What a Filing Predicts: Self-Excitation, Incremental Value and the Placebo

On what disclosure behaviour carries, and how much of it survives being tested properly

Tan Dhingra

Meridian Deal Partners · Meridian Intelligence

Working Paper III · October 2026

Companion to “The Reaggregation of Dispersed Knowledge”

ABSTRACT

A filing is an event in time, and events in time have structure a cross-section cannot see. This paper reports three results about that structure and one about the discipline required to believe any of them. First, self-excitation: modelling a filing history as a Hawkes process, in which each event raises the short-run intensity of the next and that elevation decays, improves out-of-sample Brier skill from 27.7 to 30.71 per cent against the incumbent recency summaries. Paired on 20,057 identical rows the difference is 3.01 points with a bootstrap interval of 2.64 to 3.4, excluding zero. The hypothesis and its falsifier were registered before the test.

Second, and less comfortably: assembling more evidence families does not help. A standing audit scores each challenger against the capital family alone on identical rows. The challengers correlate with the benchmark at between 0.9988 and 1.000000; the institutional family adds an amount that is statistically detectable at ninety days and economically meaningless; and the macro family is negative at both horizons. More data made the forecast worse, and the audit says so without being asked.

Third, on returns. Four filing-dated signals were measured against a universe benchmark. All four looked strong, with t-statistics between 3.1 and 11.9. Under a placebo that holds the ticker fixed and moves only the date, two of them disappear: insider cluster buying returns 9.01 per cent excess over 126 days and the same tickers held from random dates return 8.56. Two survive by a distance. I close by registering 1,264 directional predictions on events whose windows have not yet closed, resolvable without my cooperation.

Keywords: self-exciting processes · disclosure · placebo testing · incremental value · point-in-time evaluation. JEL: C53, D83, G14.

1 INTENSITY, NOT COUNT

A firm that raised capital twice in the last ninety days is not in the same state as a firm that raised the same total once, three years ago. A feature set that counts prior filings, or measures days since the last one, cannot tell those two apart: the first discards time, the second discards everything but the most recent event.

Hawkes introduced the natural formalism in 1971 for earthquake aftershocks. The intensity at time t is a background rate plus a sum of decaying kernels over prior events — one term for each thing the entity has already done, each fading on its own timescale. Written for a filing history it says something economically specific: that disclosure clusters, that a filing makes another filing more likely in the near term, and that the elevation decays rather than persisting.

Three timescales were wired at thirty, ninety and three hundred and sixty-five days, with a burstiness ratio comparing recent intensity against the long-run rate. The features sit on all 412,154 feature snapshots in the estate, written alongside the incumbent columns rather than over them, so the comparison is between two feature shapes on one population rather than between two datasets.

The falsifier was fixed before the test: reject the null if a paired bootstrap of the skill difference on a strict temporal split excludes zero.

Table 1. Self-excitation against the incumbent recency shape, scored on identical rows at a base rate of 10.63 per cent. The paired interval excludes zero, so the registered null is rejected.

Feature shape	n	Brier skill	95% interval
Incumbent recency summaries	20,057	27.7%	25.1 – 30.23
With Hawkes kernels	20,057	30.71%	28.35 – 32.99
Paired difference	—	+3.01 pp	2.64 – 3.4

Note: the level of either row is a property of the scored population as much as of the model. The paired difference is not, which is why it and not the level is the reported result.

The outcome here is a filing, not a return. Predicting that a filer files again is a weaker thing than predicting that a project is built, and this paper does not pretend otherwise. What the result establishes is narrow and real: the arrival process carries information that the incumbent feature shape was discarding, and the gain is three points of skill rather than a rounding error.

2 Assembly does not pay here

The companion paper argues that the binding constraint on public information has moved from acquisition to integration. That claim invites an obvious test: if integration is where the value is, adding an evidence family to a model should improve it.

A standing audit does exactly this, scoring each challenger against the capital family alone on identical rows with paired bootstrap intervals. The result does not support the loose version of the thesis and is reported at measured value.

Table 2. Incremental value of each additional evidence family over the capital family alone, on identical rows.

Horizon	Added to the capital stack	Prediction corr.	Incremental Brier	P(improve)
90d	capital+institutional	0.999995	+0.0000126	0.984
90d	capital+macro	1.000000	-0.0000008	0.092
90d	meta:all	0.999995	+0.0000118	0.981
180d	capital+institutional	0.998787	+0.0001097	0.678
180d	capital+macro	0.999995	-0.0000386	0.038
180d	meta:all	0.999124	+0.0000922	0.783

Note: a negative incremental Brier means the addition made the forecast worse. Each audit is stored five times in the source table; these rows are deduplicated.

The challengers correlate with the benchmark at between 0.9988 and 1.000000. The capital family carries essentially all of the signal. The institutional family adds something detectable at ninety days and nothing at one hundred and eighty, and the macro family is negative at both, with probabilities of improvement of 0.09 and 0.04.

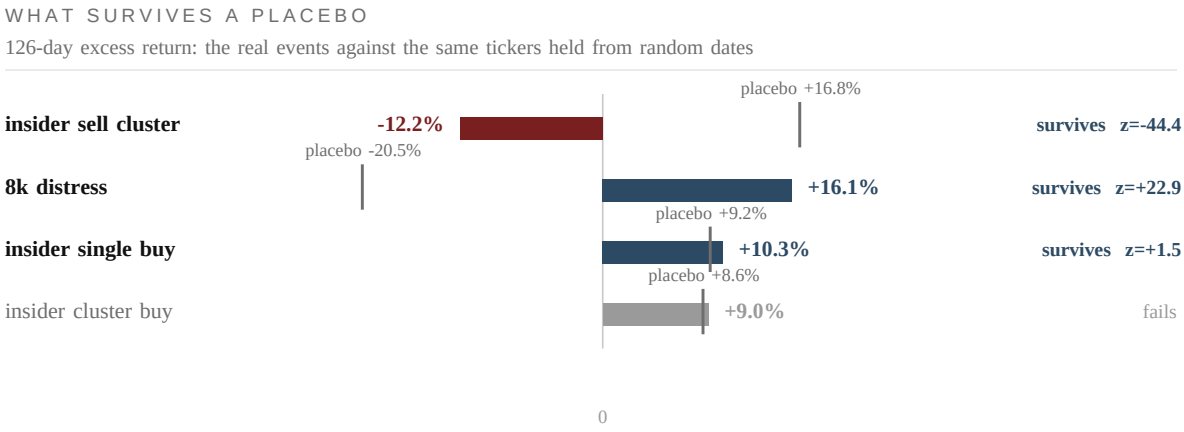
The honest reading is that assembly pays on questions no single corpus holds the state for, and not on questions one corpus already answers. Whether a filer files again is the second kind: the filing history already contains most of what is knowable about it, and bolting macro series onto that adds noise with a plausible story attached. The companion paper's physical measurements are the first kind, because no register contains both what was announced and what was built.

3 The placebo, and the two signals that failed it

The estate holds a price store of 10,422 securities and a ledger measuring forward return against a winsorised equal-weight universe index, for filing-dated events entered on the first trading day after the filing. Dating on the filing rather than the transaction matters more than any modelling choice here: an earlier study in this estate showed a large apparent premium on insider transactions that collapsed to insignificance once dated on the only day the information was public.

Four signals carry enough observations to test, and on their face all four look strong, with t-statistics between 3.1 and 11.9. The placebo holds everything constant except the thing under test: for each real

observation the event date is redrawn uniformly from the same calendar window, keeping the ticker and the per-ticker count fixed, and excess return is recomputed with identical machinery. Two hundred draws per cell.



Each bar is the mean excess return of the real events. The marker is the mean of two hundred placebo draws that hold the ticker fixed and move only the date. Where the two coincide, the signal is a property of which companies appear in the sample rather than of the event.

Source: Meridian signal ledger and price store, 200 placebo draws per cell, seed fixed.

Figure 1. Real excess return against the placebo null at 126 days. Where bar and marker coincide, the apparent edge is a property of which companies are in the sample rather than of the event.

Insider cluster buying returns 9.01 per cent excess over 126 days. The same tickers held from random dates return 8.56. The event is worth roughly half a point and the other eight and a half are a property of which companies insiders buy. Insider single buying behaves the same way at every horizon but the longest, where it clears the threshold narrowly. Both of the intuitive, publishable signals are substantially artefacts of stock selection.

The two that survive do so by a distance. Distress-flagged 8-K filers return 16.14 per cent excess against a placebo of minus 20.51, a separation of nearly thirty-seven points at z = 22.9. Clustered insider selling returns minus 12.17 against a placebo of plus 16.81, at z = −44.4. In both cases the placebo baseline is itself informative: these are populations that move sharply on their own, and the event identifies a departure from that drift rather than the drift itself.

Two limitations remain and are not corrected. Events are grouped by calendar month and dated at the month's last filing, which is not knowable in real time, so entry as specified is not implementable and the next revision must use a rolling window. And the price universe carries no delisting flag: a ticker that delists loses its forward return, which biases the real sample toward survivors, overstating the distress signal and understating the selling signal. Direction is known; magnitude is not.

4 Registered predictions

The two signals that survived the placebo are used to register directional predictions on events whose 126-day windows have not yet closed. There are 1,264 of them, spanning 2024-10-09 to 2026-09-21. The two signals that failed are excluded by construction, which is the point of having run the placebo at all.

Table 3. Registered directional predictions on unmatured events, published as a companion file with the ticker, event date and resolution date for every row.

Signal	Direction	Open predictions
Distress-flagged 8-K filers	Outperform benchmark	183
Clustered insider selling	Underperform benchmark	1,081
Total	—	1,264

Note: these are directional predictions of excess return against the universe benchmark at 126 trading days, entered on the first trading day after the filing. They are not recommendations, carry no position sizing, and take no account of costs, capacity or borrow availability.

Resolution is mechanical and needs nothing from me: each event's window closes on a known date, and the realised excess return either has the predicted sign or it does not. The aggregate claim is that these two signals, net of the placebo, continue to separate from their own null. The failure that would matter most is not a lower hit rate but convergence — the real returns drifting toward the placebo distribution, which would mean the events had stopped identifying a departure and the earlier separation was a property of one period.

A second falsifier is cheaper and sharper. If the two signals that failed the placebo in this paper were to outperform the two that passed, the placebo itself would be suspect rather than the signals, and the whole selection procedure here would need rebuilding. Both outcomes are observable from the same file.

Method and sources

Model figures come from the Meridian one-model lab: 531,416 out-of-sample holdout predictions across horizons, produced under a walk-forward protocol with purged folds, of which the ninety-day horizon carries 60,250 scored rows. The Hawkes hypothesis was registered on 18 October 2026 with the falsifier quoted in §1 and is recorded validated. Return figures come from filing-dated events in the Meridian signal ledger, scored against a winsorised equal-weight index built from the same price store; the placebo was run on 25 October 2026 with a fixed seed.

Two results in this series have been withdrawn after audit and are named here so that a reader encountering them elsewhere knows their status: a private-credit aggregate that disagreed with the filers' own published totals, and a county-level attention corpus whose geocoding proved empty. A third correction, to interconnection-queue counts inflated by counting observation rows rather than projects, is described in the second paper of this series.

Interest. The author operates the instrument described and intends to commercialise research derived from it. Nothing here is investment advice, no security is recommended, and no claim is made that any security or constraint identified is mispriced.