VIVAMEDA

Know where a company is heading,
before the numbers show it.

PREDICTING CORPORATE TRAJECTORIES FROM LONGITUDINAL EVIDENCE

Oli Nold, Founder · oli@vivameda.com

Pre-seed · Raising EUR 300,000

THE PROBLEM

You buy a company on numbers that describe last quarter.

What you can see

- Revenue, margins, debt
- Headcount today
- What management tells you

What you cannot

Whether this business is quietly becoming a different shape, and what usually happens next when it does.

By the time it reaches the accounts, you have already signed.

The team changes before the numbers do.

Companies do not decline all at once. The mix of people shifts first: which functions grow, which quietly stop being replaced. That happens years before it reaches a P&L.

3 TO 5 YEARS AHEAD

The team shifts

Engineering stops being replaced. Sales grows. The shape moves.

1 TO 2 YEARS AHEAD

Performance follows

Products slip. Costs rise. Growth stalls.

TOO LATE TO ACT

The numbers show it

It appears in the accounts. Everyone can see it.

We measure the first step. Everyone else measures the third.

A prediction engine grounded in what happened before.

An actuarial table, for companies. An insurer does not know your driver. They know what happened to a hundred thousand drivers who looked exactly like them. We do that for companies.

1. **You give us a company**

   Any company. It does not have to be one we have seen before.

2. **We find its historical twins**

   Thousands of companies once in the same position: same industry, same size, same direction of travel.

3. **We tell you what happened to them**

   Of 2,638 companies that looked like this one, 37.9% were materially smaller within four to eight years.

Today: a rate with its cohort attached. Next: regime, sector and size weighted into one calibrated probability, the way a weather forecast does.

The same panel finds the winners.

Our own research, 2026. 41,000 companies and 302,000 observations, trained on one set and scored on another held entirely out. Methodology and limitations published in full.

27x

the base rate

Companies in the top 1 per cent of scores entered hypergrowth at 27 times the rate of the population, across 12,285 held-out companies.

0.89

area under the curve

Measured on the held-out set, against a hypergrowth base rate of 0.41 per cent.

22

known unicorns checked

82 per cent scored in the top 30 per cent before they scaled. A small set, and we treat it as directional rather than as proof.

Avoiding one bad acquisition and finding one good one come from the same evidence base.

Target is hypergrowth, defined as workforce doubling within three years, which precedes unicorn status rather than being identical to it. The decline result on the next slide carries the larger sample.

Tested across 65,000 company trajectories.

The largest test we have run, and the one that carries the weight.

WITHOUT THE SIGNAL

1 in 4

of all companies in the test had materially shrunk four to eight years later.

n = 64,634

WITH THE SIGNAL

3 in 4

of the companies the signal flagged had materially shrunk.

n = 6,876 flagged

Rule fixed before any outcome was computed. Signals read in 2012, outcomes measured 2016 to 2020.

A matched random control returned no lift. The effect strengthens with company size, which reduces the likelihood that sparse-data noise explains it. Controls and full methodology available.

We built the machine that makes the history.

The scarce thing is not career records. It is seventy years of them reconstructed into a company-by-company history, with the same company recognised across decades and name changes, and the outcome attached.

Raw profiles

Millions of career records, acquired and collected over years.

Our pipeline

Resolves the employer, reconstructs each company year by year, validates the result.

The panel

Every company, every year: how many people, and which kinds.

4M
companies

48M
company-years

1950
earliest record

100+
countries

Career records can be bought. The panel cannot. Entity resolution across decades, normalisation, coverage controls and attached outcomes are the work, and they are why nobody else has one.

WHY NOW

Models are becoming abundant. Evidence is not.

What is getting cheap

- Analytical capability
- Modelling and agents
- Anyone can now analyse anything

What is not

A record of how millions of real companies actually evolved, over seventy years, with what happened next attached.

As models improve, this dataset becomes more valuable, not less. Better models interrogate it better. They cannot create it.

You cannot learn consequence from a snapshot.

Almost all company data is cross-sectional. It describes a moment. A model trained on moments learns what things look like. It never learns what they turn into.

TRAINED ON SNAPSHOTS

Learns correlation

Which features tend to appear together. It can describe a company confidently and have nothing behind the confidence.

TRAINED ON HISTORY

Learns consequence

What actually followed, across millions of company-years. It can tell you how often this ended badly, and show you the cases.

The fix for a confident wrong answer is not a better model. It is evidence of what actually followed.

The history teaches how companies move between states. Current data says which state a company is in now. That combination is what makes a forecast rather than a description.

Start with the people who buy companies for a living.

BEACHHEAD

Lower-middle-market investors

Independent sponsors, search funds and boutique private equity. Small teams, no research function, one bad acquisition matters enormously.

~2,500 organisations · EUR 25k average · EUR 60m

EXPANSION

Credit, monitoring, corporate development

Private credit and specialty lending, portfolio monitoring for funds, and corporate development teams screening acquisitions. Same instrument, larger contracts, longer cycles.

The beachhead is the wedge, not the ceiling. The same reference layer extends into private credit, portfolio monitoring, corporate development and embedded data partnerships.

~1,400 US independent sponsors (McGuireWoods 2025) · ~360 SBIC funds · 862 search funds tracked since 1984, 197 currently searching (Stanford GSB 2026) · 503 international search funds (IESE 2025)

Counted bottom-up from named segments. Boutique private equity and corporate development are additional and not included above.

Land on a deal. Expand to the pipeline.

Deal intelligence

EUR 5k to 22k

Per transaction. One target, full assessment. How we land.

Pipeline intelligence

EUR 40k to 80k

Annual, account-wide. Screening across everything they look at. How we expand.

Data partnership

Negotiated

Licensing the reference tables into someone else's product.

Today: concierge reference analysis, delivered by us. After this round: a repeatable product a deal team runs itself.

One investment decision can be worth tens of millions. Unlimited pipeline coverage is priced against that, not against a seat.

Sixty engagements. Four paying in month one.

Commercial delivery began in August 2026, after the panel and the method were already built and validated.

PAYING

4 paying organisations

2 licensees paying fees to build on the panel

Priced at the tiers on the previous slide.
Not pilots and not free trials.

IN TESTING

4 active trials

Labs and quantitative teams working with the data directly, on their own time rather than ours.

IN EVALUATION

3 partnership evaluations

Data and platform partners assessing the panel as a reference layer inside their own products.

20+ analyses delivered · EUR 40k qualified pipeline · 5 research studies published in full

What we have built.

4M+	48M	1.88B	65k
companies reconstructed	company-years	skill rows	company backtest

ON

Oli Nold

Founder

Over a decade founding and running startups, including businesses sold. Ran seven-figure performance marketing budgets before moving into data. Spent the last year building the panel, the method and the validation, then the product around it.

DE

Data engineering

Full time

Built the enrichment pipeline, the warehouse and the analytical infrastructure that produces the panel. Two independent builders work on the client-facing layer, paying licence fees rather than being paid.

Vivameda owns the panel and the pipeline outright, with no licence underneath them. The predictive layer was modelled against that asset with a technical partner, and this round builds an internally owned version. Commercial delivery does not depend on the partner.

EUR 300,000 for twelve months.

1. **Ownership** — The predictive layer rebuilt in-house and owned outright, on our own infrastructure.
2. **Product** — A self-serve company assessment a deal team runs without us in the loop.
3. **Accounts** — Six to ten institutional accounts on the annual tier.
4. **Validation** — Predictive performance independently verified, with published error rates.

The model owned outright, six to ten accounts, and verified predictive performance.

oli@vivameda.com · vivameda.com

Why the incumbents have not built this.

Workforce data vendors sell what a company looks like now and how it compares to peers. None of them publishes what historically happened next to companies in the same position.

	Current workforce	Deep longitudinal history	Outcomes attached	Historical analogues
Professional networks	Yes	Partial	No	No
Deal and funding databases	Partial	No	Yes	No
Workforce data vendors	Yes	Partial	No	No
Vivameda	**Yes**	**Yes**	**Yes**	**Yes**

The gap is not capability. A current-state product never required seventy years of reconstructed history.

Positioning framework, not a vendor audit. Verify each cell before making the claim to a specific company by name.