

# HUNTER. Methodology Brief

John Malpass · School of Economics · University College Dublin

April 2026 · v1.0

## HUNTER

**An autonomous research engine that reads across 18 professional financial silos and keeps a public ledger of what the integration reveals.**

*John Malpass · School of Economics · University College Dublin · April 2026*

### The problem

Big financial firms are full of specialists who don't cross-read. A patent lawyer reads USPTO filings; she doesn't read NAIC insurance reserves. A CMBS servicer reads Trepp and special-servicer reports; he doesn't read Federal Register rule changes. Every silo is free and public. The *implication* across silos is private, because no one reads them all. Only the price ever shows it, and the price is often wrong. That gap, the part of mispricing attributable to information that is individually free but jointly requires cross-silo expertise, is what the standard limits-to-arbitrage framework (Grossman–Stiglitz 1980; Shleifer–Vishny 1997) does not address. HUNTER is an instrument built to make the gap empirically measurable.

### How it works (five parts, three of them new)

**Implication matching** (*new*). Every fact HUNTER ingests carries an explicit note about who in other professional communities should care about it and why. Cross-silo collisions are searched across those implication fields, not across shared keywords. Two facts can collide with no words in common if their implications overlap in a specific named way.

**Model-field extraction** (*new*). Every fact is tagged with five fields: methodology, assumption, practitioner community, calibration, and disruption channel. The resulting graph has nodes like “ARGUS Enterprise DCF cap-rate assumption,” not companies.

**Differential edge** (*new*). A causal arrow only enters the graph if the kill phase has verified a named transmission pathway: the specific filing, database, or workflow through which one silo's output becomes another's input. No named pathway, no arrow. The surviving graph is made entirely of edges you could point at.

**Mechanism-verified kill phase**. Multi-round, web-searched review: fact-check, competitor-check, barrier-check, and the hard one, the mechanism check. Survivors get scored in a fresh context by an adversarial reviewer against four calibrated anchors (92, 88, 35, 25), with no memory of the generation step.

**Pre-registration**. Corpus frozen 2026-03-31. Code locked 2026-04-19 at SHA-256 `f39a2f5ff6b3e695`. Four hypotheses, three null baselines, fixed decision rules, and explicit prohibitions on post-hoc additions, weight changes, and endpoint swaps all committed in `preregistration.json`. Any drift during the study is auto-flagged and reported.

### Data (frozen v1 corpus, April 2026)

Artifact	Count	Artifact	Count
Raw facts	12,030	Cross-silo collisions	474
Professional silos	18	Multi-link causal chains	52
Distinct countries	77	Directed causal edges (named pathway)	171
Distinct entities (normalised)	11,835	Hypotheses, adversarial review (archive 263 + main 61)	324
Entity-index entries	30,967	Findings (diamond $\geq$ 65)	45
Model-field extractions	6,670	Detected epistemic cycles (Tarjan)	9
		Theory-evidence (13 layers)	1,155

Code base: 60+ modules, 43 DB tables, 26 LLM prompts. Domain distance matrix: 153 hand-calibrated pairs. Three-tier model routing: Opus 4.7 for the critical reasoning stages, Sonnet 4.5 for standard extraction, Haiku 4.5 for high-volume ingestion.

## Pre-registered summer 2026 hypotheses

**H1 (primary).** Median realised alpha over SPY total return increases monotonically across strata  $A(1 \text{ silo}) \leq B(2) \leq C(3) \leq D(\geq 4)$ , with  $D - A > 0$  at  $p < 0.05$  via 10,000-resample paired bootstrap. If it fails, a null-result paper ships. **H2.** Detected cycles (reinforcement  $\geq 0.5$ ) persist uncorrected for  $\geq 14$  days in  $\geq 2$  of the 9 detected cycles. **H3.** Cross-silo collisions (distance  $\geq 0.60$ ) score  $\geq 10$  points higher than within-silo ( $< 0.30$ ) on average. **H4.** Chain-depth-3 hypotheses outperform chain-depth-1 at Cohen's  $d \geq 0.3$ .

**Null baselines (committed in advance):** B1 random-pair · B2 within-silo · B3 shuffled-label.

*Pre-freeze patterns (the mechanism/audience kill asymmetry, the hub-and-spoke methodology graph around ARGUS cap-rate assumptions, the bimodal diamond-score distribution, the narrative/survival correlation  $r = -0.49$ , the nine Tarjan cycles) are secondary hypotheses for the summer to test, not findings.*

## Paper programme

**Paper 0 (Methods).** Instrument, pipeline, the novel methodology triad, kill-phase design · SSRN Apr 2026. **Paper 1 (Empirical).** Summer-study results, null-baseline comparisons · SSRN Sep 2026. Additional theoretical and empirical working papers ship through summer and autumn.

## Timeline

Apr 28, 2026 launch · Jun 1 – Aug 31, 2026 pre-registered 12-week out-of-sample study · mid-Jul 2026 first resolutions on the public board · Sep 2026 Paper 1 on SSRN and full summer report · autumn 2026–2027 further paper submissions to peer-reviewed venues.

## Honest limitations

One operator, one instrument. 324 hypotheses (263 archive + 61 main) in the pre-freeze corpus; summer 2026 targets an additional  $n \geq 300$  out-of-sample. About 60% of pre-freeze high-scoring hypotheses concentrate in CMBS / insurance / regulatory-transition finance. Generator and reviewer share an LLM substrate; decorrelation is attempted via fresh context, web-searched kill rounds, calibrated anchors, and three pre-registered null baselines. **An earlier retrospective pilot (v3 Golden, web-search disabled) produced Stratum D < Stratum B, contradicting the primary hypothesis. The summer study runs prospectively with web-searched kill rounds, which is the regime H1 is actually about, and commits to ship the null paper if  $D \leq B$  or monotonicity fails. See README §Pre-registered Study.** Zero predictions resolved as of launch. Formal theoretical work awaits senior-theorist collaboration.

## Artifacts and contact

Repo [github.com/Johnmalpass/hunter-research](https://github.com/Johnmalpass/hunter-research) (MIT code, CC-BY-4.0 data) · corpus Zenodo DOI (pending) · methods paper SSRN (pending) · prediction board [Johnmalpass.github.io/hunter-research/](https://Johnmalpass.github.io/hunter-research/) · pre-registration [preregistration.json](https://preregistration.json) SHA-256 locked. **John Malpass · School of Economics · University College Dublin.** Email on the Zenodo record and repository README. Critique, prior-art pointers, and replication attempts welcome.