PeerReviewAI

Journal of Medical Ethics

PEER REVIEW • AI-POWERED SCIENTIFIC PAPER REVIEW ENGINE

ZERO DATA RETENTION

Generated: July 7, 2026

MANUSCRIPT UNDER REVIEW

Relational accountability in AI-driven pharmaceutical practices: an ethics approach to bias, inequity and structural harm

JOURNAL	STUDY DESIGN	GUIDELINE
Journal of Medical Ethics	Qualitative study	SRQR (Standards for Reporting Qualitative Research)

AUTHOR INSTRUCTIONS

Knowledge Base

How to Cite PeerReviewAI See Appendix

AI-Generated Review — not verified by a human reviewer. Intended to support, not replace, expert peer review. Apply independent judgment.

Manuscript Integrity Check — Clear. Scanned for hidden text, prompt injection, and AI manipulation attempts. No issues detected.

MAJOR	MINOR
5	5

NAVIGATE REVIEW 11 sections

Summary

This manuscript applies a self-described feminist ethics framework — "relational accountability," organised around three obligations (transparency, solidarity, reparative justice) — to two illustrative cases of artificial intelligence (AI) in pharmaceutical practice: the Pfizer–IBM Watson immuno-oncology collaboration and the Google DeepMind–NHS (Streams) partnership. The author argues that corporate reliance on biased algorithms

exacerbates health inequities "by design," proposes a three-step causal pathway (biased training data algorithmic discrimination structural harm), and advances three policy solutions (an international AI ethics/audit board, equity-centred data governance, and a global access fund financed by a levy on AI-derived drug profits). The paper is a normative/theoretical contribution in medical ethics; despite its own framing, it is not an empirical study and performs no primary data collection or statistical analysis. Total length (2,706 words) and abstract length (158 words) are within the *Journal of Medical Ethics* limits.

The manuscript addresses a genuinely important and timely problem and is clearly organised around its tripartite framework. However, the review identifies serious problems that undermine confidence in the paper as currently written. The most consequential is that the paper's load-bearing "empirical" claims are either unsupported by, or contradicted by, the sources cited for them — most critically the repeated assertion that "Pfizer's AI-driven pricing model charged patients in low-income regions significantly more" (cited to an oncology-AI review that contains no such content) and the claim that DeepMind's Streams app "excluded certain communities from algorithm training, resulting in lower detection accuracy" (a factually inaccurate description of a rule-based acute-kidney-injury alert tool, cited to a data-protection ruling that says nothing of the kind).

The most critical issues are therefore: (1) central factual/empirical claims not supported by the cited references; (2) a systematic pattern of citation–claim misalignment; (3) mischaracterisation of the methodology as "mixed methods" and "empirical"; (4) overstated causal and generalisation claims; and (5) one-sided engagement with a literature that has moved substantially since the paper's most recent AI-ethics citations (2021), including evidence that directly complicates the argument.

Strengths

1. **Substantively important and under-theorised problem framing.** *(Novelty & Contribution)* — The manuscript targets a real gap: most AI-ethics scholarship the paper engages does emphasise transparency/fairness metrics at the technical level, and applying a relational/care-ethics lens specifically to *pharmaceutical* commercial practice (pricing, data ownership, benefit distribution) is a reasonable and comparatively fresh angle. The pharmaceutical framing (rather than generic "clinical AI") is the paper's most distinctive move.

2. **Accurate and well-integrated use of the Obermeyer example.** *(Clinical/Practical Relevance)* — The description of proxy bias — an algorithm using healthcare cost as a proxy for need and thereby underestimating illness in Black patients — is faithful to the source (as cited by the authors: Obermeyer et al., 2019). This is the one instance where a cited empirical finding is characterised correctly and deployed to concrete effect, and it anchors the "solidarity" analysis persuasively.

3 **A consistent analytic scaffold with actionable proposals.** *(Design & Methodology)* — Applying the same three obligations (transparency, solidarity, reparative justice) to both cases gives the analysis structural clarity, and the paper does not stop at critique: it offers three specific policy instruments. Even where the proposals need stronger evidentiary grounding (see Major Issues), moving from diagnosis to concrete governance mechanisms is a strength appropriate to this journal's readership.

Literature Context

Supporting Evidence

The core contention that healthcare algorithms can encode and perpetuate racial/ethnic inequity is broadly consistent with the weight of recent evidence. A large systematic review concludes that algorithms can mitigate, perpetuate, *or* exacerbate disparities depending on design and implementation: Siddique et al., 2024. A scoping review similarly documents that AI in healthcare can exacerbate ethnic and racial disparities: Hussain et al., 2025, as does a further systematic review: Haider et al., 2026.

The manuscript's "causal pathway" concept closely parallels the "Data Disparity Pipeline" described in Davidson et al., 2026, which traces how upstream inequities propagate through datasets into clinical AI — the same direction of effect the manuscript posits.

The ethics-of-justice orientation aligns with (as cited by the authors: Chen et al., 2021) and with a same-journal qualitative study on responsibility for algorithmic bias, Aquino et al., 2025.

Contradictory Evidence

The manuscript's "by design" framing and its confidence that de-biasing interventions (diverse-data mandates, audits) straightforwardly reduce harm are complicated by evidence that popular fairness interventions can *worsen* outcomes across all groups: Coots et al., 2025. Race-neutral ("racially blind") models can produce underprediction and reduced performance for Black patients: Brown et al., 2025, and the appropriate use of race in algorithms is genuinely contested: Basu, 2023.

The premise that AI "entrenches disparities by design" as a near-universal is challenged by a demonstration that a bias-minimising ML model was both more accurate and statistically unbiased relative to established scores: Allen et al., 2020.

Evidence Synthesis

The manuscript aligns with the general direction of the evidence base — that AI can and does propagate health inequities — but it is an outlier in its uni-directional, "by design" framing and in the vintage of its evidence. The 2023–2026 literature (including a paper on a near-identical pipeline concept and a same-journal empirical study) has converged on a more heterogeneous, tradeoff-aware picture that the manuscript neither cites nor engages, which weakens both its novelty claim and the credibility of its policy prescriptions.

Major Issues CRITICAL

1 Central factual/empirical claims are unsupported by, or contradicted by, their cited sources.

Concern: Validity Threat

Location: Abstract; Introduction (para 1, ref 2); "Ethical analysis" (Solidarity, ref 2); Discussion (para 2, refs 2 and 6).

Problem: The paper's flagship pharmaceutical claim — "Pfizer's AI-driven pricing model charged patients in low-income regions significantly more than those in high-income areas, citing 'market tolerance'" — is attributed to reference 2, which is (as cited by the authors: Lotter et al., 2024) "Artificial Intelligence in Oncology: Current Landscape, Challenges, and Future Directions," a clinical-integration review whose abstract contains nothing about pricing, market tolerance, or geographic price discrimination. Separately, the Discussion states that "Google DeepMind's Streams app excluded certain communities from algorithm training, resulting in lower detection accuracy for these groups" (ref 6). Reference 6 is a data-protection ruling about failure to inform patients; it does not address training data or detection accuracy. This claim is also factually inaccurate: as the author's own reference 21 describes (as cited by the authors: Powles et al., 2017), Streams was an acute-kidney-injury alerting application and a data-governance controversy, not a trained diagnostic classifier with differential accuracy across communities.

Impact: These are the paper's two principal instantiations of the "causal pathway" in pharmaceutical practice. If they are not supportable, the empirical spine of the argument collapses, and the abstract communicates as established fact assertions that the cited evidence does not contain. This is the single most serious barrier to the paper's credibility and requires either verifiable primary sourcing or removal/reframing as hypothetical illustration.

2 Systematic citation–claim misalignment throughout the manuscript.

Concern: Incomplete Reporting

Location: Introduction (refs 2, 3, 5, 6); Methods (case sources "8 9"); Discussion (ref 5).

Problem: Beyond Issue 1, multiple claims are bound to sources that do not support them. The statement that "clinical trial recruitment algorithms often exclude marginalised populations" (ref 3) is attributed to (as cited by the authors: Clark et al., 2019), which concerns human/structural barriers to trial diversity, not recruitment *algorithms*. The claim that "public–private partnerships … demonstrate how collaborative data sharing can enhance equity" is cited to reference 5 (Popejoy & Fullerton, "Genomics is failing on diversity"), a paper documenting the *opposite* problem. The two case studies are variously cited to refs 5–6 (Introduction) and refs 8–9 (Methods) — i.e., to a genomics-diversity commentary, a privacy review, and a general "black box" monograph — rather than to the case-specific sources (refs 20, 21, 6).

Impact: A reader cannot rely on the citations to verify the argument, and the pattern suggests references were attached to claims without confirming content alignment. In a normative paper whose persuasiveness depends on accurately reported real-world exemplars, this undermines the evidentiary basis and the scholarly reliability of the whole.

3 Methodology is mischaracterised, and required methodological transparency for the claimed design is absent.

Concern: Methodological Concern

Location: Methods ("This study employs a mixed methods approach … empirical case study analysis").

Problem: The paper is a normative/theoretical analysis of secondary material; it involves no data collection, no coding, no empirical case-study protocol, and no analytic method by which inferences were derived. Labelling it "mixed methods" and "empirical" is inaccurate. For a normative paper this is entirely legitimate work — and this journal explicitly permits philosophical papers without a research checklist — but the description must match the method. As presented, key reporting elements expected of the claimed empirical/qualitative design are missing: no rationale-based case-selection strategy beyond "public documentation," no statement of researcher positionality/reflexivity, no description of how sources were interrogated, and (see Issue 4) no limitations section.

Impact: The mismatch between claimed and actual methodology inflates the epistemic status of the argument (implying empirical demonstration where only illustration exists) and prevents assessment of how conclusions were reached.

4 Overstated causal, quantitative, and generalisation claims relative to what a two-case normative analysis can support.

Concern: Interpretation Error

Location: Abstract; Methods ("establish a causal pathway"); Discussion; Conclusion.

Problem: The manuscript repeatedly states that it "establish[es] a causal pathway from biased AI design to pharmaceutical inequities." A conceptual pathway is *asserted and illustrated*, not established; the only correctly sourced instantiation (Obermeyer) is a population-health risk tool, not a pharmaceutical pricing or drug-development system. Quantitative-sounding claims — "significantly more," "blockbuster status, generating billions," "1.6 million NHS users" — are presented without magnitude, denominator, or verifiable source. The framing that AI "exacerbates inequalities by design" generalises from two contested cases to the whole domain of "AI-driven pharmaceutical practices."

Impact: Readers and any downstream policy users would take the paper to have demonstrated a causal, field-wide phenomenon. The evidence supports, at most, a plausible normative concern illustrated by examples — an important distinction for a bioethics argument that is meant to justify concrete regulation.

5 One-sided literature engagement that ignores directly relevant recent and contradictory evidence.

Concern: Validity Threat

Location: Introduction; Discussion (policy solutions); reference list (most recent AI-ethics sources 2021).

Problem: The argument for the proposed remedies (diverse-data mandates, audits, penalties) does not engage evidence that such fairness interventions carry tradeoffs and can backfire — e.g., Coots et al., 2025, Brown et al., 2025, and Basu, 2023. Nor does it engage the recent syntheses (Siddique et al., 2024; Hussain et al., 2025) or the near-identical "pipeline" formulation of Davidson et al., 2026. The framework's claim to novelty is also weakened by absence of engagement with existing relational ethics/responsibility scholarship in this journal, e.g., Aquino et al., 2025.

Impact: The analysis reads as confirmatory rather than critical, and the policy recommendations are advanced without confronting the strongest counter-evidence — precisely the engagement a *Journal of Medical Ethics* argument requires to be persuasive.

Minor Issues ADVISORY

1 **First-person plural voice in a sole-authored paper.**

Location: Throughout ("we demonstrate," "we apply," "we propose"); Contributors states IB wrote the manuscript "in its entirety."

Observation: The consistent "we" is stylistically inconsistent with single authorship; it does not affect scientific interpretation but should be reconciled.

2 **Missing limitations discussion.**

Location: Discussion/Conclusion.

Observation: No limitations are acknowledged (e.g., two-case basis, reliance on secondary/grey-literature sources, contested empirical claims). Explicit acknowledgement of the boundaries of a two-case normative analysis would materially strengthen the paper and is expected reporting for this design.

3 **Internal inconsistency in supporting statements.**

Location: Contributors ("had access to the data") vs. Data availability ("No data are available").

Observation: These statements are contradictory as written; harmonising them would remove ambiguity about what "data" the study used.

4 **MeSH keywords not provided.**

Location: Front matter.

Observation: The journal requires Medical Subject Headings descriptors as keywords; none appear in the submission.

5 **Reference formatting/metadata errors.**

Location: Reference list (see Reference Verification).

Observation: Reference 5 gives the journal as "Nature New Biol" (canonical: *Nature*); reference 10 lists only one author before "et al." (Vancouver/journal style expects up to three); reference 12 (Benjamin) appears to carry a publisher error and matched a different work on lookup.

Statistical Evaluation

1 **Study Design and Sample Size.** No statistical sample or power calculation is applicable to a normative analysis, and none is required. The relevant analogue is case selection: the argument rests on two purposively chosen cases, justified only as having "public documentation of ethical violations," yet supports domain-wide conclusions.

Flag: Moderate — Two purposively selected (and, per Major Issue 1, partly mischaracterised) cases are an inadequate basis for the generalised claims made about "AI-driven pharmaceutical practices."

2 **Statistical Test Appropriateness.** No hypothesis tests were performed, which is appropriate. However, the term "significantly" (e.g., low-income patients "charged significantly more") implies a statistical comparison that does not exist and for which no data are presented.

Flag: Minor — Quantitative significance language is used rhetorically without any underlying test or data.

3 **Effect Sizes and Confidence Intervals.** Not applicable to normative argument; none reported. The manuscript nonetheless makes magnitude claims ("billions in revenue," "1.6 million NHS users," geographic price differentials) without values, denominators, or verifiable sourcing.

Flag: Moderate — Quantitative assertions are presented without magnitude, uncertainty, or a source that supports them.

4 **P-Value Interpretation.** No p-values are reported or interpreted, consistent with the study type.

Flag: Minor — No issue beyond the informal use of "significant" noted above.

5 **Multiple Comparisons.** Not applicable; no multiple testing is performed.

Flag: Minor — No concern.

6 **Missing Data.** The data availability statement declares no data. Because the empirical claims (pricing discrimination, differential detection accuracy) are offered as fact, the absence of any underlying data or verifiable source prevents independent evaluation.

Flag: Moderate — Empirically framed claims cannot be checked against any data or supporting source.

7 **Statistical Red Flags.** The principal red flag is the presentation of unverifiable, precise-sounding empirical assertions attached to sources that do not contain them (cross-reference Major Issues 1 and 4).

Flag: Moderate — Quantitative/empirical precision is asserted without traceable substantiation.

Conclusions vs. Evidence

1 **"Corporate reliance on biased algorithms exacerbates inequalities by design" (Abstract).** *Overstated.* The evidence base is heterogeneous (Siddique et al., 2024) and includes counter-examples where bias was reduced (Allen et al., 2020). *Real-world consequence:* a policymaker citing this could justify blanket restrictions on pharmaceutical AI as inherently inequitable, when the actual determinant is design and implementation choices.

2 **"Pfizer's AI-driven pricing model charged patients in low-income regions significantly more … citing 'market tolerance'" (Ethical analysis; Discussion).** *Not Supported by Data.* The cited source (Lotter et al.) does not address pricing; no data are provided. *Real-world consequence:* a reader could repeat a specific, damaging factual allegation about a named company that the manuscript does not substantiate.

3 **"Google DeepMind's Streams app excluded certain communities from algorithm training, resulting in lower detection accuracy" (Discussion).** *Not Supported by Data.* The cited ruling concerns consent/data protection; the description misrepresents the nature of the tool. *Real-world consequence:* the case would be taught/cited as an algorithmic-bias exemplar when its documented failure was one of data governance and consent.

4 **"We establish a causal pathway from biased AI design to pharmaceutical inequities" (Methods; Discussion).** *Overstated.* A pathway is asserted and illustrated once (Obermeyer, a non-pharmaceutical case), not established for the pharmaceutical domain. *Real-world consequence:* the paper may be cited as evidence of a demonstrated mechanism rather than a proposed conceptual model.

5 **"Relational accountability … offers a novel lens" (Abstract; Conclusion).** *Partially Supported.* The pharmaceutical application is reasonably fresh, but the tripartite obligations are asserted (and partly attributed to Gilligan 1982, which does not contain them) rather than derived, and novelty is not demonstrated against existing relational-responsibility scholarship such as Aquino et al., 2025.

6 **Policy solutions "would shift the burden … fulfilling the reparative justice obligation" (Discussion).** *Overstated.* Feasibility and net benefit are asserted without engaging evidence that mandated fairness interventions can worsen outcomes (Coots et al., 2025; Brown et al., 2025).

Clinical Impact Assessment

This is a normative bioethics/policy manuscript; implications are framed in terms of scholarly discourse, regulatory positioning, and downstream policy rather than direct patient outcomes.

Decision-Making and Policy Impact

1. **Data:** The abstract and Discussion present the Pfizer pricing claim, the Streams "detection accuracy" claim (Major Issues 1 and 2), and a "by design" causal conclusion (Conclusions vs. Evidence #1, #4) as settled findings, and build three regulatory proposals on them.
2. **Real-World Scenario:** A policymaker or guideline author reading only the abstract would take away that (a) a named manufacturer uses AI to price-discriminate against low-income regions, and (b) audits, diverse-data mandates, and a profit levy are evidence-based remedies. Acting on (a) risks propagating an unsubstantiated allegation; acting on (b) risks implementing interventions that the uncited contemporary evidence (Coots et al., 2025; Basu, 2023) shows can backfire. The evidence, as evaluated in Major Issues 1, 4, and 5, is insufficient to support these actions as written.

Generalisability Threat

1. **Data:** Conclusions about "AI-driven pharmaceutical practices" broadly rest on two cases, one of which (Streams) is mischaracterised (Major Issue 1) and one of which (Pfizer–IBM Watson) is supported chiefly by a corporate press release (ref 20).
2. **Real-World Scenario:** If cited in a standards or governance document as evidence of a field-wide mechanism, the paper would over-extend a two-exemplar normative illustration into a general empirical warrant — the same over-extrapolation flagged in Conclusions vs. Evidence #4.

Omission of Countervailing Evidence

1. **Data:** The paper presents fairness interventions as unambiguously corrective (Major Issue 5) and omits the 2023–2026 evidence on tradeoffs and potential harms.
2. **Real-World Scenario:** A committee adopting the recommended "diverse training data with penalties" without the tradeoff literature could mandate changes that reduce calibration/sensitivity for the very groups the policy intends to protect (Brown et al., 2025).

3 **Bottom Line:** Based on the implications identified above, policymakers or program designers should be aware that this manuscript's central factual exemplars are not substantiated by the sources cited for them, its causal and "by design" conclusions are asserted rather than demonstrated, and its remedies are advanced without engaging contemporary evidence that such interventions can carry significant tradeoffs — so its specific empirical claims and regulatory recommendations should not be acted upon or cited as established until the sourcing, methodological framing, and literature engagement are corrected.

Reference Verification

Citation Accuracy

Ref #	Reference (First Author, Year)	Status	Matched PMID	Notes
1	Beauchamp, 2019	Unverified	—	No matching record in PubMed or Crossref.
2	Lotter, 2024	Verified	38597966	Issue number omitted (PubMed shows issue 5). Content does not support the manuscript's pricing claim (see Major Issue 1).
3	Clark, 2019	Verified	30545650	Issue number omitted (PubMed shows issue 5). Concerns barriers to trial diversity, not "recruitment algorithms."
4	Held, 2006	Unverified	—	No matching record in PubMed or Crossref.
5	Popejoy, 2016	Verified	27734877	Journal cited as "Nature New Biol"; PubMed shows *Nature*. Issue omitted.
6	Information Commissioner's Office, 2017	Unverified	—	No matching record in PubMed or Crossref.
7	Gilligan, 1982	Unverified	—	No matching record in PubMed or Crossref.
8	Price, 2019	Verified	30617331	Issue number omitted (PubMed shows issue 1).
9	Pasquale, 2015	Unverified	—	No matching record in PubMed or Crossref.
10	Mehrabi, 2021	Unverified	—	No matching record in PubMed or Crossref.
11	Adamson, 2018	Verified	30073260	Issue number omitted (PubMed shows issue 11).
12	Benjamin, 2019	Possible Match	10.3917/res.229.0255	Crossref matched a different work (Demichelis, 2021, *Réseaux*) — a review of the book, not the book; author/year/journal differ. Manual confirmation recommended.
13	Access to Medicine Foundation, 2021	Unverified	—	No matching record in PubMed or Crossref.
14	Xiao, 2022	Verified	35435948	Issue number omitted (PubMed shows issue 6).
15	Genomics England, 2019	Unverified	—	No matching record in PubMed or Crossref.
16	The Global Fund, 2021	Unverified	—	No matching record in PubMed or Crossref.
17	Obermeyer, 2019	Verified	31649194	Issue number omitted (PubMed shows issue 6464).
18	Chen, 2021	Verified	34396058	All citation fields verified.

Ref #	Reference (First Author, Year)	Status	Matched PMID	Notes
19	Mittelstadt, 2019	Unverified	—	No matching record in PubMed or Crossref.
20	Pfizer Inc, 2019	Unverified	—	No matching record in PubMed or Crossref.
21	Powles, 2017	Possible Match	29308344	Journal cited as "Health Technol"; PubMed shows *Health Technol (Berl)*. Issue omitted.
22	Bellamy, 2019	Unverified	—	No matching record in PubMed or Crossref.

Metadata & Formatting Discrepancies

1. **Ref 5 (Popejoy):** journal given as "Nature New Biol"; canonical journal is *Nature*.
2. **Ref 21 (Powles):** journal given as "Health Technol"; canonical is *Health Technol (Berl)*.
3. **Ref 12 (Benjamin):** database lookup matched a different work (a book review); the cited publisher/entry should be manually verified against the actual monograph.
4. **Ref 10 (Mehrabi):** only one author listed before "et al."; journal style expects up to three named authors.
5. **Refs 2, 3, 5, 8, 11, 14, 17:** issue numbers omitted relative to the canonical record (minor for volume+page Vancouver style).

Self-Citation Analysis

No references are authored by the manuscript author (Irfan Biswas). 0 of 22 total references (0%) are self-citations. The level of self-citation is appropriate and typical for this field.

Key Missing References

Aquino et al., 2025 — Introduction/Discussion; a same-journal empirical study on responsibility for algorithmic bias that directly bears on the paper's accountability framework and novelty claim.

Siddique et al., 2024 — Introduction/Ethical analysis; a systematic review that would ground (and appropriately qualify) the empirical claim that algorithms drive disparities.

Davidson et al., 2026 — Methods/Discussion; presents a near-identical "Data Disparity Pipeline," essential for positioning the manuscript's causal-pathway contribution.

Coots et al., 2025 — Discussion (policy solutions); provides the counter-evidence that fairness interventions can worsen outcomes, which the proposals must address.

Reporting Guideline Compliance

Guideline applied: SRQR (Standards for Reporting Qualitative Research). Note that this journal does not require a reporting checklist for non-empirical/philosophical papers; SRQR is applied here because the manuscript itself claims an empirical, mixed-methods design.

11 of 21 items adequately reported, 6 incomplete, 4 missing (item S11 not applicable).

- **Item S5 (Qualitative approach and paradigm)** — Names "feminist ethics/mixed methods" but does not identify or justify a recognised qualitative approach/paradigm, and conflates normative with empirical inquiry.
- **Item S6 (Researcher characteristics and reflexivity)** — No positionality/reflexivity statement provided.
- **Item S7 (Context)** — Contextual detail for the two cases is minimal.
- **Item S8 (Sampling strategy)** — Case selection justified only as "public documentation of ethical violations," without a rationale-based strategy.
- **Item S10 (Data collection methods)** — Sources described generically ("corporate reports, peer-reviewed studies, regulatory filings") without a systematic procedure.
- **Item S12 (Units of study)** — The two cases are not characterised with sufficient descriptive detail.
- **Item S13 (Data processing)** — No description of how source material was processed for analysis.
- **Item S14 (Data analysis)** — No analytic process by which inferences/themes were derived is described.
- **Item S15 (Techniques to enhance trustworthiness)** — No triangulation, audit trail, or equivalent described.
- **Item S17 (Links to empirical data)** — Claims rest on secondary assertions rather than substantiated primary excerpts (compounded by Major Issues 1–2).
- **Item S19 (Limitations)** — No limitations discussion is present.

A full per-item compliance table is available in Author Review.

Review Scope & Limitations

> I have **high confidence** in the assessment of citation–claim alignment and reference accuracy: these were checked directly against the provided abstracts and an automated PubMed/Crossref verification, and the mismatches identified (Major Issues 1–2) are demonstrable from the source texts supplied. I have **high confidence** in the reporting-standard assessment against SRQR and in the observation that the manuscript's empirical framing does not match its actual (normative) method. My evaluation of the literature context and contradictory evidence draws on the supplied PubMed set and is **moderately-to-highly confident**, with the caveat that I confined citation to that provided corpus.

Confidence is **moderate** on two points that depend on information not available to me. First, several references (books and grey literature: Beauchamp & Childress, Held, Gilligan, Pasquale, and institutional reports) are legitimately difficult to index and are plausibly real works even though automated lookup did not confirm them; their "unverified" status reflects database limits, not an allegation. Second, judgments about the correctness of the framework's philosophical lineage (e.g., attribution of the tripartite obligations) reflect my reading of the argument as presented rather than adjudication of contested bioethics scholarship, which lies partly outside a single reviewer's expertise; a reviewer with a primary appointment in feminist bioethics would add value on that specific point.

Throughout, I applied standard reporting and integrity frameworks for qualitative/normative work (SRQR), the ICMJE recommendations on authorship, disclosures, and data-availability statements, and general standards for citation integrity and proportionate claims. The core empirical-sourcing problems, however, are severe enough that they would need resolution before the philosophical merits could be fairly judged.

Appendix

Citation Guide, Reference Formats & AI Disclosure

How to Cite PeerReviewAI

Most journals require disclosure of AI tool use in manuscript preparation. AI tools cannot be listed as authors. Authors bear full responsibility for the accuracy and integrity of all submitted work.

ACKNOWLEDGMENTS SECTION DISCLOSURE

PeerReviewAI (PeerReviewAI LLC; https://peerreviewai.org), an AI-powered manuscript review tool, was used to generate a structured pre-submission review of this manuscript to identify potential methodological, statistical, reporting, language, and reference issues. All revisions were made by the authors, who take full responsibility for the content and integrity of this work. PeerReviewAI was not used to write or generate the manuscript; any edits it suggested (e.g., language and clarity revisions provided as tracked changes) were reviewed and individually accepted or rejected by the authors. doi:10.5281/zenodo.21014264

Recommended by: Elsevier, Springer Nature, Wiley, ACS, IEEE, most biomedical journals

METHODS SECTION DISCLOSURE

Prior to submission, the manuscript was evaluated using PeerReviewAI (PeerReviewAI LLC; https://peerreviewai.org), an AI-powered manuscript review tool that assesses methodology, statistical analysis, reporting-guideline compliance (including CONSORT, STROBE, PRISMA, and others), reference integrity, ethical considerations, and formatting against the target journal's requirements, and that provides language and clarity edits as tracked changes. The authors reviewed all AI-generated feedback and suggested edits and independently decided which to accept; PeerReviewAI was not used to write or generate the manuscript text. doi:10.5281/zenodo.21014264

Recommended by: PLOS, Science, APA journals, journals requiring AI use in Methods

COVER LETTER DISCLOSURE

We wish to disclose that PeerReviewAI (PeerReviewAI LLC; https://peerreviewai.org), an AI-assisted manuscript review tool, was used during the preparation of this manuscript to perform a pre-submission review identifying potential issues with methodology, statistical reporting, language, and compliance with the target journal's author instructions. Any edits the tool suggested (such as language and clarity revisions provided as tracked changes) were reviewed and accepted or rejected by the authors, who take full responsibility for the manuscript. PeerReviewAI was not used to author or generate the manuscript. doi:10.5281/zenodo.21014264

Recommended by: Science, ICMJE-compliant journals, journals requiring cover letter AI disclosure

FORMAL REFERENCE CITATIONS

Vancouver / ICMJE

PeerReviewAI [software]. PeerReviewAI LLC; 2026. Accessed July 7, 2026. doi:10.5281/zenodo.21014264. Available from: https://peerreviewai.org

APA 7th Edition

PeerReviewAI LLC. (2026). PeerReviewAI [Computer software]. PeerReviewAI LLC. https://doi.org/10.5281/zenodo.21014264

AMA

PeerReviewAI [computer program]. PeerReviewAI LLC; 2026. Accessed July 7, 2026. https://doi.org/10.5281/zenodo.21014264

BibTeX

```
@software{peerreviewai2026,
  title  = {PeerReviewAI},
  author = {{PeerReviewAI LLC}},
  year   = {2026},
  doi    = {10.5281/zenodo.21014264},
  url    = {https://peerreviewai.org}
}
```

Literature Sources

All literature citations in this review were retrieved from the PubMed database (40 million records) and verified against real publication records. Each PMID link resolves to the original paper on PubMed. No citations were generated from AI training data.