VIVAMEDA LTD

ENGAGEMENTS AND FIELD EVIDENCE

Client Portfolio

TWENTY ENGAGEMENTS

February to August 2026. One longitudinal instrument, applied to twenty different economic decisions.

THE INSTRUMENT

4M+	**48M**	**1950**	**1.88B**
companies 100+ countries	observed company-year records	to 2026 coverage window	skill rows 46.5M capability buckets

Commercial intelligence from observed workforce evidence
Vivameda Ltd, Cyprus · Prepared August 2026

VMD.21958.150826.240

OVERVIEW

What this portfolio is

Twenty engagements delivered between February and August 2026, grouped by the economic decision each one was built to improve rather than by the industry the counterparty operates in.

The subject is the instrument, not the projects

Each entry records the question as it arrived, what was constructed to answer it, and the generalisable point the work established. The engagements are evidence. What they are evidence for is a single longitudinal panel and a repeatable method of applying it, which is why the same analytical moves appear in a private equity diligence question, a non-profit finance benchmark and a German freight forwarding study.

What is not here

The five published studies behind this work are in the Vivameda Research Portfolio and are not repeated. Product documentation, licensed data packages and reference assets are excluded, because they are not engagements and cannot be compared to them. Engagement structuring and pricing work is excluded for the same reason.

How it is organised

Three decision classes rather than three industries. Acquisition, diligence and allocation covers decisions where the cost of being wrong is concentrated in a single position. Advisory and intermediation covers decisions taken on behalf of a third party, where the evidence has to survive being repeated to a client who was not in the room. Operating growth and market development covers decisions about where an operating business directs scarce commercial capacity. The same instrument reads differently against each, because the decision it feeds differs in timing, reversibility and cost of error.

How each entry is built

The question as it arrived, what was constructed to answer it, and the transferable finding, followed by one exhibit and the recurring techniques visible in the work. Every figure in an exhibit is stated in the body text of the same entry.

> All engagements ran under confidentiality arrangements. Counterparties are identified by category and geography. Document titles are Vivameda's own. Where a figure would identify a specific company inside a counterparty's portfolio, client list or target set, that company is described rather than named.

SELECTION

How the set was reduced

Sixty documents were produced across the period, delivered to a smaller number of counterparties because several received more than one. Twenty engagements are reproduced here. The set is not a random draw and the criterion is stated so the reader can discount accordingly.

Inclusion criterion

Four conditions, all required. A stated question the counterparty would recognise as their own. Something actually constructed or computed, rather than assembled. A finding that generalises beyond the engagement it came from. And an evidential anchor: a population, a peer distribution, a base rate, a matched control or a measured observation window. Work that produced a usable output without a transferable finding is excluded.

Criterion rejected

Recency, taking the twenty most recent, was rejected because it selects on the availability of the record rather than its content. Volume was rejected for the same reason: the book is shorter than the record deliberately, and a complete project list would demonstrate activity rather than capability.

One entry per counterparty

No counterparty appears twice. Where several documents were produced for the same counterparty, they are consolidated into a single entry that records the sequence, because the sequence is usually the more informative object: the second document exists because the first provoked a harder question. Three entries in this portfolio cover more than one document. The set therefore understates the volume of work and shows its continuity instead.

Failed and discarded work is retained

Six entries record a signal killed against a peer distribution, a published null, a contradicted hypothesis, a collection artefact identified before it reached a document, or a framing abandoned after testing. These are not qualifications attached to successful work. They are the reason the surviving claims are usable, and an instrument that has never returned a negative result has not been tested.

How Vivameda works

Twelve techniques recur across the twenty. None was specified in advance. Each emerged from a particular build, survived contact with a buyer who had reason to object, and was then reused. The technique transfers between engagements; the individual file does not.

M01 Negative-space filtering

Define the target by what is missing rather than what is present: no marketing leader, no commercial layer above the engineering bench, no in-house engineering capacity at all. Absence is harder to fabricate than presence and almost nobody indexes for it.

M02 The qualifying gate

One binary condition, verified from named individual records rather than inferred from size or sector. It converts a plausible population into a defensible one and is the most common reason a sample survives scrutiny.

M03 Two-column evidence

Each row carries a current live signal and a historical pattern match, never merged into a composite. The first establishes that the account is real now; the second establishes that the reasoning is not a guess. Neither is persuasive alone.

M04 Deliberate disqualification

Retain the records that failed, with the reason attached. A file containing only the qualified reads as a sales document. A file showing its own rejections reads as an instrument with a stated threshold.

M05 Dated signal anchoring

Every reason to make contact carries a date and a source: a published criterion, a panel appearance, a stated position, a posted role. Undated relevance is indistinguishable from flattery and is treated as absent.

M06 Speed-to-close tiering

Split the qualified population by how quickly the decision can be taken rather than by how well it fits. A prospect with a junior person already in the role has an open budget line, which frequently outranks a better-fitting prospect who must create one.

M07 The withheld contact layer

The opening sample carries the complete reasoning and none of the actionable contact data, stated openly rather than omitted quietly. The counterparty can assess the quality and cannot execute on it, and nothing is constructed to fill the gap.

M08 Signal killing

Test each candidate signal against a peer distribution and discard the ones that fail, on the record. Two were discarded inside one engagement in this portfolio and sixteen inside another. The surviving signal is credible because the others were removed in view.

How Vivameda works

M09 Scope framing

State what the instrument cannot observe before the reader discovers it, and state it as an operating parameter rather than a concession. A constraint presented as a limitation invites negotiation; the same constraint presented as a specification does not.

M10 The checkable subject

Build the analysis on companies the reader already knows better than the analyst, using none of their materials and contacting nobody. Verification becomes immediate and the argument about test selection never begins.

M11 The paired comparison

One company against a matched peer through the same external shock, rather than a profile of the company alone. A profile is an anecdote. A pair is a measurement, because the control absorbs everything the two have in common.

M12 The published null

Report the indicators that failed, with their coefficients, before reporting anything that worked. A study that discloses its own dead ends is the only kind a technical reader can act on, because it is the only kind whose survivors mean anything.

Where each technique appears, read the other way round: by problem rather than by counterparty

M01	Negative-space filtering	17, 18
M02	The qualifying gate	17
M03	Two-column evidence	01, 04, 11, 13
M04	Deliberate disqualification	05, 08, 12, 18, 20
M05	Dated signal anchoring	05
M06	Speed-to-close tiering	18, 19, 20
M07	The withheld contact layer	12, 19
M08	Signal killing	06, 08, 09, 10, 16
M09	Scope framing	01, 02, 03, 04, 06, 09, 10, 11, 15, 16
M10	The checkable subject	01, 02, 03, 07, 14
M11	The paired comparison	07
M12	The published null	14, 15

CONTENTS

Twenty engagements, three decision classes

	ENGAGEMENT	COUNTERPARTY	DATE
CLASS ONE / ACQUISITION, DILIGENCE AND ALLOCATION			
01	**The Shape Test**	Lower middle market software private equity, United States	August 2026
02	**Portfolio and Sourcing Brief**	Venture fund, Italy	August 2026
03	**The Same Company Twice**	Venture assessment firm, Netherlands	August 2026
04	**The Fourth Document**	Sell-side M&A boutique, industrials, United States	August 2026
05	**Allocator Leads**	Quantitative currency manager, Hong Kong	June 2026
06	**Financing Signal Sample**	Commercial real estate financing marketplace, Germany	August 2026
CLASS TWO / ADVISORY AND INTERMEDIATION			
07	**Same Shock, Different Speed**	Leadership assessment platform, United States	August 2026
08	**The CPA Target List**	Exit-planning engagement, United States	July 2026
09	**The Absent Recruiter**	Recruitment intelligence, Benelux	August 2026
10	**Finance Layer Benchmark**	Fractional finance and operations advisory, United States	July 2026
11	**Client Match Sample**	Organisational consultancy, United Kingdom	June 2026
12	**The Watchlist Letter**	Growth advisory, consumer brands, United Kingdom	August 2026
13	**Account Intelligence for Complex Sales**	Channel and introduction partner	July 2026
CLASS THREE / OPERATING GROWTH AND MARKET DEVELOPMENT			
14	**The Coordination Gap**	Build partner for go-to-market systems, United States	August 2026
15	**Early Indications in Logistics**	Logistics and freight forwarding, German-speaking Europe	August 2026
16	**Segment Test, Fiduciary and Care**	Cybersecurity practice, Switzerland	August 2026
17	**Build Window Named Accounts**	Demand generation and build practice, United States	July 2026
18	**The Companies That Do Not Have a CMO**	Marketing consultancy, United States	July 2026
19	**Deal Map, Swiss AI Compliance**	Sovereign AI and EU AI Act compliance software, Switzerland	June 2026
20	**Sector Sample, Testing and Certification**	Compliance software vendor, international	June 2026

CLASS ONE / ACQUISITION, DILIGENCE AND ALLOCATION

01

The Shape Test

Lower middle market software private equity, United States · PDF, ten pages · August 2026

THE BRIEF

Twenty five platforms and seventy add-ons, roughly three bolt-ons for every platform owned. At that cadence the binding constraint stops being capital and becomes judgement speed: how quickly a target can be read before anyone opens a data room. The question was whether an outside instrument could produce that read on companies the buyer already knows better than we do.

WHAT WAS BUILT

Eight pages built entirely from panel reads on two of the firm's own companies in the years around their transaction dates, then extended to ten with a live layer covering nine portfolio companies and the trajectory model's three-estimator ensemble output with regime labels and confidence scores. No firm materials were used and nobody at any of the companies was contacted, which is stated on the cover rather than buried. Two companies returned live data with no historical panel entity and one domain resolved to an acquirer rather than the target; all three are named in the coverage section.

> **WHAT IT PROVED**
>
> Building the analysis on the buyer's own portfolio removes the argument about test selection before it starts, because the reader can falsify any claim faster than anyone else could. The strongest signal came from the live layer rather than the history, which establishes that the two layers carry independent information rather than one confirming the other.

The sharpest live finding, one portfolio company measured against its own peak

The only company of the nine posting no open roles. Decline accelerating, not flattening.

METHODS The checkable subject · Two-column evidence · Scope framing

02

Portfolio and Sourcing Brief

Venture fund, Italy · PDF, three versions · August 2026

THE BRIEF

Five funds, roughly five hundred million under management, about fifty eight portfolio companies, and portfolio monitoring software already in place. The question a document has to answer in that situation is not whether the data is good. It is what a longitudinal record adds to a monitoring stack that already exists.

WHAT WAS BUILT

Three versions. The first covered seven portfolio companies across thirty five months each with named historical analogues drawn from the panel. The second added nine off-book Italian and Spanish sourcing candidates, none of them held by the fund, so the document served origination as well as monitoring. The third was restructured around four concrete findings, with one line retained on a collection artefact that was caught and excluded, and the two-layer architecture stated as a design decision rather than a caveat.

> **WHAT IT PROVED**
>
> A fund that already monitors its portfolio does not need a second monitoring tool. It needs the comparison set its own stack structurally cannot hold, because a monitoring system observes the companies it has been pointed at and a panel observes the ones it has not. The nine off-book candidates are what converted a portfolio review into an origination argument.

What the fund already sees, and what the brief added

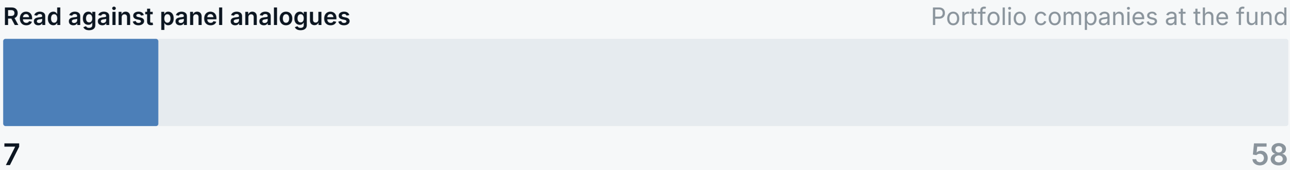

Nine off-book sourcing candidates were added in version two. None is held by the fund and none appears in its monitoring stack.

METHODS The checkable subject · Scope framing

03

The Same Company Twice

Venture assessment firm, Netherlands · Two documents · August 2026

THE BRIEF

The firm grades every dimension of a venture on an evidence scale that separates assumption from anecdote from proof. Applied honestly, the dimensions where the money actually dies, team, organisation and scale readiness, run on the weakest evidence in the system: a deck slide, a management interview, a headcount nobody can check. The question was what an observed workforce record adds to that architecture.

WHAT WAS BUILT

Four European companies read live, with no illustrations and no composites, and a closing section mapping the panel onto their own dimension and gate scheme rather than proposing a parallel one. The firm's technical evaluator then returned four questions about what sits underneath the reported figures, which produced a second document. Every claim in it was computed directly rather than restated: the base rate was measured on the correct within-table denominator across 2,622,245 company-years, the pattern flags were confirmed not to exist as a production field anywhere in the schema, and the coverage curve was rebuilt from whole-panel queries rather than from nine hand-picked companies.

> **WHAT IT PROVED**
>
> Two companies can carry identical headcount, product and revenue and be different businesses, because one grew into that number and one shrank into it. A snapshot cannot separate them and a longitudinal record can. The follow-up established the second half of that claim: a figure is only interpretable against the population it was measured on, and naming the denominator ended the technical objection rather than opening it.

Flag precision against the underlying base rate, measured on 2,622,245 company-years

A lift of 3.0x on this population. The separate pre-hypergrowth classifier is measured on a different population and is not comparable to it.

METHODS Scope framing · The checkable subject

04

The Fourth Document

Sell-side M&A boutique, industrials, United States · PDF · August 2026

THE BRIEF

An advisory firm produces three documents on every mandate. The question was whether a fourth, built from the workforce record rather than the financials, changes how a transaction is read. Four of the firm's own completed transactions were the test set.

WHAT WAS BUILT

All four companies were placed on a single workforce taxonomy with identity confirmed by domain matching against each acquisition announcement before any figure was computed. The central construction was a two-variable diagnostic: average tenure direction read jointly with headcount direction. Each company's observation depth is stated alongside its reading, and one company in the original candidate set was too thin to carry an analysis and was named as such rather than dropped.

> **WHAT IT PROVED**
>
> A single variable read alone is ambiguous and read in pairs is decisive. Falling tenure means expansion when headcount is rising and replacement when headcount is flat, and those are opposite facts about the same company. The financial record does not make that distinction, which is precisely why the workforce record is not redundant with it.

The diagnostic the four transactions were placed on

Two variables, each with two directions, produce four distinct readings of an identical tenure movement.

METHODS Two-column evidence · Scope framing

05

Allocator Leads

Quantitative currency manager, Hong Kong · Workbook · June 2026

THE BRIEF

A capital allocator target book for a manager who has to reach a small and well-defended audience, where a generic approach is not merely ineffective but disqualifying. The commercial question was what constitutes a legitimate reason to make contact at all.

WHAT WAS BUILT

Every signal column was rebuilt from dated public activity rather than from profile attributes: one allocator's published allocation criteria, another's recent conference panel appearance, a third's stated position on externalised verification, a fourth's stance on digital assets, a fifth newly in seat. Two allocators who do not post publicly were marked as such rather than given a constructed hook. Tier one and tier two engagement steps were rewritten to reference those specific dated signals, so the sequencing is not a template with names inserted.

> **WHAT IT PROVED**
>
> In a small and defended audience the date is the signal. Undated relevance is indistinguishable from flattery, and marking the two non-posters as non-posters was more useful to the manager than inventing a reason to approach them, because it told him where to spend the effort of a warm introduction instead.

Hooks that survived the dating rule

Seven allocators. Nothing was softened into a hook; the two without dated activity are stated as such in the file.

METHODS Dated signal anchoring · Deliberate disqualification

06

Financing Signal Sample

Commercial real estate financing marketplace, Germany · PDF, twenty four pages plus data annex, German · August 2026

THE BRIEF

A financing marketplace needed to identify lenders building capacity in a specific credit segment and investors approaching a financing window, before either side begins to search. The counterparty asked for financial-grade rigour rather than a target list.

WHAT WAS BUILT

Sixteen companies read against the full universe of German banks above one hundred employees, with the universe counted before selection so the base rate under the signal could be stated rather than implied. Five specific gaps were closed on request: the base rate, the universe size, how the approach would be backtested, whether the lead time is measurable, and the pilot economics. A monthly job-posting count was tested as a candidate signal and rejected as too noisy to carry weight, with the failing series shown in the annex rather than removed.

> **WHAT IT PROVED**
>
> Naming the universe is what converts a signal into an economic argument. A finding without a denominator cannot be priced, and a signal that has not been tested against the population it was drawn from cannot be distinguished from a coincidence.

Companies read against the addressable universe they were drawn from

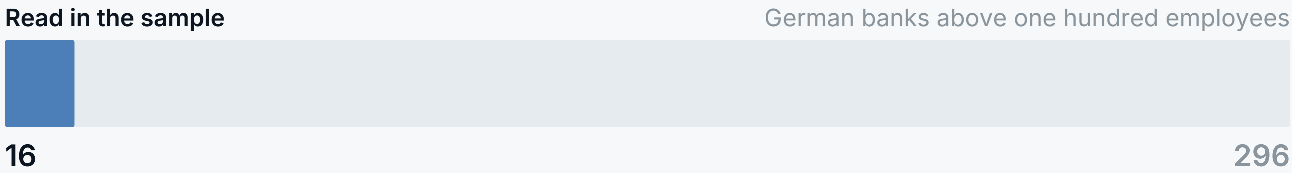

Universe counted before selection. One candidate signal was tested against it and rejected.

METHODS Scope framing · Signal killing

07

Same Shock, Different Speed

Leadership assessment platform, United States · Three documents · August 2026

THE BRIEF

The counterparty scores individual leader capacity. The complementarity is exact: they measure capacity, the panel measures organisational demand. The sequence began with an unsolicited note ahead of a first call, and then two follow-on requests, each one narrower than the last.

WHAT WAS BUILT

The first document read a single listed company chosen because its leadership outcomes are publicly verifiable, and graded the panel's own 2020 regime classification against what actually followed, including where it missed. The second was requested on a named company the counterparty was already in conversation with. A single-company profile was rejected internally as too thin to carry a conclusion and rebuilt as a matched-peer pair read through the same external shock, as a difference-in-differences design with robustness checks. The third moved to person level, sizing a senior function against nine peers and a cross-section of forty six comparable companies.

WHAT IT PROVED

One company through a shock is an anecdote and the same company against a matched peer through the same shock is a measurement. The deceleration in the first study was in the record for years before the reversal arrived, which is the general claim: organisational demand changes before the outcome that makes it visible.

The deceleration sequence visible in the record before the peak

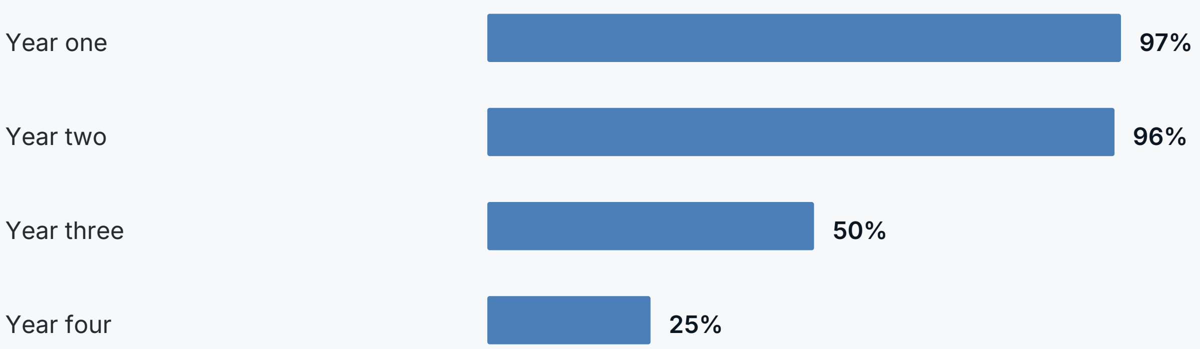

Year-over-year workforce growth. The observed workforce peaked in the same month the founder chief executive resigned, then fell 46 per cent across three leadership regimes.

METHODS The paired comparison · The checkable subject

08

The CPA Target List

Exit-planning engagement, United States · Workbook, twenty five firms, six drafts · July 2026

THE BRIEF

Find owner-operated accounting firm founders approaching an exit decision, for a twenty five thousand dollar engagement sold through a third party. The reader of the file is a salesperson rather than an analyst, which set the presentation constraint: reasoning without methodology.

WHAT WAS BUILT

Eleven query passes against the panel. The target definition combined the bottom quartile of decade headcount growth with the top half of a partner-layer thickening measure, scored inside a sixteen-point exit window carrying thirty five per cent weight in a four-component composite. Two candidate signals were killed during the work. The first senior finance hire and its co-occurrence variant are both structurally meaningless in this sector, because the finance role bucket at an accounting firm captures the billable delivery workforce rather than a back-office function. A replacement signal was derived from the change in executive share of headcount between two windows and validated against a peer distribution of 1,002 firms.

> **WHAT IT PROVED**
>
> Killing two signals on the record is what made the third one believable. A signal that has never been tested against a peer distribution cannot be distinguished from a definition, and the confirmed targets are only interesting because the distribution they sit outside was measured first.

Where the confirmed targets sit against the peer distribution

Percentile of the replacement signal, peer distribution of 1,002 firms

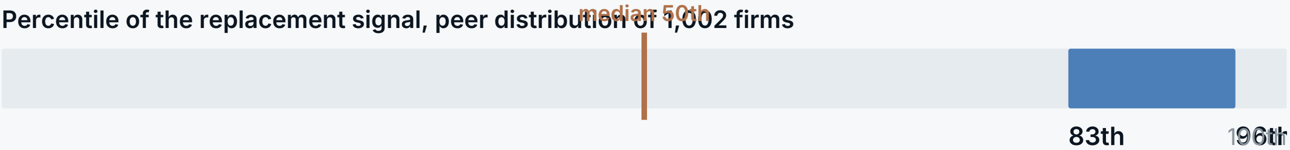

The median firm in the same distribution shows minus 0.3 points of change. Two prior candidate signals were discarded before this one was derived.

METHODS Signal killing · Deliberate disqualification

09

The Absent Recruiter

Recruitment intelligence, Benelux · PDF, ten pages · August 2026

THE BRIEF

A written question sent as a follow-up after a call: can an employer's likelihood of using external agencies be scored from workforce structure alone, using hiring velocity measured against in-house talent acquisition headcount.

WHAT WAS BUILT

The buyer's ratio was the right anchor and was corrected twice before it was used. Hiring velocity has to be gross hires rather than net headcount growth, because a company flat at five hundred people with twenty five per cent churn is still hiring one hundred and twenty five a year, and that backfill volume is exactly where agencies operate. In-house capacity has to be weighted by the difficulty of what is being hired, because a full-cycle recruiter closes far fewer senior engineering roles than standard ones. Six structural indicators in total, including one measuring the lag between growth onset and internal recruiting capacity arriving. Six Benelux companies carried the evidence across 2016 to 2020, with a live layer for validation and nineteen scaling companies sampled to establish the base rate the lag indicator is measured against. Every overclaim was softened to propensity language on review.

> **WHAT IT PROVED**
>
> A buyer's own metric is usually right and incomplete. Correcting the denominator of a ratio the buyer already trusts is a stronger opening than introducing a new metric, because it starts from something they have already accepted rather than something they have to be persuaded of.

What the study was built from

The base rate sample exists so the lag indicator has a distribution to be measured against rather than a threshold to be asserted.

METHODS Scope framing · Signal killing

10

Finance Layer Benchmark

Fractional finance and operations advisory, United States · PDF, eleven pages · July 2026

THE BRIEF

An advisory firm placing finance leadership into mid-sized non-profits wanted to identify organisations carrying a finance function out of proportion with their size. The first framing, finding organisations with no finance function at all, was tested and abandoned.

WHAT WAS BUILT

Observation density on organisations of this size runs at roughly five to ten per cent of true headcount, which makes an absolute count of a function unusable: a missing function reflects unobserved staff rather than absent ones. The measure was changed to relative composition against peer cohort medians, which is a ratio and therefore invariant to the sampling rate. Three organisations were shortlisted, confirmed for current identity, staffing and open roles, then written up. A longitudinal case traced one organisation's finance share across sixteen observed years against outcome-grouped cohort reference paths.

> **WHAT IT PROVED**
>
> The observation density determines which questions the instrument can answer, and the correct response is to change the question rather than to weaken the claim. An absence test fails at five to ten per cent density; a ratio between functions does not, because both numerator and denominator are sampled at the same rate.

Observation density at this organisation scale, and what it permits

Observed share of true headcount at this organisation scale

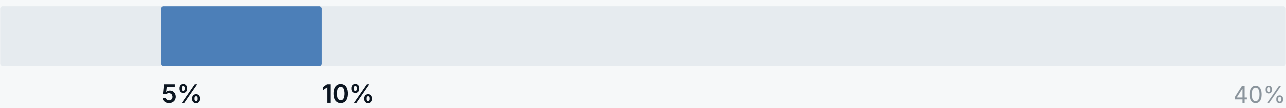

A ratio between two functions is invariant to the sampling rate. An absolute count of a single function is not.

METHODS Scope framing · Signal killing

11

Client Match Sample

Organisational consultancy, United Kingdom · Workbook, seven companies · June 2026

THE BRIEF

The founder held sixteen years of pattern recognition in his own head and wanted a pipeline matched to the moment his work lands hardest, rather than a list of companies that fit a description.

WHAT WAS BUILT

Seven UK firms across a wide size band, after three rounds of live pulls. Each row carries two evidence columns that are never merged: a current live workforce signal with year-on-year movement, and a historical pattern match from the panel. The historical column was labelled as pattern inference rather than a live panel query, because the full run happens inside the engagement and the sample exists to show the format. Coverage was spread deliberately across sustained scale-up, pre-listing maturation, strategic reset and lean reset, so the file demonstrates the method reading contraction as clearly as growth.

> **WHAT IT PROVED**
>
> Neither column is persuasive alone. The current signal proves the account is real now; the historical match proves the reasoning is not a guess. Labelling the inferred column as inferred cost nothing and removed the only objection available to a reader who wanted one.

The size band the seven companies were spread across

Employees per company, seven United Kingdom firms

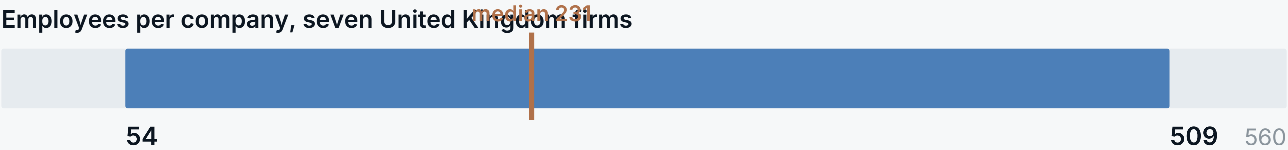

Spread deliberately across four trajectory types, including two distinct forms of contraction.

METHODS Two-column evidence · Scope framing

12

The Watchlist Letter

Growth advisory, consumer brands, United Kingdom · PDF, nine pages plus a three-page watchlist · August 2026

THE BRIEF

A growth adviser working with founder-led consumer brands needs to reach companies at the point where outside capital has just landed or is about to. The question was whether the moment is observable from outside.

WHAT WAS BUILT

The people layer was built to a verification standard printed in the document: a profile link appears only where the exact profile was confirmed, and a name otherwise stands alone with the confirming step indicated. Nothing was inferred from a name match. One funding signal was upgraded during verification from an anonymous private equity round to a named acquirer with a dated announcement, which sharpened the reason to make contact rather than merely the fact of it.

> **WHAT IT PROVED**
>
> A stated verification standard carries more than a uniformly filled column. The reader can see which rows are actionable immediately and which require a confirming step, which is an instruction a complete-looking table cannot give and a reader cannot reconstruct.

The people layer under the stated verification standard

Seven named individuals across six companies. Each row states which of the two categories it falls into.

METHODS The withheld contact layer · Deliberate disqualification

13

Account Intelligence for Complex Sales

Channel and introduction partner · PDF · July 2026

THE BRIEF

A partner who already delivers introductions and needed to see what sits underneath them: which accounts to approach, why now, and who to speak to. A partner cannot resell a capability they cannot restate without the author present.

WHAT WAS BUILT

Rather than a spread of illustrative accounts, the document walks a single realistic campaign end to end: one named seller, a three-country target definition, the account logic shown before any name appears, and eight decision makers pulled live while the document was being written. The choice of one worked example over many partial ones was deliberate, because a worked example is repeatable by a third party and a capability list is not.

> **WHAT IT PROVED**
>
> For an intermediary the reasoning is the deliverable and the names are evidence that the reasoning runs. A document that can be repeated by someone who did not build it is a different object from a document that demonstrates what the builder can do.

One worked campaign rather than a spread of illustrations

One seller	Three countries	Account logic	Eight contacts
one real proposition	target definition	shown before any name	pulled live

A capability list cannot be restated by a partner in a client meeting. A worked example can.

METHODS Two-column evidence

14

The Coordination Gap

Build partner for go-to-market systems, United States · PDF, twenty one pages · August 2026

THE BRIEF

A partner building operating systems for mid-market clients needed to identify the moment a company outgrows its own coordination capacity, which is the moment their engagement lands. The specification came from a recorded call rather than a written brief.

WHAT WAS BUILT

Five construction companies read longitudinally with a sixth as control, on three co-moving observables: tenure decay, specialist compression and executive absorption. Panel trajectories were combined with job postings and capability reads, then corroborated against county planning records, property transactions and trade press. One hypothesis carried into the work, that a company's first technology hire starts the vendor clock, was contradicted by the data and published as a numbered finding rather than removed. Seven findings in total.

> **WHAT IT PROVED**
>
> Third-party revenue data on private mid-market companies is unreliable in a way that is demonstrable rather than arguable. One company appeared in a major commercial aggregator at nineteen million dollars in the same year it completed a land transaction of one hundred and eighteen point eight million, which is a matter of public record.

One company, as reported by a commercial aggregator and as evidenced by a single public transaction

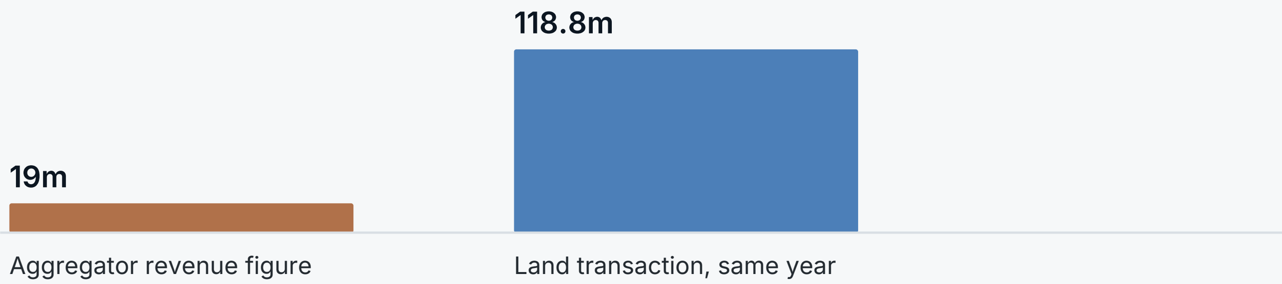

United States dollars. The aggregator figure was in active commercial use at the time of the analysis.

METHODS The checkable subject · The published null

15

Early Indications in Logistics

Logistics and freight forwarding, German-speaking Europe · PDF, twenty three pages plus a briefing sheet, German · August 2026

THE BRIEF

A client briefing asked which workforce patterns separate compounding freight forwarding companies from stagnating ones, to be used as lead indicators by a sales function.

WHAT WAS BUILT

Panel coverage for the cohort looked too thin and a current-layer pull produced 1,256 usable firms instead. Sixteen indicators were tested across five lead distances and none reached a correlation of 0.10. An apparent seventy eight per cent cohort growth rate was identified as coverage drift rather than hiring before it reached the document. The study also sets out why a workforce-only insolvency indicator cannot be constructed backwards at all: German insolvency notices are public only briefly, so individual firms cannot be attributed as failed after the fact. The closing section is a costed forward design monitoring the same 1,256 firms for twelve to eighteen months, which yields roughly twenty five events.

> **WHAT IT PROVED**
>
> The null result is the deliverable. Sixteen failed indicators published with their coefficients, followed by the only construction that can actually answer the brief, is a more useful commercial position than a manufactured finding, because the client can now spend on the design that works rather than on the one that reads well.

Why the backward-looking version cannot be built

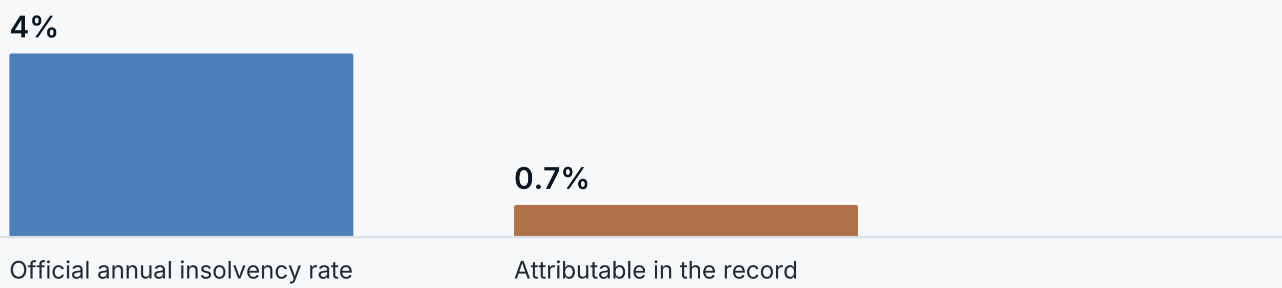

The events exist. They are not attributable to named firms after the fact, which is why the proposed design runs forward rather than backward.

METHODS The published null · Scope framing

16

Segment Test, Fiduciary and Care

Cybersecurity practice, Switzerland · PDF, German · August 2026

THE BRIEF

A security practice selling to small and mid-sized organisations wanted signals that identify a buyer before a procurement process opens. Coverage in this market is densest at enterprise scale, so the segment had to be selected before any signal could be specified.

WHAT WAS BUILT

Two adjacent segments were tested against the observation base, fiduciary accounting firms and residential care. The fiduciary segment carries the required density and produced three confirmed signals: a technical headcount that stays near nine tenths of one per cent regardless of firm size, an onboarding burden metric, and a merger spike visible within a fortnight. Seasonality was tested as a fourth signal and cleanly falsified. Delivered first in the branded format and then rebuilt in plain black and white on request.

> **WHAT IT PROVED**
>
> A constant that holds across firm size is a stronger commercial instrument than a correlation, because it converts any observed deviation into a named quantity rather than a rank. The onboarding finding is the clearest instance: one hundred and twenty eight new hires across nine firms served by twenty one HR roles, with one firm carrying neither an HR nor a technical function.

Onboarding load against the internal capacity available to absorb it, nine fiduciary firms

Four candidate signals were tested in this segment. Three held across firm size and one, seasonality, was falsified.

METHODS Signal killing · Scope framing

17

Build Window Named Accounts

Demand generation and build practice, United States · PDF and workbook · July 2026

THE BRIEF

The client sells fixed-scope software builds in the low six figures to operations-heavy service businesses. The commercial question was timing: reach them at the moment the build exists and before the tender does, because that is the only moment a large fixed-scope engagement is an easy yes.

WHAT WAS BUILT

A five-signal model was written first, covering capital and capacity, the first systems hire, tooling admissions, volume shock and a new operator at the top. Each surviving account passed a binary gate confirming zero software engineers, developers or technical leads on staff, read from named employee records rather than inferred from size or sector. One contact had no locatable public profile and was flagged as such, with outreach redirected rather than a link constructed. A follow-on workbook extended the work upward into twenty two enterprise accounts, adding an interpretation column alongside eight standard signal columns.

WHAT IT PROVED

The interpretation column is the only part of the file that cannot be reproduced by another vendor. Signal columns are commodity output; the historical context that explains why a signal matters at this company is not, and it is the only column a seller can open a conversation with.

Where the defensible part of the enterprise file sits

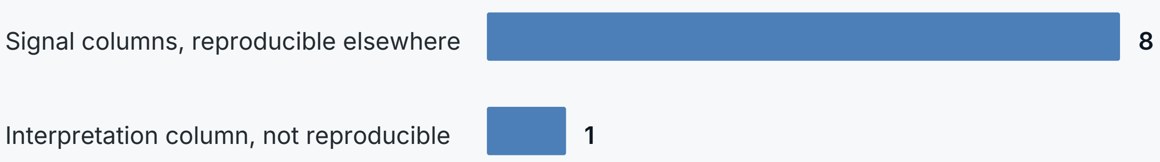

Nine columns per account across twenty two accounts. Named-account selection was gated on zero in-house engineering capacity, verified from individual records.

METHODS The qualifying gate · Negative-space filtering

18

The Companies That Do Not Have a CMO

Marketing consultancy, United States · PDF brief and workbook · July 2026

THE BRIEF

The brief arrived as an unusually precise written profile: software development agencies, revenue between two and twenty five million, growing, with nobody holding a marketing director, vice president or chief marketing officer title. Her entire target was defined by an absence.

WHAT WAS BUILT

A live scan of eleven United States agencies produced seven qualified, three disqualified because marketing leadership was already in seat, and one below the revenue floor. All four outcomes were kept in the document. The qualified list was split into two tiers separating firms with no marketing function at all from firms carrying a junior marketer and no leader above them, on the argument that the second group closes faster because the budget line already exists. The filter was then run backwards against her own client list, which showed three former clients now carrying marketing leadership, one of them with three people reporting to him.

> **WHAT IT PROVED**
>
> The finding was structural rather than a list. Her own success systematically converts qualified prospects into disqualified ones, which replaces the case for a one-time pull with the case for a standing feed. An absence closes without warning, so a list defined by absence decays from the day it is delivered.

Eleven agencies scanned, all four outcomes retained in the file

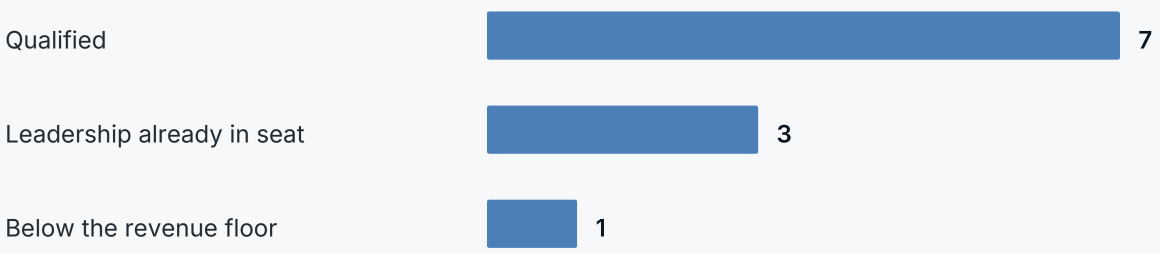

A file containing only the qualified reads as a sales document. A file showing its own rejections reads as an instrument.

METHODS Negative-space filtering · Deliberate disqualification · Speed-to-close tiering

19

Deal Map, Swiss AI Compliance

Sovereign AI and EU AI Act compliance software, Switzerland · Workbook, German · June 2026

THE BRIEF

A seat-based Swiss compliance product sold domestically, where the founder confirmed during the qualifying exchange that the product functions across the wider German-speaking market and that the regulatory driver is not Swiss-specific. The unit economics were confirmed directly, which set the commercial frame of the sample.

WHAT WAS BUILT

Seven decision makers across insurance, banking, healthcare, trust services, energy, industry and retail, with no sector doubled. Each carries a buyer fit score, a seat and deal estimate, a why-now trigger, the likely objection with its counter, an opening line and an outreach sequence. Contact data was withheld and stated as withheld rather than omitted silently. The file was then restructured on instruction from a lead list into a deal map, with an executive summary, a commercial case showing that a single mid-size close covers the engagement, and a referral sheet. All seat figures are marked as estimates derived from company size.

> **WHAT IT PROVED**
>
> The same seven accounts read entirely differently once the structure changed from a list to a map. Spreading one account per sector rather than concentrating in the densest sector is what makes the file a market description instead of a target list, and a market description is what supports a pricing conversation.

Seven accounts, one per sector, no sector doubled

Seven sectors	Per account	Contact layer	Commercial case
one account each	trigger, objection, counter	withheld and stated	one close covers it

Seat and deal figures are marked as estimates derived from company size, not as observed values.

METHODS The withheld contact layer · Speed-to-close tiering

20

Sector Sample, Testing and Certification

Compliance software vendor, international · Workbook, thirteen leads · June 2026

THE BRIEF

Reach the people who feel a scheduling and resource-planning problem daily inside the largest testing, inspection and certification groups, alongside the directors who hold the budget for solving it. In a category with a handful of serious organisations, the difficulty is locating the function rather than the company.

WHAT WAS BUILT

Two personas were targeted in parallel rather than sequentially: front-line planners and resource coordinators who experience the constraint, and operations and regional directors who can authorise a purchase. Six of the eight largest global groups returned verified contacts. Two returned nothing usable even after the searches were re-run in German and French respectively, and were left visibly empty in the file rather than padded with weak matches. Each row carries workforce and operational scale for the parent group, the persona, the reason the person is a prospect, and an opening drawn from the client's own positioning.

> **WHAT IT PROVED**
>
> Thirteen verified leads read better than twenty with filler, because the visible gap where two organisations should appear is itself a statement about the verification standard. Where a category has eight serious buyers, depth inside each organisation is worth more than breadth across many.

Coverage across the eight largest global groups

Searches re-run in two additional languages before either organisation was recorded as empty.

METHODS Deliberate disqualification · Speed-to-close tiering

What repeats across the twenty

Seven statements hold across more than one entry in this portfolio. Each is carried by at least two engagements in unrelated markets, which is the only reason it appears here rather than in the entry it came from.

Longitudinal comparison outperforms snapshot description

Two companies carrying identical present-day figures can have arrived there by opposite routes, and no snapshot separates them. This is the recurring structure behind the private equity, venture assessment and sell-side entries.

A figure without a denominator cannot be priced

Base rate against a stated population appears as the decisive move in the venture assessment, exit planning and financing marketplace entries. In each, the estimate was already available and the population was not.

Peer-relative evidence survives thin observation where absolute values do not

Where the observed share of true headcount runs at five to ten per cent, a ratio between two functions is invariant to the sampling rate and a count is not. This determined the instrument in the non-profit finance and accounting firm entries.

Two variables read jointly resolve what one variable leaves ambiguous

Falling tenure means expansion or replacement depending on headcount direction; hiring velocity means agency reliance or internal capacity depending on the recruiting layer beneath it. Single-variable readings were rejected in both cases.

The current and historical layers are kept adjacent and never blended

They rest on different observation bases and different counting rules. Every entry that reports figures from both states which instrument produced each, and no entry computes across them.

Negative evidence carries commercial weight

Sixteen indicators published with their coefficients, two signals discarded against a peer distribution, one candidate signal rejected as noise, one hypothesis contradicted and published as a finding. In each case the disclosure is what made the surviving claim usable.

The strongest subject is one the reader can independently falsify

Analyses built on the reader's own portfolio, own transactions or own client list removed the argument about test selection entirely, and in one case the counterparty verified the central claim against their own published testimonials.

THE INSTRUMENT

What the evidence rests on

Every figure in this portfolio was produced without access to a counterparty's systems, without contacting anyone at the companies analysed, and from two data layers that are held separate by design.

SPECIFICATION

The annual panel covers more than four million companies across more than one hundred countries, with 48 million observed company-year records, 1.88 billion skill rows and 46.5 million capability buckets. It runs 1950 to 2020 and is held there. Model training and backtesting require a stationary historical corpus; appending live observations would render the training set non-stationary and contaminate every backtest run against it.

The live enrichment layer runs 2018 to present and refreshes every three to four weeks. The overlap between the two is the calibration window. The layers rest on different observation bases and are not joined, netted or interpolated across. No figure in this portfolio crosses between them, and every entry reporting from both states which layer produced which number.

ESTABLISHED ACROSS THE TWENTY

That a specified commercial question can be answered from observed workforce evidence before the subject discloses anything. That the finding survives inspection by the party best placed to test it. And that the instrument returns negative results, which is the condition under which its positive results are worth acting on.

OUTSIDE THE SCOPE OF A SAMPLE

A universe specified jointly with the counterparty, so the base rate is drawn from their population. Signals validated against their own closed and lost record rather than a public proxy. Row-level source, observation date and confidence, so any single result is independently checkable. And a cohort large enough for a finding to be measured rather than illustrated.

COMPANION DOCUMENTS

Vivameda Research Portfolio 2026 · Five studies, one panel. What the instrument establishes, and where it stops.

Oliver Nold Founder Portfolio · The operator behind the panel, and the decade of owned companies that produced it.

Vivameda Ltd, Cyprus · Commissioned research, scored populations and diligence briefs for investors, acquirers and advisers.
Prepared August 2026. Figures current to the collection dates stated in each entry. Reference VMD.21958.150826.240.