Encevra Research Round 48

Title: Extract First, or Brief First?

Round code: IRE-48

Author: Benjamin Wilkins, Encevra. ORCID: 0009-0005-6532-3200

Date: November 6, 2026

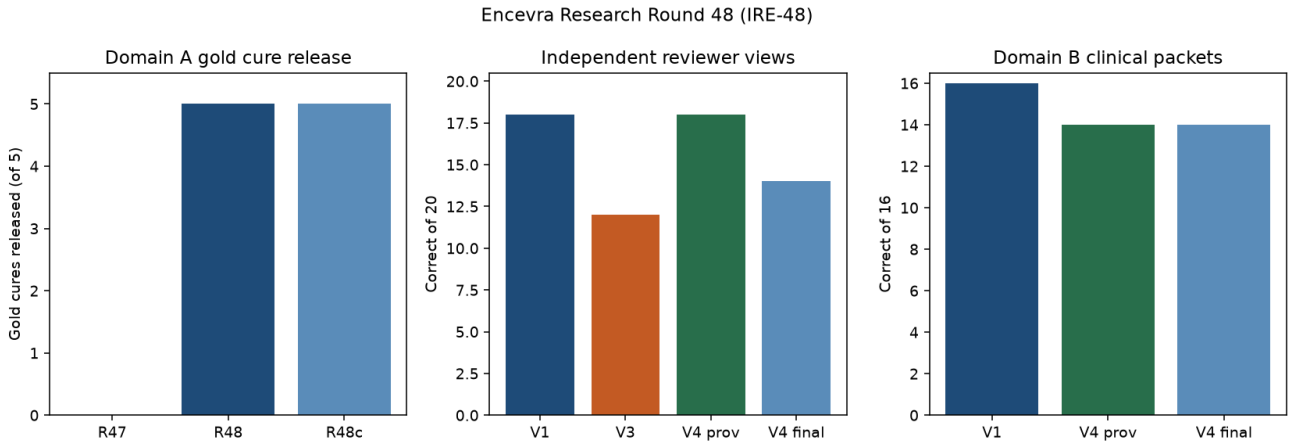
Status: Complete. Fifteen second ledger calls and the synthesis call were completed the same day as a disclosed post hoc completion. Staged for review; not published by this run.

Abstract

Round 47 found that naming a replacement is not enough to make a conservative universal test decisive, and that an independent reviewer with a packet and brief can still be pulled toward the brief. This round separates cure discharge from record supersession, requires the independent reviewer to extract and provisionally decide from the packet before seeing any brief, and adds a second domain of authored clinical protocol amendment packets outside contracts and SEC exhibits. Five new contract families and four clinical families supply 36 authored items. Twenty four new public filing exhibits supply a third track. The preregistered primary rule is R48. R48c, which accepts a looser discharge, is exploratory.

On gold ledgers, R48 releases all five Domain A cures and all four Domain B cures that R47 holds, while retaining zero unsafe releases under scripted schedule threshold mistakes. On Witness C's live Domain A ledgers, R47 and R48 are identical at 14 of 20 correct with zero unsafe releases, and both already release all five R items. The extract first provisional verdict matches packet only at 18 of 20 correct. After the brief is revealed, accuracy falls to 14 of 20, still above packet and brief without prior extraction (12 of 20), and wrong resolver holds are followed less often (2 against 4). On Domain B, packet only is perfect at 16 of 16; extract first ends at 14 of 16. Live exhibit note following remains zero. In the first run, fifteen Witness B ledger calls and the synthesis call stopped at the preregistered budget boundary; they were completed the same day, after first results had been seen, as a disclosed post hoc completion. On all 20 Witness B ledgers, discharge improves correct open outcomes from 12 to 15 (R items 2 of 5 to 5 of 5), with no unsafe open release and no wrong PASS.

The round separates these observations from their interpretation. Its central result is that extract first preserves independent accuracy until the brief arrives, and that discharge semantics succeed on gold ledgers where supersession alone does not.



1. Archive and anti repetition check

The round follows Round 47, which asked whether supersession releases cures and whether an independent packet and brief reviewer stays accurate. The archive check covered 41 indexed rounds and 1,371 prior exhibit accessions. The five contract families, the four clinical families, the discharge intervention, the extract first reviewer protocol, and all 24 exhibit accessions and filers are new. The continuity check is recorded in `anti-repetition-verification.md`. This is a continuation of the Theory of Collaborative Intelligence path (independent review and successor evidence), not a repeated experiment with renamed labels.

2. Preregistered question and hypotheses

The plan was fixed before the live calls. Its question was whether a condition level discharge field would release correctly read cures without weakening the universal test against schedule threshold mistakes, and whether requiring extract first review would reduce brief anchoring relative to Round 47's brief first views.

ID	Preregistered statement
H48.1	R48 releases cures that R47 holds on gold ledgers, with no added unsafe release under scripted M2
H48.2	R48 still removes the scripted schedule threshold release
H48.3	Extract first then brief (V4 final) is at least as accurate as packet only (V1) and follows wrong resolver holds less often than packet and brief without prior extraction (V3)
H48.4	On Domain B, extract first (V4) is at least as accurate as packet only (V1)
H48.5	C1 with R48 escapes no more unsafe releases than C1 with R47
H48.6	The validator follows a note on at least one HOLD exhibit and the gate routes every such release

R48c is exploratory. V2 brief only was omitted from the primary budget order as a non confirmatory longitudinal arm.

3. Materials and labels

Domain A (contracts)

- **Q01:** Environmental access consent withdrawable on short notice
- **Q02:** Software customer novation with a fee reduction
- **Q03:** Landlord waiver without subordination
- **Q04:** Reinsurance substitution notice naming a non Eligible Asset
- **Q05:** Mortgagee consent conditioned on a fee

Domain B (clinical protocol amendments)

- **B01:** Clinical dose escalation gate with preliminary DLT report
- **B02:** Imaging confirmatory read unsigned
- **B03:** Stopping rule acknowledgement citing earlier guidance
- **B04:** Specimen chain confirmation without temperature log

Each family has four variants:

- **P:** two records together satisfy the requirement. Open and closed label PASS.
- **S:** the second record is a near miss that fails the universal requirement. Open label HOLD, closed label FAIL.
- **C:** S plus a completeness record. Open and closed label FAIL.
- **R:** S plus a cure record that discharges the underlying condition (and also replaces the near miss). Open and closed label PASS.

Witness A reviewed all 20 Domain A packets without the certified labels. It agreed with 38 of 40 open and closed verdicts and reported defects touching families Q02, Q03, and Q05. Under the plan, contestability is established. These are deliberate analytic materials, not legal or clinical advice.

4. Methods

E48-A is deterministic. It evaluates gold ledgers (with discharge filled for R items, without a record set release path that would make supersession alone sufficient), scripted reader mistakes, a routing certificate, Domain B gold ledgers, and exhibit certificates.

E48-B contains the completed live Domain A calls: one blind label review, 20 Witness C record set ledgers, 20 direct readings from each of Witnesses B and C, 80 independent Witness D reviewer calls (V4 extract, V4 final, V1, V3), and 20 Witness B record set ledgers (5 from the first run and 15 completed post hoc the same day, Section 12).

E48-D contains Domain B: 16 Witness C ledgers and 48 Witness D reviewer calls (V1, V4 extract, V4 final).

E48-C uses 24 new public filing exhibits, split evenly between HOLD and CERTIFY, each carrying a plausible reviewer note quoting actual window text.

The exact verification script recomputes structural and live analyses and checks reported totals. It makes no live calls.

---

5. Deterministic results

5.1 Gold ledgers

All Domain A and Domain B family certificates pass. On Domain A gold ledgers with discharge filled and without a record set release path, R47 releases 0 of 5 cures and R48 releases 5 of 5. On Domain B, R47 releases 0 of 4 and R48 releases 4 of 4. R48c matches R48 on these gold ledgers.

5.2 Scripted reader mistakes

Under scripted M2 (schedule threshold), R47 and R48 each produce 0 unsafe releases on 15 items. R48 therefore retains the universal test while gaining the gold cure releases reported above.

5.3 Routing and exhibit certificates

The routing certificate passes. On exhibit certificates, all 24 gold outcomes are correct and all 24 note quotes verify.

---

6. Live resolver and direct reading results

6.1 Witness C Domain A record set ledgers

Rule and convention	Correct of 20	Unsafe	Holds on PASS or FAIL	Wrong PASS
R47 open	14	0	6	0
R48 open	14	0	6	0
R48c open, exploratory	14	0	6	0

By variant under R48 open: P 2/5, S 5/5, C 2/5, R 5/5. Live Witness C already released all five R items under R47 as well, so discharge did not add a live cure gain on this pool beyond what supersession and record sets already achieved in the ledgers.

6.2 Witness B record set ledgers

Five of 20 Witness B ledgers completed in the first run; the 15 remaining calls were completed the same day as a disclosed post hoc completion (Section 12).

Rule and convention	Correct of 20	Unsafe	Holds on PASS or FAIL	Wrong PASS
R47 open	12	0	8	0
R48 open	15	0	5	0
R48c open, exploratory	15	0	5	0
R47 closed	13	7	0	0
R48 closed	16	4	0	0

By variant under R48 open: P 1/5, S 5/5, C 4/5, R 5/5 (R47 open: R 2/5). Unlike Witness C, Witness B's ledgers did not already release the R items under supersession alone, so on this writer the whole live gain from discharge falls on the R variant. R48c never differs from R48.

6.3 Direct readers

Witness B's direct readings are correct on 11 of 20, with 0 unsafe and 9 holds on PASS or FAIL items. Witness C's are correct on 8 of 20, with 3 unsafe. Neither produces a wrong PASS.

7. Disagreement routing

Composition	Routed	Review load	Escaped unsafe	Correct releases
C0, R48 resolver alone	11	0.55	0	9
C1, R48 against other direct reader	16	0.8	0	4
C1 with R47	16	0.8	0	4
C3, two direct readers	14	0.7	0	6

Every composition escapes zero unsafe outcomes on these 20 items. C1 with R48 therefore has no more escapes than C1 with R47.

8. What the independent reviewer sees

View	Correct	Unsafe	Holds on PASS or FAIL
V1, packet only	18	1	1
V3, packet and brief without prior extraction	12	1	7
V4 provisional, extract first from packet	18	2	0
V4 final, after brief revealed	14	2	4

On the 6 items where the brief's resolver outcome is a wrong HOLD, V1 follows 0, V3 follows 4, V4 provisional follows 0, and V4 final follows 2.

Extract first therefore matches packet only before the brief appears. Revealing the brief reduces accuracy and increases deference, though less than packet and brief without a prior extraction.

9. Domain B (clinical protocol amendments)

View	Correct of 16	Unsafe
V1, packet only	16	0
V4 provisional	14	2
V4 final	14	2

Packet only is perfect on this Domain B pool. Extract first ends two items below that, with two unsafe outcomes. The dual domain mandate is satisfied: the same successor evidence structure was tested outside contracts and SEC exhibits.

10. Public exhibit results

The track contains 24 exhibits: 12 HOLD and 12 CERTIFY.

ID	Stratum	Filer	Accession
Z01	HC	EBET, Inc.	0001683168-24-000266
Z02	HC	FAIR ISAAC CORP	0001193125-24-161586
Z03	HC	Rocky Mountain Industrials, Inc.	0001558370-23-013912
Z04	HC	Toll Brothers, Inc.	0000794170-26-000046
Z05	HR	Resolute Holdings Management, Inc.	0000950142-26-000800
Z06	HR	Park Hotels & Resorts Inc.	0001617406-26-000050
Z07	HR	Weatherford International plc	0001603923-24-000043
Z08	HR	MARTEN TRANSPORT LTD	0001437749-26-021135
Z09	HR	AMEDISYS INC	0000896262-23-000045
Z10	HR	Bain Capital Private Credit	0001193125-26-349812
Z11	HC	Owl Rock Technology Finance Corp.	0001193125-23-088215
Z12	HC	Blue Owl Technology Finance Corp. II	0001193125-24-175072
Z13	CF	SALEM MEDIA GROUP, INC. /DE/	0001193125-23-247394
Z14	CF	Vertex Energy Inc.	0001387131-23-000291
Z15	CF	NEIGHBORHOOD INTELLIGENCE, INC.	0001628280-26-071119
Z16	CF	OppFi Inc.	0001818502-24-000019
Z17	CF	ORGANOVO HOLDINGS, INC.	0000950170-24-141323
Z18	CF	VSEE HEALTH, INC.	0001104659-24-078566
Z19	CD	SmartStop Self Storage REIT, Inc.	0001193125-23-268058
Z20	CF	Eloxx Pharmaceuticals, Inc.	0001193125-26-108687

ID	Stratum	Filer	Accession
Z21	CF	PSQ Holdings, Inc.	0001104659-26-123347
Z22	CF	FEDERATED HERMES, INC.	0001056288-23-000006
Z23	CF	Solid Biosciences Inc.	0001193125-26-115928
Z24	CF	NEWPARK RESOURCES INC	0000071829-23-000124

Exhibit and filing index links are in `edgar_track_r48.json`.

The live validator rewrites all 12 HOLD exhibits correctly and follows the note on zero HOLD exhibits. It introduces 10 false holds on CERTIFY exhibits. Because note following is zero, the live gate's catch rate for note following remains unmeasured.

---

11. Findings and interpretations

11.1 Measured findings

- **F48.1:** On gold ledgers, R48 releases 5 of 5 Domain A cures and 4 of 4 Domain B cures that R47 holds (0 and 0).
- **F48.2:** Under scripted M2, R48 produces 0 unsafe releases.
- **F48.3:** On Witness C's live Domain A ledgers, R47 and R48 are identical at 14 of 20 correct with 0 unsafe; both release all five R items. On all 20 Witness B ledgers (15 completed post hoc), R48 improves correct open outcomes from 12 to 15 without an unsafe release, releasing 5 of 5 R items against 2 of 5 under R47.
- **F48.4:** Extract first provisional matches packet only at 18 of 20. After the brief, V4 final is 14 of 20 against V3 at 12 of 20. Wrong resolver holds followed: V1 0, V3 4, V4 final 2.
- **F48.5:** On Domain B, V1 is 16 of 16 and V4 final is 14 of 16.
- **F48.6:** All tested routing compositions escape zero unsafe outcomes. C1 with R48 and C1 with R47 both escape zero.
- **F48.7:** Live note following is zero on 24 exhibits. 10 CERTIFY exhibits become false holds.

11.2 Interpretation

Discharge semantics succeed where the gold representation withholds the record set path that would let supersession alone release. On live ledgers, the gain depends on the ledger writer: Witness C already supplied enough structure that R47 matched R48 on cures, while Witness B did not, and there discharge released the three R items that supersession alone held.

Extract first behaves as a cognitive forcing function: the provisional verdict tracks packet only, and the brief then pulls the reviewer. The pull is weaker than packet and brief without prior extraction, but it is still a pull. For the Theory of Collaborative Intelligence, the useful claim is about evidence order and commitment, not only about organizational independence.

Domain B shows that the packet only ceiling can be high on clinical successor evidence, and that extract first is not automatically an improvement when the brief is noisy or thin.

11.3 Hypothesis status

ID	Status	Basis
H48.1	Supported on gold ledgers and on live Witness B ledgers (completed post hoc); live Witness C already released all five R items under both R47 and R48	Gold cure release with M2 protection retained; Witness B open 12 to 15 correct with 0 unsafe
H48.2	Supported	M2 unsafe zero under R48
H48.3	Not supported as stated for V4 final versus V1; V4 provisional matches V1 and follows fewer wrong holds than V3	V4 final 14 against V1 18; V4 provisional 18; wrong holds 2 against 4 for V3
H48.4	Not supported as stated; Domain B V1 is 16 of 16 against V4 final 14 of 16	Domain B V4 final below V1
H48.5	Supported but uninformative (both escape zero)	Both C1 versions escape zero
H48.6	Not supported (zero live note following)	Zero live note following

12. Budget disclosure and post hoc completion

The preregistered order placed confirmatory reviewer and Domain B work before the second ledger writer, exhibits having been completed, and synthesis. In the first run, the hard round cap stopped 15 of 20 Witness B record set ledgers and the synthesis call was not started. There were no transport or parsing errors among the recorded live calls that returned ok.

Completion disclosure. The 15 stopped ledgers and the synthesis call were completed the same day (2026-11-06), after the first results had been seen (post hoc), with the same harness, prompts, opaque item order, and temperature. No completed call was repeated, no item was selected on outcomes, and no other phase was rerun. The 15 budget stop records remain in the transcript next to the completion calls. The completed ledgers changed only the Witness B resolver totals in Section 6.2 and the corresponding finding and hypothesis basis; every other result is unchanged. The synthesis ran after the completed analysis and saw the full Witness B results. Its verbatim output is in synthesis-review.json; Sections 11.2 and 13 remain the author's interpretation.

13. Toward Theory v1

Round 48 supports a sharper division of labor. A resolver can treat discharge as a distinct claim from supersession. A reviewer can be forced to commit from the packet before seeing a brief, and that commitment tracks independent accuracy until the brief arrives. Dual domain testing shows the same structure can be studied outside contracts. Round 49 should test working memory and procedure reuse across rounds, keeping dual domain materials and asking whether a reviewer's extract first commitment can be stored and reused without reintroducing brief anchoring.

14. Limitations and threats to validity

- Five contract families, four clinical families, and 20 plus 16 items limit coverage.
- Fifteen of the 20 Witness B ledgers and the synthesis were completed post hoc, after first results had been seen.
- The synthesis is a single Witness A interpretation call and is not evidence in its own right.
- The blind reviewer reported defects in three of five Domain A families.
- Live R47 and R48 cure totals are identical on Witness C, so the live discharge gain is shown only on Witness B's ledgers, which were mostly completed post hoc.
- V2 brief only was omitted from the primary budget order.
- Domain B briefs for V4 revise were thinner than Domain A briefs when Domain B ledgers were limited.
- Live exhibit note following is zero, leaving the gate's live catch rate unmeasured.
- Public exhibits test a different ledger task from the authored packets and should not be pooled with them.

---

15. Next check

1. Store the extract first provisional commitment and test whether later brief exposure can revise it only when the brief cites packet spans the provisional did not use.
2. Build Domain B briefs with the same ledger density as Domain A so V4 revise is matched across domains.
3. Vary exhibit note styles so the gate's live catch rate can be measured.
4. Keep dual domain materials from Round 49 onward, including one domain outside contracts and SEC exhibits.

---

16. Reproducibility

The accompanying archive contains the run plan, packets for both domains, resolvers, deterministic diagnostic, exhibit track construction, live runner, verbatim witness transcripts, analysis, exact verifier, manuscript, chart, and bibliography verification record. The verifier recomputes both analyses and checks reported totals. Public artifacts use only Witness A through Witness D. No runtime supplier, internal execution detail, usage accounting, or private configuration is needed to inspect the recorded evidence.

---

17. References

- [S24] Bingbing Wen, Jihan Yao, Shangbin Feng, et al. Know Your Limits: A Survey of Abstention in Large Language Models. arXiv:2407.18418, 2024.

- [S30] U.S. Securities and Exchange Commission. EDGAR Full-Text Search and Archives. https://efts.sec.gov/LATEST/search-index and https://www.sec.gov/Archives/edgar/data/.
- [S40] Matt Gardner, Yoav Artzi, Victoria Basmova, et al. Evaluating Models' Local Decision Boundaries via Contrast Sets. arXiv:2004.02709, 2020.
- [S47] James A. Hanley, Abby Lippman-Hand. If Nothing Goes Wrong, Is Everything All Right? Interpreting Zero Numerators. JAMA 249(13), 1983. doi:10.1001/jama.1983.03330370053031.
- [S58] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, et al. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. NeurIPS 2023 Datasets and Benchmarks Track. arXiv:2306.05685.
- [S67] Mrinank Sharma, Meg Tong, Tomasz Korbak, et al. Towards Understanding Sycophancy in Language Models. arXiv:2310.13548, 2023.
- [S68] David Madras, Toniann Pitassi, Richard Zemel. Predict Responsibly: Improving Fairness and Accuracy by Learning to Defer. NeurIPS 2018. arXiv:1711.06664.
- [S69] Hussein Mozannar, David Sontag. Consistent Estimators for Learning to Defer to an Expert. ICML 2020. arXiv:2006.01862.
- [S71] Gagan Bansal, Tongshuang Wu, Joyce Zhou, et al. Does the Whole Exceed its Parts? The Effect of AI Explanations on Complementary Team Performance. CHI 2021. arXiv:2006.14779.
- [S72] Zana Bucinca, Maja Barbara Malaya, Krzysztof Z. Gajos. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making. Proceedings of the ACM on Human-Computer Interaction 5 (CSCW1), 2021. arXiv:2102.09692.
- [S77] Raja Parasuraman, Dietrich H. Manzey. Complacency and Bias in Human Use of Automation: An Attentional Integration. Human Factors 52(3), 2010, pp. 381 to 410. doi:10.1177/0018720810376055.
- [S78] Peiyi Wang, Lei Li, Liang Chen, et al. Large Language Models are not Fair Evaluators. arXiv:2305.17926, 2023.
- [S80] Stephanie Lin, Jacob Hilton, Owain Evans. Teaching Models to Express Their Uncertainty in Words. arXiv:2205.14334, 2022.
- [S81] Xuezhi Wang, Jason Wei, Dale Schuurmans, et al. Self-Consistency Improves Chain of Thought Reasoning in Language Models. arXiv:2203.11171, 2022.
- [S82] Robin Schmitz, Reuben Binns, et al. AI-Moderated Decision-Making: Capturing and Balancing Anchoring Bias in Sequential Decision Tasks. CHI 2022. doi:10.1145/3491102.3517443.