

Peer Review

Review of: "Visible but Unverified: How Public Signals Diverge from Formal Records in Israeli Digital Health Startups"

Evgenii Arsentev¹

1. Independent researcher

1. What the work claims

The paper asks a simple question that I have not seen put this way before. When a digital health startup is publicly tied to named people and institutions, does that tie tell us the same thing about two different kinds of evidence? One kind is news: retrieved articles that describe clinical use, hospital pilots, health-sector deployment, or regulatory activity. The other is formal records, meaning an accepted match in ClinicalTrials.gov, in openFDA device records, or in OpenAlex (Section 3). The author takes 526 Israeli digital health startups from Startup Nation Finder, and the 330 with at least one retrieved article become the main sample (Table 1). A firm counts as connected when it shares a named person or institution with another firm in the set; 168 firms do, with 402 edges between them (Figure 1). The same logit model with the same controls is then run once for each outcome (Section 4, Equation 1).

The answer is a split. Connected firms have an odds ratio of 2.16 [1.24, 3.77] for medical news evidence and only 1.12 [0.60, 2.08] for formal records. Formal records follow company age (2.41) and funding (1.87) instead, all in Table 3. In the full sample of 526, the two connection estimates are 2.27 and 1.20 (Table 4). From this, the author concludes that public connections measure visibility more than they measure formal innovation records (Abstract).

2. Whether the conclusion holds on the data

I will start with what I liked, because there is a lot of it. Both outcomes are measured on the same 330 firms with the same controls, so the two columns of Table 3 can be read side by side, and that is the right design for this question. The counts also add up. I recomputed Figure 2 from Table 1: there are $330 - 146 = 184$ firms without a connection, 56.5% of 184 is 104, 80.1% of 146 is 117, and $104 + 117 = 221$, which is the

medical-news count in Table 1. For formal records, 29.9% of 184 is 55 and 35.6% of 146 is 52, and together they give the 107 in the same table. And I respect that the author writes that Hypothesis 2 stays "unresolved in these data" (Section 5.1) and does not claim a null effect.

Now the main thing. The headline is a divergence between two outcomes, but I could not find a test that the two PublicConnection coefficients differ from each other. In Table 3, the intervals are [1.24, 3.77] and [0.60, 2.08], and they overlap between 1.24 and 2.08. In Table A1, the two estimates are 14.1 and 2.0 percentage points, each with a standard error of 5.3. One coefficient being significant and the other not is not yet evidence that they are different. A joint model for both outcomes with a test of equality would settle it.

My second concern is about where the two variables come from. PublicConnection is built from people and institutions "observed in public sources" (Section 3), and MedicalNewsEvidence is read from retrieved public articles. The paper does not say whether the entities come from company profiles, from the same articles, or from both. If the same articles feed both variables, then part of the 2.16 is one text counted twice. Table A3 hints at this: shared institutions give 1.91 [1.06, 3.44], while shared persons give 1.36 [0.59, 3.15]. ArticleCount is controlled, yes, but it measures volume and not overlap of sources. To be fair, Table A4 speaks against a purely mechanical link, because PublicConnection does not predict partnership news (0.88 [0.44, 1.77]) or financing news (1.30 [0.74, 2.27]). I would bring that result into the main text.

Having more connections does not add much. Section 5.2 calls Table 5 "a consistent difference across outcomes," but for medical news, NeighborDegree is 1.19 [0.91, 1.56] and WeightedDegree is 1.27 [0.95, 1.69], and both intervals include one. Only ConnectionPageRank, 1.41 [1.06, 1.88], does not. So the signal sits in having any connection at all.

The wording also goes further than the design. Hypotheses 1 and 2 say connections "increase the probability," and the Abstract speaks of "the effect of public connections." The data are one cross-section frozen on June 26, 2026 (Section 3), with no event dates. An older, funded firm may simply attract both press and partners. What the data do support is an association between a shared public connection and medical news, and no detectable association with formal records at this sample size.

3. Reproducibility

A supplementary archive is attached on the Qeios page (visible-but-unverified-supplementary-data-v3-20260812.zip). The Data Availability statement says the replication materials hold the analysis code,

derived tables and figures, model outputs, and documentation. The company-level dataset itself is not public because of source-access and redistribution constraints and is available from the author for research review. I should be honest here: I read the article as rendered on Qeios and did not audit the archive, so some of what I ask for below may already be inside it. Dates are given, which I appreciate. The database was frozen on June 26, 2026, and the articles were collected through June 13, 2026 (Section 3). The weighted-degree construction is described with its constants, a factor of 0.9 for institutions, and a cap of 2.5 on firm-entity weights.

Several things I looked for in the text and did not find. How are articles coded into MedicalNewsEvidence: by hand, by keywords, or by a language model? There is no codebook and no agreement check on a hand-labelled subsample. How were the articles retrieved, from which source, with which queries, and in which languages? ArticleCount (mean 3.63, median 2.00, Table 2) and the $526 - 330 = 196$ firms with no article at all depend on this step. How were named people and institutions extracted and merged, and what does "source confidence" in the firm-entity weights mean? FormalRecord matching uses standardized names and "a limited set of company aliases" (Section 3), but no check of matching accuracy on a sample is reported. The statistical software is not named either.

One practical suggestion: ClinicalTrials.gov, openFDA, and OpenAlex are open sources, so the firm-level outcome indicators and an anonymized firm-to-firm edge list could be released even if the profiles and the articles cannot.

4. What specifically to fix

1. Table 1 against Section 3. Table 1 counts 221 firms with medical news evidence among the 330 firms with articles, and 228 among all 526. So $228 - 221 = 7$ firms have medical news evidence and no retrieved article, while Section 3 defines the variable from "a firm's retrieved articles". Please explain where the evidence comes from for these seven, or correct the counts. It matters for Table 4, where AnyArticle enters the medical-news model with 8.64 [2.94, 25.38].
2. Add a test that the PublicConnection coefficient differs between the two models of Table 3, and report it under the table.
3. Section 3: describe how MedicalNewsEvidence is coded and check it on a hand-labelled subsample. I would do the same for FormalRecord name matching.
4. Section 3: say which sources the named people and institutions come from, and add to Table A3 a connection indicator built without the retrieved articles. Table A4 already helps here, since

PublicConnection does not predict partnership news (0.88 [0.44, 1.77]) or financing news (1.30 [0.74, 2.27]). That deserves a sentence in the main text. It also sits a little uneasily with the argument in Section 2 that named counterparties publicize joint activity, so I would discuss why partnership news does not move.

5. Section 4, last sentence, cites "Tables A4 and A1" for the financing-news and partnership-news outcomes. Table A1 holds the linear probability estimates, so I think Figure A1 is meant.
6. The Abstract says "Among clinically oriented firms with retrieved news," but Section 3 defines no clinical-orientation filter; the sample is firms "whose primary products serve human health care." In the same Abstract, "the effect of public connections" should become association wording.
7. FormalRecord pools three different things: a trial registration, an FDA device record, and any research-indexed work in OpenAlex (Section 3). A paper with a company affiliation is a much lower bar than a 510(k) clearance. I would report the three components separately, at least as counts, and rerun Table 3 without OpenAlex.
8. A short limitations paragraph would help; the paper has none now. The trial registry and the device database are US ones, so Israeli and European device registrations are not covered. CompanyAge has a mean of 10.54 and a median of 9.00 years (Table 2), so "young" for these firms in the Abstract and Section 6 needs a qualifier. Funding is undisclosed for a 0.28 share of firms (MissingFunding, Table 2). And it is not stated whether closed firms stay in Startup Nation Finder.

5. Recommendation

Minor revisions. I liked the question, and the counts in Table 1 agree with Figure 2, which I was happy to see. I would add a test comparing the two PublicConnection coefficients (Table 3), and a short description of how the news outcome is coded (Section 3). It would also help to use association wording in the Abstract. None of this needs new data collection, so I think the changes are small. I would be glad to read the revised version.

Declarations

Potential competing interests: No potential competing interests to declare.