What This Is

A plain explanation of the Meridian research programme, what it found, and what it did not

Meridian Deal Partners · November 2026

Most research you are sent shows you what worked. This shows you what was tested, what survived and what did not, because the second list is the one that tells you whether to believe the first.

1 The idea in one paragraph

Four decades of disclosure law pushed an enormous quantity of economic detail into the public record: who owns what, who crossed which threshold on what date, which power station exists and in what state of construction, which company told a regulator it was in distress. None of it is secret. All of it is fragmented — each register built by a different institution for its own purpose, none designed to be read against another. The opacity is not secrecy. It is a filing problem, and a filing problem is an engineering problem.

Meridian assembles those registers into one estate and asks questions no single register can answer. It currently holds 5,058,265,273 records, of which 1,864,734,940 are joined to a single identity spine. The rest are held for reference and are never mixed into an analytical claim.

2 The first thing we did was try to fail

A research programme that only reports its successes cannot be checked. So before publishing any finding, the pipeline was pointed at a result somebody else had already published, to see whether it could recover it.

Brav, Jiang, Partnoy and Thomas reported in the Journal of Finance in 2008 that when an activist investor files an initial Schedule 13D, the target's shares show roughly a seven per cent abnormal return around the announcement, and that the gain is not subsequently given back. We ran the same measurement on our own independently assembled filings and our own price store.

Table 1. The replication test. Recovering both the magnitude and the absence of reversal, from a different period than the original study, is stronger evidence than matching the headline alone.

Window	Events	Abnormal return	t	What it shows
(–20,+20)	266	+8.16%	2.63	published benchmark ≈ +7%
(+1,+60)	287	+0.21%	0.09	no reversal — their central claim

Note: our abnormal return is measured against a universe index rather than a risk model, and our events are 2025–26 against their 2001–06. The test is whether the pipeline can recover a number it was never tuned to recover. It can.

That is the whole purpose of the exercise. If a machine cannot reproduce a known answer, nothing it produces about unknown ones counts.

3 What we then tested, and what happened to it

Four signals were measured on filing-dated events — dated on the day the information became public, never the day the underlying act occurred, because the difference between those two dates is where most false results in finance come from.

On first measurement all four looked strong, with t-statistics between 3.1 and 11.9. They were then put through four tests, each designed to kill them:

A placebo. The same stocks held from random dates. If the event carries no information, the real result should sit inside that distribution. Two of four did.

Characteristic matching. Distressed firms and firms with insider activity are more volatile than average. Each event was matched to five control firms on pre-event volatility and momentum, so the comparison is against similar firms rather than the market.

A skewness adjustment. Long-horizon returns are badly skewed and the ordinary t-test over-rejects. We use the skewness-adjusted statistic.

A clustering correction. Events sharing a date are not independent evidence. p-values come from resampling whole dates, not observations.

Table 2. All four signals after every correction. Three of the four edges shrank by 75 to 82 per cent once matched against similar firms.

Signal	Thesis	vs index	vs matched	p	Verdict
8k distress	short	-17.65%	-15.15%	0.0005	Survives
insider single buy	long	+9.26%	+2.33%	0.0115	Survives
insider sell cluster	short	+11.19%	+2.01%	0.0195	Survives As A Long
insider cluster buy	long	+8.06%	+1.85%	0.1569	Fails

Note: insider sell cluster was carried as a short. Matched, those stocks outperform — so it survives with its sign reversed, and is reported as it measured rather than as it was framed.

Read that table honestly and it says that most of what looked like insider information was the characteristics of the firms insiders buy. Insider single buying fell from 9.26 per cent to 2.33. Cluster buying fell from 8.06 to 1.85 and out of significance entirely. Those are not edges; they are small-cap and volatility exposure wearing an event's clothing.

One result held. Shorting companies that file a distress-flagged 8-K returns 15.15 per cent against matched controls, with t = –6.03 and a bootstrap p of 0.0005, on 847 events across 326 separate dates. It shrank only fourteen per cent under matching, where the others lost three quarters of their apparent edge.

It was then tested the way a desk would test it. Half of all distress filings are repeats by the same company; removing the overlap costs 1.7 points and the result holds. Delaying entry does not kill it — minus 12.8 per cent entering the day after the filing, minus 14.3 entering five days later, minus 13.2 at ten days — so the underperformance accrues over months rather than being a same-day reaction, and being first is not the edge. On a borrowable universe above five dollars it is minus 11.8 per cent at t = –3.6.

After seven tests built to destroy it, that is the one thing here we would call an edge — and it is a short, which is not the finding anybody hopes for. What we have not modelled is the cost of borrowing these names, which is exactly where a short like this is hardest to implement, and nothing above establishes that the spread survives it.

4 What the physical work found

The same method applied to infrastructure produces results that need no statistics at all, because they are counts rather than estimates.

46.5% OF ANNOUNCED POWER CAPACITY ARRIVED

Of 202.2 GW proposed to the US federal government in 2022, 94.0 GW was operating three years later. Two independent cohorts agree within 2.2 points, so it is a rate rather than a bad year.

285 sites WHERE THE GRID CONNECTION ALREADY EXISTS

136.9 GW of retired thermal capacity sitting in counties with live 345 kV substations and no announced reuse. A developer entering an open queue waits five to seven years; one starting where headroom exists waits twelve to eighteen months.

63.7% OF REQUESTED INTERCONNECTION WITHDRAWN

1,066 GW of 1,672 GW ever requested across four system operators. The queue keeps its failures, which is what makes it measurable.

None of this is a forecast. It is what the files already record, read against each other.

5 What we have had to withdraw

Four things have been published and then corrected. They are listed because a programme that only shows corrections it was not forced to make is not showing you very much.

A private-credit aggregate.
Position counts and a fair-value total were published, then withdrawn when an outside reader asked whether they reconciled against the filers' own statements. They did not — one large lender reconstructed to three to four times its reported portfolio. Found by a stranger, not by us.

Interconnection queue counts.
Published roughly sixfold too large because observation rows were counted instead of projects, and the inflation was then described as growth in the dataset. It was double counting.

A sign convention.
The signal ledger stores each result already multiplied by the direction of its thesis. We read those as the stocks' own returns, which inverted the meaning of both short signals. Distressed filers were reported as rising when they fall.

A county-level attention corpus.
A 113-million-event news dataset was to be joined to project outcomes by county. Its latitude column turned out to hold the date and its longitude column was empty. Abandoned before it produced a number rather than after.

Three of those were caught before publication and one after. That ratio, stated rather than rounded up, is the only track record this programme currently has.

6 How to check any of it

Every figure above comes from public files. Nothing requires our permission to verify.

The power conversion rates	EIA Form 860, annual vintages 2022–2025, matched on plant code and generator identifier.
The 285 sites	EIA Form 860 for retirements, HIFLD for substations. The full screen and its definition are published as a CSV.
The queue figures	The published interconnection queues of MISO, CAISO, NYISO and SPP.
The activist replication	SEC Schedule 13D filings and any daily price source.
The signal results	SEC Forms 3/4/5 and 8-K, dated on filing. The method is stated in full in the papers.
The forty forecasts	Published before their outcomes, resolvable against the 2028 EIA vintage without our involvement.

7 What we do not claim

- That any of this is investment advice. It is not, and no security is recommended.
- That the 285 sites have usable interconnection rights. The join is by county, not by verified electrical connection. It tells you where to look, not what you can connect to.
- That the distress signal is implementable as measured. Events are currently dated at the last filing of a calendar month, which is not knowable in real time.
- That our price universe is free of survivorship bias. It carries no delisting flag; the direction of the bias is known and its size is not.
- That we have a performance record. We have measurements, published forecasts, and a list of things we got wrong.