespecified as tests. The author's company operates a press-release aggregator.

Conclusions

Drug-trial press releases report a trial's failure mainly at the first announcement. Later releases, including those announcing conference presentations and journal publications, mostly concern trials that succeeded and usually do not say whether the primary outcome was met. Readers of those releases, and journalists who report on them, should look for the first announcement or the registry record.

Article Information

Author Contributions: Émile Bolduc had full access to all of the data and takes responsibility for the integrity of the data and the accuracy of the analysis; conceived and designed the studies; acquired and analyzed the data; and drafted and revised the manuscript.

Conflict of Interest Disclosures: Émile Bolduc is the owner of Bollé Communications inc., which operates PPN World, a press-release aggregator whose archive is study 1's source, and Hastewire, which sells AI writing tools. No pharmaceutical, biotechnology or press-release distribution company funded or reviewed this work.

Funding/Support: Bollé Communications inc., which operates PPN World and PPN Source.

Role of the Funder/Sponsor: The funder is the author's company; see Conflict of Interest Disclosures.

Data Sharing Statement: The registrations, deviations, codebook and Supplement 1 are published with the paper at ppnworld.com/research/clinical-trial-press-releases. Release identifiers, the registry records used, all codings with their quoted evidence, prompts and code will be deposited on Zenodo with the paper's DOI.

Use of AI: Release coding used Claude Sonnet 5.5, Claude Opus 5.5 and Claude Haiku 5.5 (Anthropic) as described in Methods, at a cost of $9.54 (study 1) and $11.84 (study 2) at batch prices. The author used Claude Opus 5.5 through Claude Code to write and run the analysis code and to draft the manuscript, and checked every statement against the registered outputs. The tokens used in that work were not counted apart from the author's other work.

References

1. Purohit U, Sacks CA, Raissi A, et al. Press releases of drug-related randomized trial results prior to publication in high-impact journals: an observational study. *J Gen Intern Med*. 2023;38(14):3107-3114. doi:10.1007/s11606-023-08313-1
2. Rothenstein JM, Tomlinson G, Tannock IF, Detsky AS. Company stock prices before and after public announcements related to oncology drugs. *J Natl Cancer Inst*. 2011;103(20):1507-1512. doi:10.1093/jnci/djr338
3. Sumner P, Vivian-Griffiths S, Boivin J, et al. The association between exaggeration in health related science news and academic press releases: retrospective observational study. *BMJ*. 2014;349:g7015. doi:10.1136/bmj.g7015
4. Schwartz LM, Woloshin S, Andrews A, Stukel TA. Influence of medical journal press releases on the quality of associated newspaper coverage: retrospective cohort study. *BMJ*. 2012;344:d8164. doi:10.1136/bmj.d8164
5. Boutron I, Dutton S, Ravaud P, Altman DG. Reporting and interpretation of randomized controlled trials with statistically nonsignificant results for primary outcomes. *JAMA*. 2010;303(20):2058-2064. doi:10.1001/jama.2010.651
6. Chiu K, Grundy Q, Bero L. 'Spin' in published biomedical literature: a methodological systematic review. *PLoS Biol*. 2017;15(9):e2002173. doi:10.1371/journal.pbio.2002173
7. Yavchitz A, Boutron I, Bafeta A, et al. Misrepresentation of randomized controlled trials in press releases and news coverage: a cohort study. *PLoS Med*. 2012;9(9):e1001308. doi:10.1371/journal.pmed.1001308
8. Dwan K, Altman DG, Arnaiz JA, et al. Systematic review of the empirical evidence of study publication bias and outcome reporting bias. *PLoS One*. 2008;3(8):e3081. doi:10.1371/journal.pone.0003081
9. Goldacre B, Drysdale H, Dale A, et al. COMPare: a prospective cohort study correcting and monitoring 58 misreported trials in real time. *Trials*. 2019;20:118. doi:10.1186/s13063-019-3173-2
10. De Angelis C, Drazen JM, Frizelle FA, et al. Clinical trial registration: a statement from the International Committee of Medical Journal Editors. *N Engl J Med*. 2004;351(12):1250-1251. doi:10.1056/NEJMe048225
11. von Elm E, Altman DG, Egger M, et al. The Strengthening the Reporting of Observational Studies in Epidemiology (STROBE) statement: guidelines for reporting observational studies. *PLoS Med*. 2007;4(10):e296. doi:10.1371/journal.pmed.0040296
12. Bolduc É. The anatomy of a press release: a pre-registered census of 246,777 English-language releases on 1,151 websites. PPN World Research Paper No. 2. Zenodo; 2026. doi:10.5281/zenodo.24005947
13. US Securities and Exchange Commission. Form 8-K: current report; Form 6-K: report of foreign private issuer. Accessed November 12, 2026. https://www.sec.gov/files/form8-k.pdf; https://www.sec.gov/files/form6-k.pdf
14. Tasneem A, Aberle L, Ananth H, et al. The database for aggregate analysis of ClinicalTrials.gov (AACT) and subsequent regrouping by clinical specialty. *PLoS One*. 2012;7(3):e33677. doi:10.1371/journal.pone.0033677
15. Gilardi F, Alizadeh M, Kubli M. ChatGPT outperforms crowd workers for text-annotation tasks. *Proc Natl Acad Sci U S A*. 2023;120(30):e2305016120. doi:10.1073/pnas.2305016120
16. Hayes AF, Krippendorff K. Answering the call for a standard reliability measure for coding data. *Commun Methods Meas*. 2007;1(1):77-89. doi:10.1080/19312450709336664