The Copy Gap

Case Studies in False Consensus, Effective Objectivity, and Independent Recoverability

J. Farrow

Draft case-study manuscript - May 2026

Abstract

Consensus is often treated as evidence. Yet agreement can arise either from independent access to a target or from repeated copying of a small number of sources. This paper applies the Copy Gap framework to three case studies: the AI/data-center nuclear-power narrative, the emerging GLP-1/alcohol-use narrative, and the historical Iraq WMD consensus. The measure compares naive functional information, computed from nominal fragments, with effective functional information, computed from independent evidence clusters or target-conditioned fragment independence. Across cases, the method distinguishes real but copy-amplified evidence from unsupported overclaiming. The AI+nuclear case shows genuine independent support for a broad infrastructure narrative but a large gap between article volume and independent deal clusters. The GLP-1/alcohol case shows a real early signal but no effective support for the stronger claim that GLP-1s are established alcohol-use-disorder treatments. The Iraq WMD case shows that the weak claim of unresolved disarmament uncertainty retained partial objectivity, while the strong claim of active deployable WMD stockpiles collapsed under effective-evidence correction and later failed against historical ground truth. The contribution is not that redundancy matters, but that the difference between apparent redundancy and effective independent recoverability can be measured, interpreted, and used as a diagnostic of copied consensus.

Keywords: copy gap; false consensus; objectivity; information theory; evidence aggregation; narrative analysis; WMD; GLP-1; nuclear energy

1. Introduction

The most dangerous consensus is not always the consensus that is false. It is the consensus that is copied. When many outlets, analysts, models, or records repeat the same source, public agreement can become numerically large while informational independence remains small. A society, institution, or market may then mistake echo for objectivity.

This paper applies a compact information-theoretic diagnostic to that problem. The proposed Copy Gap is the difference between naive redundancy and effective independent recoverability. It asks whether fragments are adequate readouts of a target and whether those fragments are independent enough to count as multiple records rather than copies of one source.

The case-study question is practical: can the measure separate (a) real distributed evidence, (b) early but over-amplified evidence, and (c) copied consensus that later fails against ground truth? To explore this, we examine three cases with different epistemic structures: an ongoing market/infrastructure narrative, an emerging biomedical narrative, and a retrospective historical consensus failure.

2. Method: Naive FI, Effective FI, and the Copy Gap

Let X be a target claim and let F1, ..., FN be fragments: articles, documents, clinical records, model outputs, reports, or sensors. A fragment subset is adequate if it recovers the target to a specified tolerance. The minimal adequate subset size is k*.

$$FI_{\text{naive}} = \max\left(0, \log_2 \frac{N}{k^*}\right)$$

$$FI_{\text{eff}} = \max\left(0, \log_2 \frac{N_{\text{eff}}}{k^*}\right)$$

$$\textit{Copy Gap} = FI_{\text{naive}} - FI_{\text{eff}}$$

N is the nominal number of supportive fragments. N_eff is the effective number of independent fragments or evidence clusters. In a full dataset, N_eff may be estimated from target-conditioned residual correlations. In these manual pilots, N_eff is approximated by coding independent evidence clusters and discounting fragments that are duplicates, summaries, or amplifications of the same source.

The measure is target-relative. It does not decide whether a target is true by itself. It measures how independently recoverable the target is from available fragments. Truth-sensitive use requires external validation, causal grounding, or later ground truth.

For cross-case comparison, we report values at k*=2 where possible. This means that at least two independent evidence clusters are required before the target is treated as adequately objectified. Sensitivity curves are provided in the underlying data files.

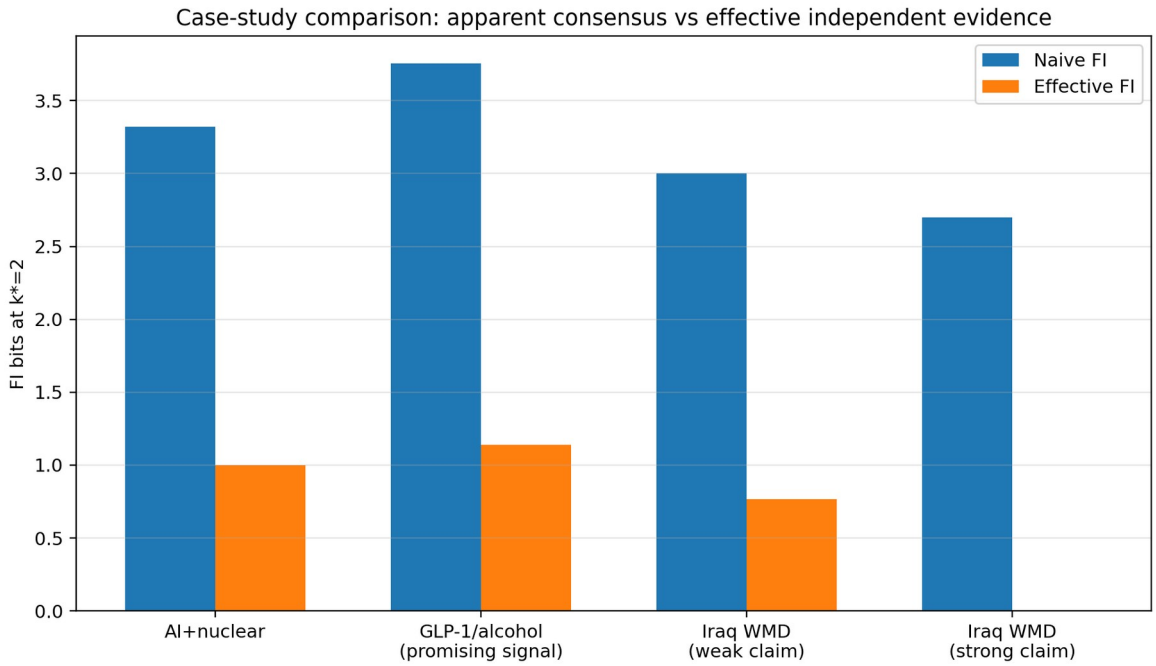

Figure 1. Across the three case studies, naive fragment counts systematically exceed effective independent evidence. The Iraq strong-stockpile claim is the clearest copied-consensus failure: naive FI is high while effective FI is zero.

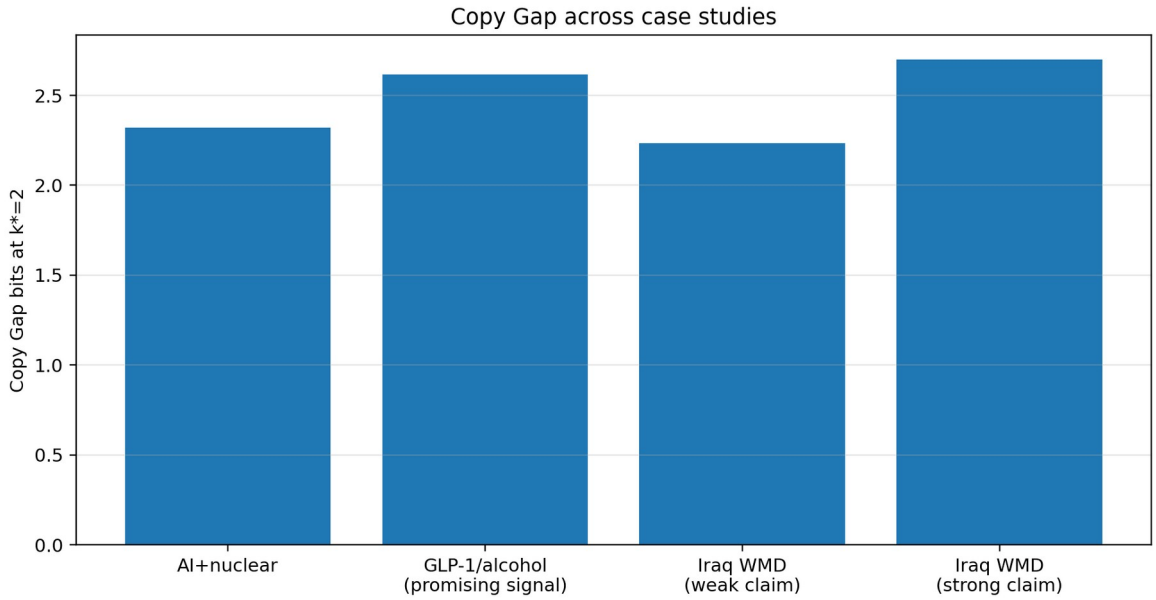

Figure 2. Copy Gap values at k=2. The Copy Gap measures how many bits of apparent consensus disappear after correcting for dependence, duplication, and weak target grounding.*

Table 1 summarizes the case-study results at k*=2.

Case	Naive FI	Effective FI	Copy Gap	Overstatement factor
AI+nuclear	3.322	1.000	2.322	5.0x
GLP-1/alcohol (promising signal)	3.755	1.138	2.617	6.1x
Iraq WMD (weak claim)	3.000	0.766	2.234	4.7x
Iraq WMD (strong claim)	2.700	0.000	2.700	6.5x

3. Case Study 1: AI Data Centers and Nuclear Power

The first case asks whether the market/media narrative that AI data centers are pushing Big Tech toward nuclear energy is independently grounded or mostly copied amplification. The target is deliberately narrow: not that all nuclear equities will rise, but that multiple Big Tech firms have made distinct nuclear-energy commitments or solicitations linked to data-center and AI power demand.

Four independent commitment clusters were identified: Microsoft/Constellation/Three Mile Island, Google/Kairos Power, Amazon/X-energy/Energy Northwest, and Meta nuclear-energy solicitation. Constellation described a Microsoft-linked agreement to restart the Crane Clean Energy Center, and Google announced an agreement with Kairos Power for up to 500 MW of 24/7 carbon-free nuclear power. Amazon announced small modular reactor agreements and X-energy financing, while Meta released an RFP targeting 1-4 GW of new nuclear generation for AI and sustainability objectives [5-10].

The public narrative had many article-level fragments, but many were repeated coverage of the same few commitments. For k*=2, the pilot coded 20 supportive article fragments but only four independent commitment clusters.

Metric at k*=2	Value
N_article	20
N_eff independent clusters	4
FI_naive	3.322
FI_eff	1.000
Copy Gap	2.322 bits (~5.0x overstatement)

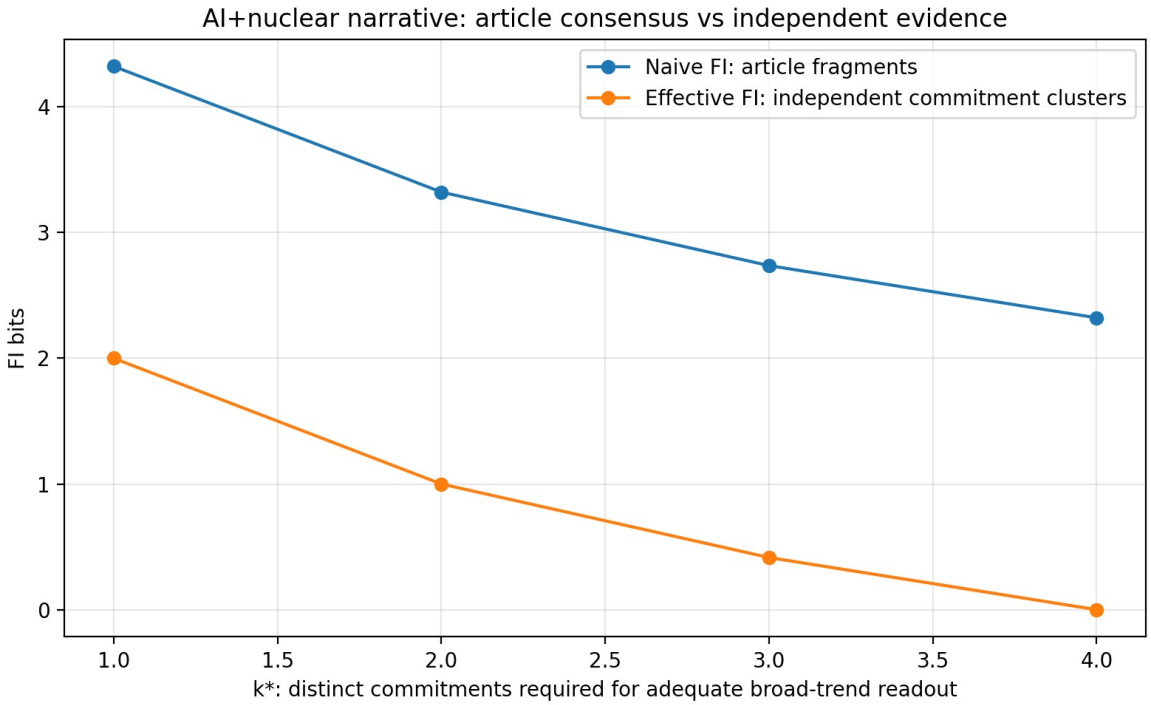

Figure 3. AI+nuclear case. The broad infrastructure narrative has effective support, but article-count consensus overstates independent evidence by roughly fivefold at k=2.*

- Finding: not false consensus. There are multiple independent primary/deal clusters.
- Copy Gap: large. Article volume materially overstates the independent evidence base.

- Boundary: the stronger stock-market translation - that broad nuclear/uranium/SMR baskets are all objective winners - is not established by the same evidence.

4. Case Study 2: GLP-1 Drugs and Alcohol Use

The second case examines an emerging biomedical narrative: GLP-1 receptor agonists may reduce alcohol craving or alcohol use. This case is useful because the weak claim and the strong overclaim diverge sharply.

The cautious target is that GLP-1s show a plausible early signal for reducing alcohol craving or use. Evidence clusters include a small randomized semaglutide trial, a mixed exenatide trial, observational/register studies, electronic-health-record analyses, and preclinical mechanisms. The semaglutide trial reported initial prospective evidence but explicitly justified larger trials rather than establishing a treatment standard. The exenatide RCT did not significantly reduce the primary heavy-drinking outcome overall, although subgroup and neuroimaging results were suggestive. Observational studies reported associations between GLP-1 use and lower AUD-related outcomes but require trials for causal confirmation [11-14].

The stronger target is that GLP-1s are already established AUD treatments. In the pilot, effective FI for this stronger claim is zero: the evidence supports promising investigation, not standard-care certainty.

Target claim at k*=2	Naive FI	Effective FI	Copy Gap
Promising signal	3.755	1.138	2.617
Established AUD treatment	2.700	0.000	2.700

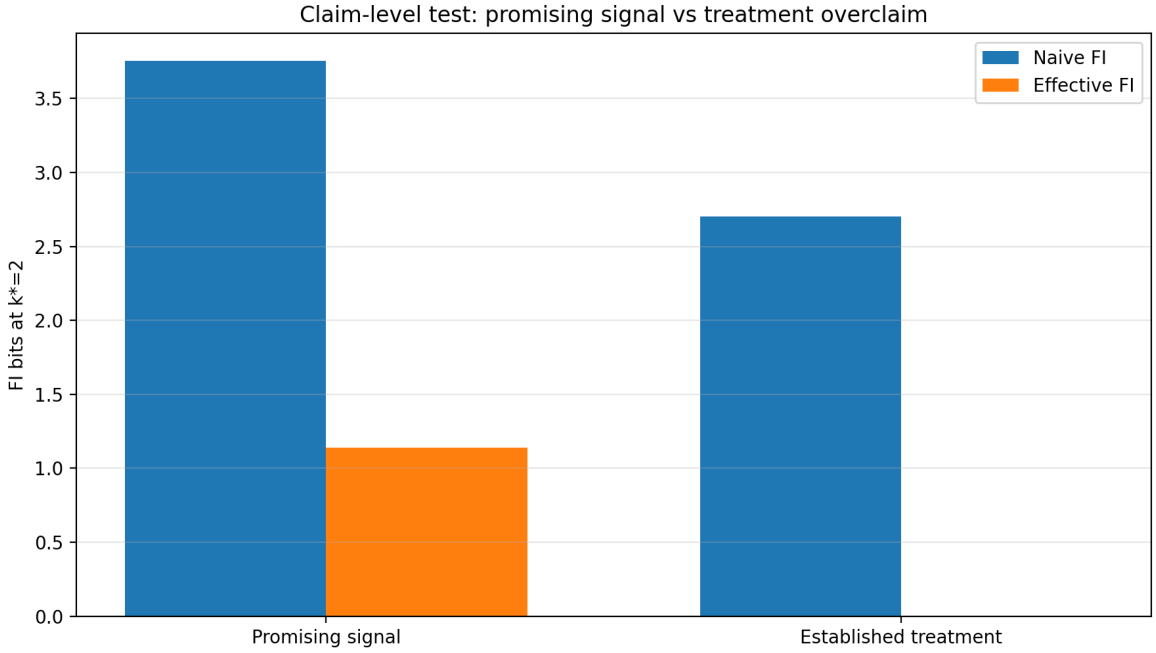

Figure 4. GLP-1/alcohol case. Effective evidence supports a cautious promising-signal claim but not the stronger treatment-established claim.

- Finding: real early objectivity for a cautious claim.
- Copy Gap: public-fragment consensus overstates independent evidence by about sixfold for the promising-signal claim.
- Boundary: the established-treatment claim is not effectively objectified; it is premature consensus inflation.

5. Case Study 3: Iraq WMD Consensus, 2002-2003

The historical case is the most decisive because later ground truth is available. The target must again be split by claim strength. The weak claim is that Iraq had unresolved WMD-related questions, past programs, and disarmament uncertainty. The strong claim is that Iraq possessed active, deployable WMD stockpiles or ongoing WMD production in 2002-2003.

The 2002 U.S. National Intelligence Estimate judged that Iraq continued WMD programs, and public presentations amplified that case. Yet contemporaneous inspection and dissent fragments already weakened the strong stockpile claim. The IAEA reported in March 2003 that the Niger uranium documents were not authentic and that the aluminum-tube evidence did not support the centrifuge claim. Postwar reviews later found the strong prewar WMD judgments wrong; the WMD Commission wrote that the Intelligence Community was "dead wrong" in almost all prewar Iraq WMD judgments, and the Iraq Survey Group did not find militarily significant stockpiles [15-18].

The Copy Gap method captures the distinction. The weak uncertainty claim retains partial effective objectivity. The strong stockpile claim has high naive FI but effective FI collapses to zero once copied streams, unreliable sources, and counter-evidence are accounted for.

Target claim at k*=2	Naive FI	Effective FI	Copy Gap
Unresolved WMD questions	3.000	0.766	2.234
Active deployable WMD stockpiles	2.700	0.000	2.700

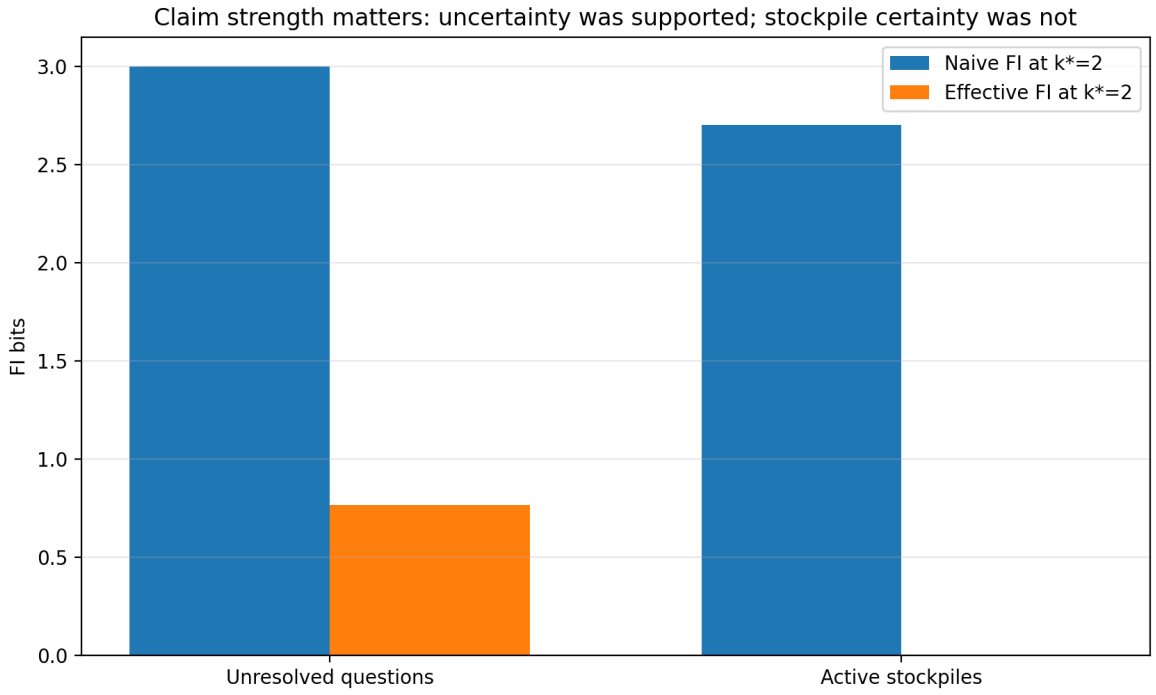

Figure 5. Iraq WMD case. Unresolved questions retained partial effective support; the stronger active-stockpile claim collapsed after effective-evidence correction.

- Finding: the method does not say there was no WMD-related uncertainty. It says the stockpile certainty claim was not effectively objectified.
- Copy Gap: the weak claim was over-amplified but not baseless; the strong claim was copied consensus and later failed against ground truth.

- Historical value: the case shows why claim strength matters. A weaker target can be partly objective while the stronger policy-relevant claim is not.

6. Cross-Case Findings

Across the three cases, the Copy Gap identifies different forms of overstatement. The AI+nuclear case is a real narrative with copy-amplified media volume. The GLP-1/alcohol case is a real early biomedical signal that becomes false consensus if translated into established-treatment certainty. The Iraq WMD case is a historical consensus failure in which the weak uncertainty claim survives but the strong stockpile claim fails.

The common pattern is not simply that high Copy Gap means false. High Copy Gap means that apparent consensus contains much less effective independent evidence than its nominal fragment count suggests. Whether the claim is true depends on the target and external validation.

Pattern	Case	Diagnostic signature	Interpretation
Real but amplified	AI+nuclear infrastructure	High effective FI, high Copy Gap	Narrative is real; article volume overstates independent evidence.
Early signal over-amplified	GLP-1/alcohol	Moderate effective FI for cautious claim; zero for treatment overclaim	Promising research becomes false consensus if upgraded too soon.
Historical consensus failure	Iraq active WMD stockpiles	High naive FI, zero effective FI	Policy-relevant consensus lacked adequate independent support.

7. Validation and Limitations

The case studies are manually coded pilots, not exhaustive corpus studies. They demonstrate use of the framework, not final empirical estimates. A full implementation should use reproducible source collection, preregistered coding rules, sensitivity analyses, and independent annotators.

The most important limitation is target dependence. The measure can show that a target is independently recoverable, but it cannot by itself prove that the target is true. A false target can become highly objectified if many independent fragments are captured by the same wrong structure. This is why later ground truth, causal access, or adversarial reference channels are required.

The second limitation is cluster coding. In these case studies, N_eff is approximated by independent evidence clusters. In larger datasets, N_eff should be estimated algorithmically using target-conditioned residual correlations, source lineage detection, citation or syndication graphs, and text reuse analysis.

The third limitation is k*. The value of k* encodes how many independent fragments are required for adequacy. The sensitivity curves in the underlying files show how scores change with k*, but applied use requires domain-specific justification of this threshold.

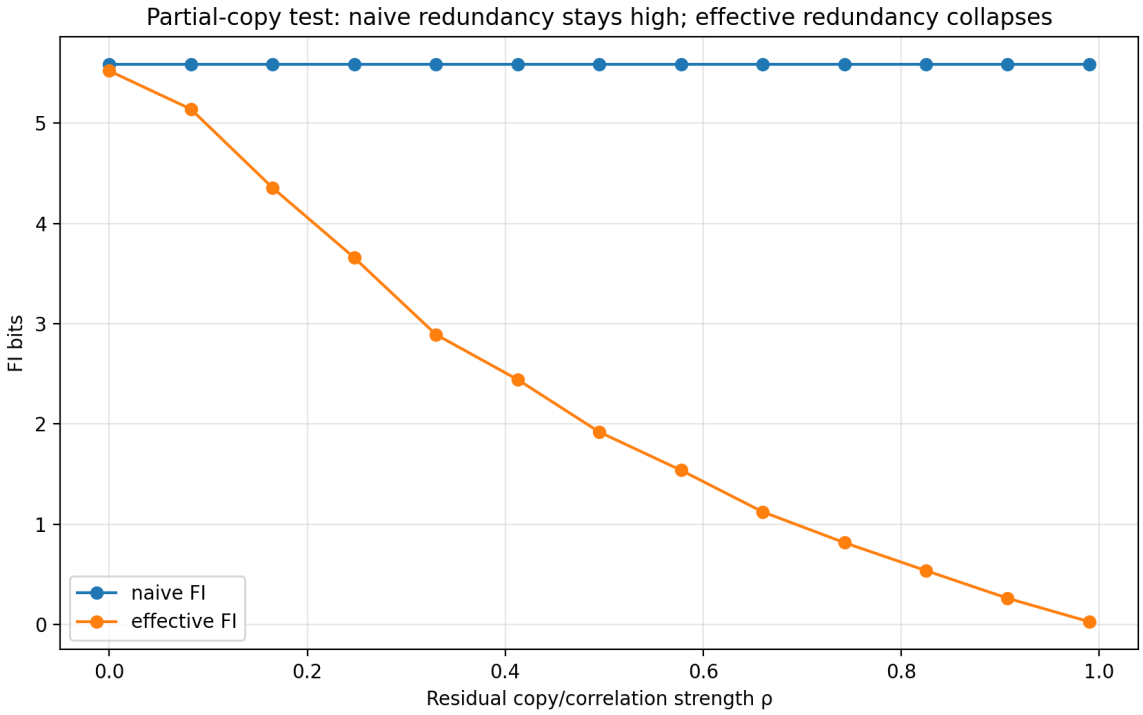

Figure 6. Simulation guardrail. Naive FI remains high under copied/correlated fragments, while effective FI collapses as residual dependence rises.

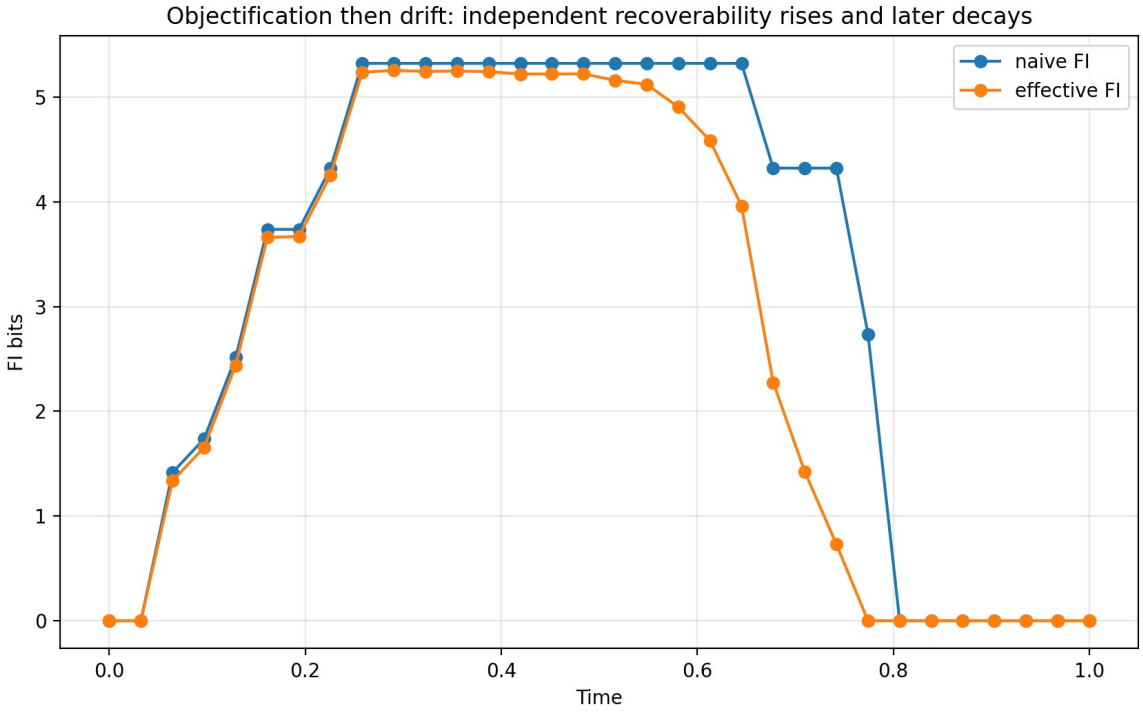

Figure 7. Simulation guardrail. Effective FI rises when independent recoverability appears and falls when drift/correlation destroys it.

8. Empirical Roadmap

Domain	Implementation	Falsifiable prediction
News/narrative corpus	Use GDELT or equivalent to collect articles; cluster by provenance, outlet family, quoted	Copy Gap predicts which narratives later

	source, and text reuse.	reverse or fail.
Scientific literature	Cluster papers by lab, dataset, method lineage, and citation dependence; compare with replication or meta-analysis outcomes.	High naive/low effective consensus predicts fragile fields.
Markets	Combine filings, earnings calls, analyst notes, flows, options, and news; compare Copy Gap with later fundamentals or returns.	High Copy Gap flags crowded narrative without independent grounding.
Clinical documentation	Compute effective recoverability across notes, target data, plans, supervision, and audit outcomes.	Falling FI_eff predicts documentation drift or compliance risk.

9. Conclusion

The Copy Gap is a way of asking whether consensus is public evidence or public echo. The case studies show that the same diagnostic can separate several epistemic states: real distributed evidence, copy-amplified early signals, premature treatment narratives, and historical consensus failure.

The central contribution is not a universal truth oracle. It is an operational distinction: naive consensus counts fragments; effective objectivity counts independent adequate readouts. The difference between them is often where the risk lives.

References

1. Quantum Darwinism / functional information framing: recent formulations define objectivity through redundant environmental fragments carrying pointer information. See arXiv:2509.17775.
2. Stan Development Team. Stan Reference Manual, effective sample size. https://mc-stan.org/docs/2_19/reference-manual/effective-sample-size-section.html
3. Information-theoretic redundancy and partial information decomposition; see recent redundancy-bottleneck formulations in Entropy.
4. Huh et al. The Platonic Representation Hypothesis. arXiv:2405.07987.
5. Constellation. Constellation to Launch Crane Clean Energy Center. 2024.
6. Microsoft. Accelerating the addition of carbon-free energy. 2024.
7. Google. New nuclear clean energy agreement with Kairos Power. 2024.
8. Kairos Power. Google and Kairos Power partner to deploy 500 MW of clean electricity generation. 2024.
9. Amazon and X-energy announcements on small modular reactors and carbon-free energy. 2024.
10. Meta. Accelerating the next wave of nuclear to power AI innovation. 2024.
11. Hendershot et al. Once-Weekly Semaglutide in Adults With Alcohol Use Disorder. JAMA Psychiatry, 2025.
12. Klausen et al. Exenatide once weekly for alcohol use disorder investigated in a randomized, placebo-controlled clinical trial. JCI Insight, 2022.
13. Lahteenvuo et al. Repurposing Semaglutide and Liraglutide for Alcohol Use Disorder. JAMA Psychiatry, 2025.
14. Wang et al. Associations of semaglutide with incidence and recurrence of alcohol use disorder in real-world population. Nature Communications, 2024.
15. National Intelligence Council. Iraq's Continuing Programs for Weapons of Mass Destruction: Key Judgments. 2002 NIE.
16. IAEA. The Status of Nuclear Inspections in Iraq: An Update. March 2003.

17. Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction. Report to the President. 2005.
18. Iraq Survey Group / Duelfer Report. Comprehensive Report of the Special Advisor to the DCI on Iraq's WMD. 2004.

Data and Reproducibility Note

The case-study coding tables and plots are provided as separate ZIP files generated during the pilots. The present manuscript is a synthesis paper; the coding is illustrative and should be replicated with preregistered criteria before making strong empirical claims.