

Research Portfolio

FIVE STUDIES, ONE PANEL

*Predicting growth. Identifying structural vulnerability. Measuring
transformation. Reconstructing long-term organisational readiness.
Detecting capability decay.*

THE FIVE STUDIES

STUDY ONE	The Unicorn Signature	Pre-hypergrowth patterns
STUDY TWO	The Defensible Shelf	Structural business vulnerability
STUDY THREE	The Closed Circuit	Capability and talent-network lock-in
STUDY FOUR	The 1998 Head Start	Long-term organisational readiness
STUDY FIVE	The Portable Asset	Capability decay before failure

INTRODUCTION

Five applications of one dataset

These five studies demonstrate five applications of the same proprietary longitudinal workforce panel: predicting growth, identifying structural vulnerability, measuring transformation, reconstructing long-term organisational readiness, and detecting capability decay before failure. They are not five separate research projects. They are five question types asked of one substrate.

The substrate is the Vivameda Longitudinal Workforce Panel: more than four million companies, 48 million observed company-years, 1.88 billion skill rows and 46.5 million capability buckets across more than 100 countries, spanning 1950 to 2026. For every company and every year, the panel carries an observed workforce size, a role composition and a capability composition, all on one taxonomy. That last point is what makes the panel a research instrument rather than a database: a company of 2013 is directly comparable with another of 2016, and a company that no longer exists is comparable with one that does.

4M+

companies observed

48M

observed company-years, 1950
to 2026

1.88 bn

skill rows on one taxonomy

100+

countries

Why workforce structure answers questions financial data cannot

An annual report gives a headcount. It does not say how many of those people were technical, how many were on the shop floor, how many were in buying. No reporting standard requires that breakdown, and for privately held companies no figure is published at all. Workforce composition is therefore one of the few variables that is observable at company level, comparable across peers, available for firms that have since disappeared, and deep enough in time to test whether a pattern actually preceded an outcome.

That combination is what all five studies exploit. Each one takes a state of the workforce at a point in the past, fixes the classification before looking at what followed, and then checks the outcome. The Unicorn Signature does it with a model trained on companies it is never tested on. The Defensible Shelf does it with signals drawn only from 2012 to 2016 and outcomes read from the public record for 2017 to 2026. The Closed Circuit does it with a scale validated against nine reference companies whose expected values were written down before the query was run. The 1998 Head Start does it with three classes formed from a single base year and a permutation test across 100,000 draws. The Portable Asset does it with a cohort and a classification fixed in writing before the composition was measured.

What connects the five

Read together, the five papers map the same variable across the full life of an organisation. One study asks which companies are structurally ready to scale before the growth curve becomes visible. One asks which are structurally exposed before the market removes them. One asks whether an industry that says it is transforming is actually transforming. One asks what an organisation already was, decades before the question mattered. One asks what a company was made of in the years before it emptied, and where those people went.

The same panel answers all five because the underlying question is identical in each case: what does the composition of a workforce reveal about what an organisation is capable of becoming, and how early does it reveal it. Growth data ranks companies. Workforce structure explains the ranking, and it is visible earlier.

How these papers are written

Each study states its headline finding first, shows the evidence, and then consolidates its limitations into a single section before the appendix. Those sections are not boilerplate. They state the operating mode the result was built for, the alternative explanations that were tested and survived, and the boundary beyond which the reading does not carry. A reader who takes a headline number from this portfolio without reading the corresponding scope section will overstate what the panel shows.

Every classification rule and query behind these studies is reproducible and is disclosed on request. Where a figure comes from a second source with a different collection methodology, the paper says so and does not net it against panel figures.

OVERVIEW

What each study detects

STUDY	WHAT IT DETECTS	RELEVANT DECISIONS
The Unicorn Signature	Pre-hypergrowth patterns	Sourcing, investment screening, account prioritisation
The Defensible Shelf	Structural business vulnerability	Due diligence, credit risk, portfolio review
The Closed Circuit	Capability and talent-network lock-in	Transformation, recruiting strategy, competitive analysis
The 1998 Head Start	Long-term organisational readiness	Resilience, internationalisation, operating model assessment
The Portable Asset	Capability decay before failure	Acquisition diligence, portfolio review, technical-team assessment

OVERVIEW

The five studies in one line each

STUDY I THE UNICORN SIGNATURE

The top 1 per cent of scored companies enter hypergrowth at 27 times the base rate. Of six unicorns held entirely out of training, 100 per cent scored in the top 30 per cent in the year before they scaled.

STUDY II THE DEFENSIBLE SHELF

Of 225 large retailers, 72 are gone. Those with an online-substitutable assortment failed at 42.1 per cent against 7.6 per cent for the rest, a factor of 5.5, and seven workforce signals measured before 2017 all pointed the same way.

STUDY III THE CLOSED CIRCUIT

Across seventy analysed personnel-exchange relationships spanning six companies, not one software house appears. Eight of nine automotive units moved by less than 0.10 index points in the decade they spent describing themselves as software businesses.

STUDY IV THE 1998 HEAD START

Across 41 multinational corporations, those already distributed in 1998 reduced workforce concentration almost ten times as much as concentrated ones over the following twenty years, at an identical rate of entity growth.

STUDY V THE PORTABLE ASSET

Six AI companies, one identical window. The four that later contracted were thinning their engineering share by 2020; the two that persisted were thickening theirs. Six of six in the expected direction.

Reading order

The studies are independent and can be read in any sequence. The order used here runs from the earliest signal in a company's life to the longest horizon: growth readiness, then vulnerability under a channel shift, then capability lock-in inside an existing industry, then structural inheritance across two decades, then capability decay in the years before an organisation empties. The closing section sets out how the same instrument is applied to a specific company list, market, portfolio or acquisition universe.

PREDICTIVE WORKFORCE SIGNAL

The Unicorn Signature

Future unicorns concentrate in the top score bands years before they scale. Reviewing 30 per cent of a universe captures 82 per cent of them.

HEADLINE

The top 1 per cent of scored companies enter hypergrowth at 27 times the base rate

VALIDATION

50 named technology unicorns; of the 6 held entirely out of training, 100 per cent scored in the top 30 per cent at their pre-scaling year

TEST DESIGN

12,285 held-out companies, 90,233 observations, time-safe features, company-level split, AUC 0.889

SECTION 01

01 The result

Most company scoring systems see growth after it is already visible. By the time a headcount curve bends upward sharply, the company is on the radar of every investor, platform and competitor in its market. This model ranks companies before that point, using only data observable at the time.

27x

the base rate for hypergrowth in the top 1 per cent of scores, on companies the model never saw

100%

of the six unicorns held entirely out of training landed in the top 30 per cent at their pre-scaling year

0.889

AUC on a held-out test set of 12,285 companies and 90,233 observations

GROWTH RANKS. WORKFORCE STRUCTURE EXPLAINS.

Headcount and growth tell you what is moving. Workforce structure and capability composition tell you whether the organisation is built for the move. Vivamedata supplies both as a pre-computed signal layer that an AI agent or a scoring pipeline ingests directly.

Four findings

1 HELD-OUT PREDICTION

The top 1 per cent of scored companies hypergrow at 27 times the base rate. That 1 per cent contains 27 per cent of all hypergrowth outcomes in the held-out set. AUC 0.889.

2 UNICORN CONCENTRATION

Of 22 unicorns with a visible pre-scaling year, 82 per cent landed in the top 30 per cent of scores. Of the 6 held entirely out of training, 100 per cent did.

3 INDEPENDENT SIGNAL

A model using workforce composition alone, with no growth, size, industry or age data, predicts hypergrowth at 4.4 times the base rate in the top 1 per cent.

4 EARLY VISIBILITY

The prediction window is three years. The unicorn validation scores companies in the year before their early-scaling pattern fires. The signature is present in workforce structure before headcount expansion is visible.

SECTION 02

02 Validation against named unicorns

Each of the 50 unicorns in the validation cohort was scored at its pre-scaling year, defined as the year before the early-scaling pattern first fired. 22 of the 50 had a pre-scaling year visible in the scoring window. The remaining 28 either entered the panel already in a scaling state or did not have the pattern fire.

Of those 22, 36 per cent landed in the top decile of model scores for their year, 59 per cent in the top 20 per cent, 82 per cent in the top 30 per cent and 86 per cent in the top 50 per cent. Random expectation would be 10, 20, 30 and 50 per cent respectively.

Six of the 22 were held out of the training set entirely; the model never saw their data. In this cleanest subset, 50 per cent landed in the top decile and 100 per cent in the top 30 per cent.




FIGURE 1 Unicorns concentrate in top scores. The held-out cohort (right) is the cleanest validation: the model placed all six in the top 30 per cent at their pre-scaling year.

WHAT THIS BUYS A REVIEWER

Screening only the top 30 per cent of model scores would have surfaced 82 per cent of the unicorns with a visible pre-scaling year, and all six of the held-out unicorns. Random review of the same 30 per cent captures 30 per cent.




FIGURE 2 Capture rate by review depth. The horizontal axis is the share of the universe reviewed, the vertical axis the share of unicorns captured at their pre-scaling year.

Companies that later became unicorns carried, in the year before their measurable scaling event, workforce signatures the model recognised as pre-hypergrowth.

SECTION 03

03 Performance on the held-out set

The unicorn cohort is a showcase. The operating number comes from 12,285 companies the model never saw during training.

The base rate for hypergrowth within three years is 0.41 per cent in the test set. The full model achieves AUC 0.889. At the top 1 per cent of scores, precision is 11.1 per cent, a lift of 27x. At the top 5 per cent, precision is 4.5 per cent, a lift of 11x. At the top 10 per cent, precision is 2.8 per cent, a lift of 7x.

THE OPERATIONAL METRIC

The top 1 per cent of the test set contains 902 observations. 100 of those experienced hypergrowth, out of 366 hypergrowth events in the held-out set. Reviewing 1 per cent of the universe captures 27 per cent of the eventual scalars.




FIGURE 3 Model performance with ablation, held-out test set. Lift is measured against the 0.41 per cent base rate.

Workforce structure carries signal on its own

To isolate the contribution of the workforce layer, a model was trained using only capability composition, role diversity, role coverage, primary-role concentration and capability coverage. No growth rate, no headcount, no industry, no company age. That model achieves AUC 0.658 on the held-out set, with a lift of 4.4x in the top 1 per cent and 2.6x in the top 10 per cent.

A baseline using headcount alone achieves AUC 0.536. Adding growth features reaches 0.884. Adding industry and company age reaches 0.891. The full model, with workforce structure added, is at 0.889. Growth and demographic features carry most of the ranking power; the workforce layer is what explains the ranking, and it stands up as an independent predictor when everything else is removed. Section 06 sets out what follows from the overlap.

SECTION 04

04 Where it is used

Investors

Funnel compression and attention allocation. A reviewer screening the top 30 per cent of scores would have captured 82 per cent of the unicorns with visible pre-scaling years and all six of the held-out unicorns, against 30 per cent under random review. The model concentrates probability rather than guaranteeing outcomes. For a fund facing thousands of inbound companies, that compression is the value.

Accelerators and platforms

Application screening, founder benchmarking and programme selection. Companies in the top 1 per cent of scores enter hypergrowth at 27 times the base rate, at a precision of 11.1 per cent. The signal is sharpest at the top of the distribution and softens progressively below it, which makes ranked review the correct operating mode rather than rejection at a threshold.

AI agents and GTM platforms

Account scoring, research automation and explainable company intelligence. The model logic and pre-computed signals ship as a knowledge file, so an agent can assess scaling readiness for any company in a CRM, portfolio or research universe and state why a company scores where it does. The file works in ChatGPT, Claude, custom agent frameworks and standard RAG infrastructure such as Pinecone, Weaviate and Chroma. It contains the signal definitions, the score interpretation logic and pre-computed signals on the validation cohort; panel-scale signal layers are available under licence.

Typical questions an agent answers once the file is loaded:

- Which companies in my pipeline resemble historical pre-hypergrowth companies?
- Which accounts show structural scaling readiness?
- Which companies have momentum but weak workforce structure?
- Which companies have organisational maturity before the growth curve becomes obvious?
- Why is this company scoring high or low, and which features drive the score?

SECTION 05

05 How it was built

PREDICTION TARGET

Hypergrowth: year-over-year workforce growth above 100 per cent within the next three years. Chosen over unicorn status because valuation events are externally reported, delayed and unevenly captured across regions and sectors. Workforce hypergrowth is observable directly in the panel. Test-set base rate 0.41 per cent.

DATASET

41,000 companies from the Vivameda panel in technology and technology-adjacent industries (computer software, internet, IT services, computer and network security, financial services, real estate, computer games, design, automotive, logistics, health and wellness, entertainment, media production, broadcast media, leisure and travel). 302,000 company-year observations across 2001 to 2016.

SPLIT AND LEAKAGE CONTROL

Split by company, not by year. 70 per cent of companies train the model (28,664 companies, 211,736 observations); 30 per cent are held out (12,285 companies, 90,233 observations). No company appears in both sets. Company identifiers are not used as features. Every feature is computed only from data available at or before year T, which prevents both time leakage and company-level memorisation.

FEATURES

Observed headcount, year-over-year growth rate, previous-year growth, growth acceleration, role diversity, role coverage, capability composition (which of 10 capabilities appear in the top 3), capability coverage, primary role concentration, industry and company age.

MODEL

Gradient boosted tree ensemble, 200 trees, maximum depth 3, learning rate 0.05, trained without class balancing. Four baselines are compared against the full model on the same held-out set, which contains 366 positive hypergrowth outcomes out of 90,233 observations.

SECTION 06

06 Scope of application

Four points define how the model is operated and what it is for. They set the operating mode; they do not qualify the result.

RANKED REVIEW, NOT A THRESHOLD	Precision in the top 1 per cent is 11.1 per cent against a base rate of 0.41 per cent. The 27x figure is a concentration of probability, which is what compresses a review funnel. The model is built for ranked review, not for rejection at a cut-off.
THE TARGET IS HYPERGROWTH	The model predicts more than 100 per cent year-over-year workforce growth within a three-year window. Hypergrowth is observable directly in the panel, which is why it was chosen over externally reported valuation events. It precedes unicorn outcomes without being identical to them.
WHAT THE WORKFORCE LAYER ADDS	Growth and demographic features carry most of the ranking power, at AUC 0.891. The workforce layer holds standalone predictive signal at AUC 0.658 and a top-1 per cent lift of 4.4x, and it supplies the reason a company scores where it does. Ranking and explanation are delivered together; each is usable on its own.
COHORT AND COVERAGE	The 50-unicorn cohort was selected for reaching unicorn status before 2024 with sufficient panel depth, which weights it toward US technology-adjacent companies. It demonstrates that unicorns carried the signature. The operating numbers come from the held-out set of 12,285 companies, not from the cohort.

The full constraint set, the cohort selection rule and every query behind this paper are documented and disclosed on request.

SECTION 07

07 Delivery formats

The Unicorn Signature is one product built on the Vivameda Longitudinal Workforce Panel. The panel is available in four configurations, scaled to the use case.

RAG-READY SIGNAL FILE	The knowledge file drops into ChatGPT, Claude or an agent framework. The agent assesses scaling readiness for any company, using the validated unicorn signatures as analogical evidence. Suited to research workflows and AI-agent products.
CUSTOM SCORING RUN	A company list (CRM, portfolio, research universe) is scored with the Unicorn Signature model and returned as a labelled feature file with model probabilities, decile placements, pattern flags and explanation fields. Suited to production-scale screening on a defined universe.
PANEL-SCALE SIGNAL LICENCE	Pre-computed scores and pattern flags across the full panel of 4 million companies, delivered as a labelled feature file. Suited to ongoing enrichment of a data platform or alternative-data product.
FULL DATASET FOR ML	The underlying longitudinal workforce dataset, licensed for training proprietary models, building signal layers or running custom research at seven decades of time depth. Suited to quantitative funds, AI research teams and academic work.

4M+	48M	1.88 bn	1950 to 2026
companies in the panel	observed company-years	skill rows	coverage window, 100+ countries

APPENDIX A

Cohort, test set and model details

The 50-unicorn validation cohort

Airbnb	GitHub	Samsara
Asana	HashiCorp	Segment
Atlassian	HubSpot	ServiceNow
Box	Instacart	Shopify
Coinbase	Jamf	Slack
Confluent	Klarna	Spotify
CrowdStrike	Lyft	Splunk
Databricks	MongoDB	Square
Datadog	MuleSoft	Stripe
DataRobot	Opendoor	Tanium
DocuSign	PagerDuty	Uber
DoorDash	Pinterest	Unity
Dropbox	Redfin	WeWork
Elastic	Rivian	Workday
Epic Games	Robinhood	Zillow
Fastly	Roblox	Zscaler
Flexport	Salesforce	

Test set summary

90,233 company-year observations across 12,285 companies, none of which appear in the training set. 366 positive hypergrowth outcomes, a base rate of 0.41 per cent. The top 1 per cent of the test set contains 902 observations and 100 positive outcomes, so 1 per cent of reviewed observations captures approximately 27 per cent of all positives in the held-out test set.

Model details

Gradient boosted tree classifier, 200 trees, maximum depth 3, learning rate 0.05, trained without class balancing. All metrics in this paper are from the held-out test set.

Data basis: Vivameda Longitudinal Workforce Panel. 4 million companies, 48 million observed company-years, 1.88 billion skill rows, 46.5 million capability buckets, 100+ countries, 1950 to 2026. Modelling window 2001 to 2016, time-safe features only, company-level train and test split. Vivameda Research, vivameda.com. This analysis is provided for information and does not constitute investment, legal or tax advice.

WORKFORCE DATA 2012 TO 2016 · OUTCOMES 2017 TO 2026

The Defensible Shelf

225 large retailers, ten years, 72 are gone. The outcome depended less on size, age or balance sheet than on a single question.

CORE FINDING

Retailers with an online-substitutable assortment failed at 42.1 per cent, the rest at 7.6 per cent. A factor of 5.5

DESIGN

Signals drawn exclusively from the window 2012 to 2016 and ranked without knowledge of the outcome. Outcomes from the public record 2017 to 2026

WHAT IT CHANGES

Of the 225 retailers, 72 no longer exist. Their workforce structure before the event survives only in the panel

SECTION 01

01 Size did not decide it. The assortment did.

For 225 large retailers we reconstructed how their workforce was composed between 2012 and 2016. Then we looked at what happened afterwards.

The cohort was split by a single criterion: was the core assortment in 2016 substitutable online without a major physical obstacle. Perishable goods, immediate need, prescription dispensing and bulky goods with a collection channel do not count as substitutable.

42.1%

assortment substitutable online
67 of 159 failed

7.6%

assortment not substitutable online
5 of 66 failed

5.5x

the difference in failure probability between
the two groups

Five and a half times. Whoever was exposed to the channel shift disappeared with five and a half times the probability. The question was never whether a retailer was well run. The question was whether its shelf remained defensible. The complete cohort, all 225 names and all 72 outcomes, is in the appendix.

SECTION 02

02 Seven signals, measured before 2017

All of them point in the same direction. The 159 retailers with an online-substitutable assortment, split into 92 later survivors and 67 later failures. Every value comes from the window 2012 to 2016, that is before any event. Reported is the median of both groups.

EARLY INDICATOR 2012 TO 2016	SURVIVORS	FAILED	DISCRIMINATION
Workforce development 2012 to 2016	+5.6 %	-3.4 %	0.61
Technical share, change	+0.6 pp	+0.3 pp	0.59
Technical share 2016	3.1 %	2.3 %	0.59
Software capability 2016	1.0 %	0.9 %	0.59
Size 2016 (records)	2,510	2,012	0.57
Data capability 2016	2.1 %	2.0 %	0.57
Floor area per technical role 2016	13.0	16.6	0.57

Scale from 0.50 (coin toss) to 0.75 (clear separation). A value of 0.61 means: draw one later survivor and one later failure blindly, and the survivor comes out ahead in six cases out of ten.

Seven metrics, seven times the same sign: the later survivors grew, built technical capability and carried less floor area per technical role. None of them separates decisively on its own. Taken together they form a pattern that runs through the cohort.

What happens when all seven are combined

The 159 exposed retailers were sorted into five equally sized groups by their combined 2016 signal value. The sorting was done without knowledge of the outcome. Only afterwards did we look at who disappeared between 2017 and 2026.

RANK GROUP BY 2016 SIGNAL VALUE	SHARE THAT LATER FAILED
Rank 1 to 32 best signal values	34.4 %
33 to 64	31.2 %
65 to 96	40.6 %
97 to 128	50.0 %
129 to 159 worst signal values	54.8 %

159 exposed retailers, five groups of 32. Shown is the share that later failed.

In the group with the best 2016 signal values a good third failed; in the group with the worst, more than half. The increase runs in the same direction across all five groups. The distance here is a factor of 1.6, against a factor of 5.5 for the assortment. Anyone who wanted to know in 2016 who was at risk had to look at the shelf first and at the workforce only after that. Section 06 states how far this ranking carries.

SECTION 03

03 By 2016 the lead was already built

Share of technical roles in the total workforce, as at 2016. Four years before the gap became visible in industry figures, it was finished in the personnel records.

RETAILER	TECHNICAL SHARE 2016
Amazon	32.0 %
Sears	8.1 %
Costco	6.1 %
The Home Depot	5.6 %
Target	4.9 %
Macy's	4.1 %
Best Buy	3.4 %
JCPenney	2.9 %
Walmart	2.7 %
Rite Aid	1.7 %

Copper: companies that filed for insolvency or were liquidated by 2026.

In 2016 Amazon employed proportionally six times as many technical roles as Target and twelve times as many as Walmart. No annual report disclosed this. It stood in the workforce. Within the established retailers, however, this share did not separate survivors from failures: it describes the distance to the attacker, not the fate of the attacked.



FIGURE 1 Share of technical roles in the workforce, start and end values for 2010 and 2016, on a uniform role taxonomy. Scale 0 to 9 per cent. The dagger marks later insolvency or liquidation. Amazon lies outside the scale.

Three patterns, and one of them does not fit. Nordstrom, Best Buy and Target built continuously and survived. JCPenney reduced and failed. Toys R Us stayed flat and failed. Sears, by contrast, had the highest technical share of all classical retailers in 2016, had raised it by half since 2010, and was liquidated anyway. Part of that rise comes from everything else shrinking faster than the technical roles.

Amazon lies outside the scale. Between 31.6 and 33.4 per cent, stable across all seven years. Amazon did not have to rebuild. The company was already built that way before the competition started measuring.

Four years in which the ranking turned

RETAILER	2012	2016	CHANGE	OUTCOME TO 2026
Amazon	22,139	52,900	+139 %	survived
Ulta Beauty	5,088	8,773	+72 %	survived
Sephora	5,893	7,923	+34 %	survived
Costco	12,001	13,399	+12 %	survived
Target	72,486	56,779	-22 %	survived
Best Buy	41,730	30,299	-27 %	survived
Sears	22,461	15,928	-29 %	insolvency 2018, liquidation
JCPenney	21,907	15,458	-29 %	insolvency 2020
Rite Aid	10,359	9,490	-8 %	insolvency 2023 and 2025

In 2012 Target was three times the size of Amazon. By 2016 the gap was almost closed. Reported are observed workforce records in the panel, not reported corporate employment.

SECTION 04

04 It is happening again. This time around AI.

The following positions were open in July 2026. All of them at retailers that survived the first upheaval. The same companies are simultaneously looking for floor staff at minimum wage and platform engineers at six times that.

COMPANY	POSITION	STATUS	ANNUAL SALARY, USD
Macy's	Principal Engineer, Data Engineering	posted 24 July 2026, repeat	180,600 to 300,900
Macy's	Principal Engineer, MarTech	posted 7 July 2026, twice	180,600 to 300,900
Best Buy	Senior Engineering Manager, Data Store Platforms	posted 10 July 2026	121,000 to 216,600
Target	Senior Engineering Manager, Advanced AI	posted 14 July 2026	not disclosed
Best Buy	Manager, Yield and Inventory Analytics	posted 8 July 2026, repeat	90,200 to 160,000
Macy's	Lead BI Developer, Stores	open since 22 June 2026, repeat	90,100 to 150,300
The same firms	Retail Sales Associate, shop floor	ongoing, all three companies	15.00 to 20.57 per hour

Three of the six technical positions are marked as repeats, meaning they had already been posted once and were not filled. A department store chain that is currently closing branches offers up to 300,900

dollars for a data engineer and does not get one. The capability that made the adaptation lead visible in 2016 is no longer simply purchasable on the labour market today.

THE SURROUNDING FRAME

31 billion dollars of AI spending in retail for 2026, against 7.3 billion four years earlier. 74 per cent of retail executives intend to roll out AI at scale within twelve months. Home Depot installed a new chief technology officer with an AI mandate in April 2026; Marks and Spencer is appointing a chief digital and technology officer for the first time in the summer of 2026.

The first upheaval took ten years to remove 72 of 225 retailers from the market. The second began in 2024.

SECTION 05

05 Why this cannot be reconstructed anywhere else

The question is: how was Sears composed in 2013? It is not a curiosity. It is precisely the question one would have had to ask in 2013 in order to see something five years before the liquidation. Sears Holdings filed for insolvency in 2018 and no longer exists.

SOURCE	WHAT IT ANSWERS TO THE QUESTION	VERDICT
Annual report 2013	A total figure for the group. No breakdown by function or capability.	partial
Professional network today	Company page without an active workforce. Profiles show the state of 2026, not of 2013.	not available
Web archive	Snapshots of web pages. No structured workforce composition.	not available
Vivameda panel	15,943 observed records. 1,082 technical roles, 6.79 per cent. Split across 11 role and 15 capability classes. Comparable with any other company of the same year.	complete

The difference is not in the quantity but in the structure. An annual report gives a headcount. It does not say how many of those were technical, how many were floor staff, how many were buying. And it exists only for listed companies. Of the 72 failed retailers in this study, several were privately held at the time of the upheaval and never published a workforce figure at all.

None of this can be collected after the fact, for four independent reasons. Profiles are overwritten rather than versioned: when someone updates their title, the old one is gone. Company pages disappear with the company; at Sears, Toys R Us, Bed Bath & Beyond and RadioShack, liquidation also ended the anchor that employment relationships hang on. No reporting standard requires a breakdown of the workforce by function or capability, and where research and development figures exist, they are costs rather than people. And a substantial part of retail was privately held or owned by investment firms and published no personnel figure in any year.

The panel recorded these states continuously rather than reconstructing them later, and brought them onto a uniform role and capability taxonomy. Only that standardisation makes a company of 2013 comparable with another of 2016, and a defunct one with a surviving one.

PUBLICLY REPRODUCIBLE

- Outcomes 2017 to 2026 from court records and reporting, documented per company with date and event type
- Assortment assignment based on the 2016 business model, rule and individual decisions in writing
- Current job postings from July 2026, every figure individually verifiable
- Industry metrics on AI spending, investment intentions and board appointments

The most laborious of these parts, the event assignment for 225 companies across ten years, costs a practised analyst about a week. It is reproducible.

EXCLUSIVELY FROM THE PANEL

- Role composition per company and year, eleven role classes, window 2000 to 2020
- Capability composition per company and year, fifteen capability classes, window 1990 to 2020
- Observed personnel stock per company and year, window 1950 to 2020
- Comparability across companies and years, the same taxonomy for 305,107 companies
- States of defunct companies: of the 225 retailers, 72 no longer exist

Around twelve thousand individual values were read for this analysis. The same reconstruction is possible for any industry and any competitor, backwards to 1950.

4M+	48M	1.88 bn	305,107
companies in the panel	observed company-years	competency entries	companies with a capability layer

SECTION 06

06 Scope of application

STRUCTURAL RISK, NOT AN EVENT FORECAST	The statement is exposure, not date. Individual companies are not predicted, and between signal and outcome one to ten years elapsed in this cohort. The data describe preconditions rather than causes: Sears held the highest technical share of all classical retailers in 2016 and was liquidated anyway.
THE ASSORTMENT RANKS FIRST, THE WORKFORCE SECOND	The assortment criterion separates by a factor of 5.5; the combined workforce signal separates by a factor of 1.6 across five rank groups, in the same direction throughout. The correct order of review is therefore the shelf first and the workforce after it, and the workforce signal is a ranking within an exposed segment rather than a verdict on a name.
WHERE IT APPLIES AND WHAT IS COUNTED	For groceries, convenience, pharmacy and home improvement the relationship is explicitly not demonstrated; that was measured, not assumed. Panel figures are observed workforce records on a uniform taxonomy, not reported corporate employment: they show structure and direction. The cohort comprises retailers with at least 750 observed records each in 2016.

SECTION 07

07 Running this against a portfolio

The cohort in this paper was built from a segment definition and a written threshold. Given a list of holdings or borrowers instead, the same construction runs against it.

Each name is assigned against a stated criterion before any outcome is examined, exactly as the assortment rule was applied here, and that split is checkable line by line. The seven signals then run per company and year on one taxonomy: workforce trajectory, technical share and its change, software and data capability, size, and floor area per technical role, combined into a single ranked value. Every holding is measured against its own past and against peers that no longer exist, which is the comparison a current data source cannot supply. What comes back is a labelled feature file, one row per name with the measured fields, the exposure flag and the explanation columns, or a written analysis in the form of this paper.

APPENDIX A

The cohort, complete

225 retailers. Copper marks insolvency, liquidation or wind-down between 2017 and 2026, with the year of the event.

7-Eleven	Boscov's	Ethan Allen Global	IKEA
AAFES	Brooks Brothers 2020	Express 2024	Inditex
Abercrombie & Fitch	Brookshire Grocery	Family Dollar 2025	Ingles Markets
Academy Sports + Outdoors	Brookstone 2018	Fanatics	J.Jill
Ace Hardware	Build-A-Bear Workshop	Festival Foods	J.Crew 2020
Advance Auto Parts	Burlington Stores	Finish Line	JCPenney 2020
Aeropostale	Cabela's	Five Below	Jo-Ann 2024
Ahold	Camping World and Good Sam	Floor & Decor	Jos. A. Bank Clothiers 2020
Albertsons	CarMax	Follett Higher Education	Journeys
Aldi USA	Carter's OshKosh B'gosh	Foot Locker	Kirkland's
Aldo 2020	Cato	Forever 21 2019	Kohl's Department Stores
Amazon	Century 21 Department Stores 2020	Francesca's 2020	Kroger
American Eagle Outfitters	Charlotte Russe 2019	Fry's Electronics 2021	Kwik Trip
Ann Inc. 2020	Charming Charlie 2017	FTD 2019	L Brands
Ascena Retail Group 2020	Chico's FAS	GameStop 2025	L.L.Bean
AutoZone	Circle K	Gander Mountain 2017	Lids Sports
B&H Photo Video	Claire's 2018	Gap	Lord & Taylor 2020
Barnes & Noble	Conn's HomePlus 2024	Giant Eagle	Love's Travel Stops
Barneys New York 2019	Cost Plus World Market	GNC 2020	Lowe's Companies
Bass Pro Shops	Costco Wholesale	Golfsmith International	Lowe's Foods
Batteries Plus Bulbs	Crate and Barrel	Gordmans 2017	Lululemon Athletica
BCBGMaxazriagroup 2017	David's Bridal 2018	Guitar Center 2020	Macy's
Bealls	Defense Commissary Agency	Gymboree 2017	Mattress Firm 2018
Bed Bath & Beyond 2023	Destination Maternity 2019	H&M	Meijer
Belk 2021	Dick's Sporting Goods	Harbor Freight Tools	Men's Wearhouse 2020
Best Buy	Dillard's	H-E-B	Menards
Big 5 Sporting Goods	Dollar General	hhgregg 2017	Micro Center
Big Lots 2024	Dollar Tree	Hibbett Sporting Goods	Modell's Sporting Goods 2020
Big Y World Class Market	DSW	Hobby Lobby	Navy Exchange Service
BJ's Wholesale Club	Eddie Bauer 2026	Hot Topic	Command
Bloomingdale's	Edible Arrangements	HSN 2026	Nebraska Furniture Mart
Bluestem Brands 2020	Empire Today	Hudson's Bay 2025	Neiman Marcus 2020
Books-A-Million		Hy-Vee	New York & Company 2020

Nordstrom	Restoration Hardware	Stein Mart 2020	Trader Joe's
Nordstromrack.com	Rite Aid 2023	Stitch Fix	TravelCenters of America
O'Reilly Auto Parts	Ross Stores	SuperValu	True Value 2024
OfficeMax	rue21 2017	Sur La Table 2020	Trunk Club 2020
Orchard Supply Hardware 2018	Saks Fifth Avenue 2026	Talbots	Tuesday Morning 2020
Overstock.com	Sally Beauty	Target	Ulta Beauty
Pacific Sunwear	Savers	The Body Shop 2024	Uniqlo
Paper Source 2021	Scheels	The Bon-Ton 2018	Urban Outfitters
Party City 2023	Schnuck Markets	The Buckle	Von Maur
Payless ShoeSource 2017	Sears Holdings 2018	The Children's Place	Wakefern Food Corp.
Pep Boys	Sephora	The Container Store 2024	Walgreens
Petco Animal Supplies	Sheetz	The Disney Store	Walmart
PetSmart	Shoe Carnival	The Fresh Market	Wawa
Pier 1 Imports 2020	Shopko 2019	The Home Depot	Wayfair
Pilot Flying J	Signet	The TJX Companies	Wegmans Food Markets
Price Chopper Supermarkets	Sleep Number 2026	The Vitamin Shoppe	Weis Markets
Publix Super Markets	Smart & Final	Things Remembered 2019	West Marine
QuikTrip	Southeastern Grocers 2018	Tiffany & Co.	Wet Seal 2017
QVC 2026	SpartanNash	Tillys	Whole Foods Market
RaceTrac Petroleum	Speedway	Tops Markets 2018	Williams-Sonoma
RadioShack 2017	Spencer Gifts	Torrid	WinCo Foods
Raley's	Sports Authority	Total Wine & More	Zulily 2023
Raymour & Flanigan	Sprouts Farmers Market	Toys R Us 2017	Zumiez
REI	Stage Stores 2020	Tractor Supply	

Data basis: longitudinal workforce panel, 4 million companies, 48 million observed company-years, 1.88 billion skill rows, 100+ countries, 1950 to 2026. Historical signals exclusively from the window 2012 to 2016. Outcomes from the public record 2017 to 2026. Job posting figures July 2026, individually verifiable. Cohort: 225 retailers, each with at least 750 observed records in 2016. Selection rule, segment assignment and individual evidence fully available for inspection on request.

WORKFORCE ANALYSIS

The Closed Circuit

Across seventy analysed exchange relationships spanning six companies, not one software house appears. That explains twenty years of standstill behind twenty years of transformation rhetoric.

CORE FINDING

The labour market both manufacturers and dealers recruit from is closed. 77 to 92 per cent of personnel exchange takes place inside the industry

SCOPE

Six manufacturers, four dealer groups, 2,325 classified job postings, historical series 2000 to 2020

WHAT IT CHANGES

The question is no longer who catches up faster, but who leaves the closed circle first

SECTION 01

01 The finding

Between 2000 and 2010 the German automotive groups raised their software capability relative to the market. Since 2010 that movement has stood still. The standstill begins precisely in the decade in which the industry started to describe itself publicly as a software business.

0 of 70

exchange relationships across six companies lead to a software house

8 of 9

units change by less than 0.10 index points between 2010 and 2020

0.88

Volkswagen's software index in 2020 against 5.27 for SAP, where the market average is 1.00

Only the two tier-one suppliers achieved a clear lead over the market as a whole. BMW and Daimler sit just above it. The three Volkswagen Group brands did not reach the market average in a single year. Today those same manufacturers post a quarter to a third of their vacancies in the digital field; the dealer groups post three to eight per cent.

THE ONE NUMBER

In 2020 Volkswagen AG reached a software share of 0.88 relative to the panel mean. Thyssenkrupp stood at 0.56, BASF at 0.27, SAP at 5.27. An automotive group was therefore closer to a chemicals group than to a mechanical engineering firm, and a factor of six away from a software house.

The common narrative holds that the manufacturers have completed the transition and the dealer trade is lagging behind. This study does not support it. It shows an industry that fixed its self-image early and then confirmed rather than rebuilt it for two decades, and that is only now beginning the rebuild in its job postings. The dealer trade has not begun at all. In both cases the reason is the same, and it is in the next section.

SECTION 02

02 The mechanism: the closed circle

For every company it was recorded which firms the workforce comes from and which firms it goes to. The ten strongest inflows and outflows in each case were analysed.

COMPANY	FIRMS IN EXCHANGE	OF WHICH AUTOMOTIVE SHARE	OUTSIDE THE INDUSTRY
Bosch	12	11 92 %	Siemens
AUDI AG	10	9 90 %	Siemens
AMAG Group	16	13 81 %	Swisscom, armed forces, AXA
Emil Frey Switzerland	15	12 80 %	armed forces, AXA, Credit Suisse
Porsche Holding	14	11 79 %	EY, Aldi Süd, Red Bull
BMW Group	13	10 77 %	TU Munich, Siemens




FIGURE 1 Share of exchange relationships that take place within the automotive industry. The ten strongest inflows and outflows per company were analysed.

Across seventy analysed exchange relationships spanning six companies, not one software house appears.

No SAP, no Microsoft, no digital consultancy, no technology group. The only technology-adjacent address is Siemens, and Siemens sits at index 1.16 in the same measurement, barely above the market average.

Audi posts 29.9 per cent of its positions in the digital field. Ninety per cent of the people who fill those positions come from within the Volkswagen Group, Audi Hungary, Audi China, CARIAD, Porsche, Mercedes-Benz, SEAT and BMW. An industry that recruits its software capability from itself can rename positions without any new capability entering the system.

For the dealer trade this turns the question around. It is not behind because it is slower. Both sides sit in closed circles, simply in different ones. AMAG exchanges staff with Emil Frey, Emil Frey with AMAG, both additionally with Binelli, Franz AG, Grosspeter and the Swiss armed forces. Audi exchanges within its own group. A dealership that wants to build a digital function is therefore searching in a set in which the capability sought does not occur. The problem is not speed. It is visibility in the wrong market.

SECTION 03

03 Twenty years without movement

The table below shows, for every unit, the software share divided by the panel mean of the same year. An index of 1.00 means the market average, measured against the whole economy rather than against the industry.

Only company-years with at least 100 observed people are reported; this threshold was set before the analysis. The analysis window begins in 2000, and section 07 explains why the available data for the 1990s were discarded.

UNIT	2000	2010	2020	DELTA 2000 TO 2010	DELTA 2010 TO 2020
Bosch Engineering, India	n.a.	7.07	5.11	n.a.	-1.96
Bosch India	n.a.	n.a.	1.80	n.a.	n.a.
Bosch, Germany	0.65	1.57	1.67	+0.92	+0.10
Continental	0.67	1.46	1.51	+0.79	+0.05
Daimler AG	0.59	1.17	1.20	+0.58	+0.03
BMW Group	0.66	1.09	1.12	+0.43	+0.03
Volkswagen AG	0.74	0.79	0.88	+0.05	+0.09
Porsche	n.a.	0.38	0.82	n.a.	+0.44
Audi AG	n.a.	0.58	0.55	n.a.	-0.03
Schaeffler	0.18	0.34	0.43	+0.16	+0.09
Bosch Rexroth	0.51	0.33	0.43	-0.18	+0.10

"n.a." means the unit did not reach the minimum base of 100 observed people in that year.




FIGURE 2 Software index for the units with a continuous series, 2000, 2010 and 2020. The entire movement sits in the first decade.

Four observations

First. In 2000 every single German automotive unit in this set sat below the market average. By 2020 Bosch and Continental had reached a clear lead, with BMW and Daimler just above the line. Volkswagen, Audi and Porsche did not reach it in any year.

Second, and this is the actual finding. The entire movement sits in the first decade. Between 2010 and 2020 practically nothing changes for eight of nine units: the range runs from minus 0.03 to plus 0.10 index points across ten years. The only exception is Porsche at plus 0.44, and that series begins only in 2007. The decade of standstill is also the decade in which software capability became the dominant theme of industry communication. In the observed workforce composition that decade left no trace.

Third. Such build-up as exists did not happen in Germany. Bosch's Indian captive unit stood at index 7.07 in 2010, that is at the level of Microsoft, and it did so a decade before the public debate about software capability in the automotive industry began. Since then it has fallen by almost two index points.

Fourth, and this is the cleanest counter-example. Bosch and Bosch Rexroth belong to the same house and run in opposite directions. Bosch rises from 0.65 to 1.67, Rexroth first falls from 0.51 to 0.33. Anyone wishing to dismiss the finding on the grounds that it reflects the industry or the period has to explain why two units of the same group diverge over the same span.

SECTION 04

04 Why the scale holds up

Before an index value is interpreted, it must be established that the scale scales sensibly. Nine reference companies outside the automotive industry were therefore measured. The expectation was recorded in writing before the query was run: software houses well above index 3, consulting in between, chemicals and retail clearly below 1.




FIGURE 3 Software index 2020 for the nine reference companies. Siemens, in copper, is the only deviation from the expectation recorded in advance.

REFERENCE	2020 SOFTWARE SHARE	INDEX	EXPECTATION MET
Microsoft	55.6 %	7.21	yes
Oracle	49.5 %	6.42	yes
SAP	40.7 %	5.27	yes
Accenture	28.4 %	3.68	yes
Siemens	8.9 %	1.16	no, see below
Thyssenkrupp	4.3 %	0.56	yes
Walmart	3.6 %	0.47	yes
Lufthansa	3.1 %	0.41	yes
BASF	2.1 %	0.27	yes

Eight of nine references sit in the expected range. The one deviation is Siemens, and it is the most revealing line of the study. Siemens falls from index 2.05 in 2000 to 1.16 in 2020, where a value between 1.5 and 3 had been expected. The decline coincides with the sale of Siemens IT Solutions to Atos. The measure therefore captured an externally documented corporate event that nobody planted in the data and that the expectation did not contain.

A scale that makes a real corporate divestment visible without anyone having looked for it is not measuring noise.

SECTION 05

05 The present position, 2026

For the present, all open positions at the companies under study were collected and classified into four substantive fields under identical rules fixed in advance. Contract form was recorded as a separate and independent dimension. 2,325 open positions were collected.

COMPANY	POSTINGS	ENTRY LEVEL	IT AND DIGITAL	WORKSHOP, PRODUCTION	SALES, MARKETING
AUDI AG	97	94.8 %	29.9 %	14.4 %	1.0 %
BMW Group	1,000	75.2 %	25.2 %	7.2 %	8.0 %
Bosch	793	56.0 %	18.0 %	10.8 %	10.7 %
Porsche Holding	209	13.4 %	8.1 %	34.0 %	20.6 %
AMAG Group	147	14.3 %	5.4 %	32.0 %	22.4 %
Emil Frey Switzerland	66	0.0 %	3.0 %	43.9 %	16.7 %

The Weller group was collected but not reported because of an insufficient case count. At 13 open positions, a single role shifts the result by eight percentage points.




FIGURE 4 Share of open positions in IT and digital against workshop and production, as of August 2026.

Emil Frey is looking for workshop staff in 43.9 per cent of its postings and for digital capability in 3.0 per cent. BMW is looking for workshop staff in 7.2 per cent and for digital capability in 25.2 per cent. The two sides of the market are not filling the same rooms more slowly or more quickly, they are filling different rooms. Within the dealer trade the spread is considerable as well: between Emil Frey at 3.0 per cent and Porsche Holding at 8.1 per cent lies a factor of 2.7. The dealer trade does not stand still as a block.

WHAT THE DEALER TRADE LOOKS FOR WHEN IT LOOKS DIGITAL

AMAG's seven IT positions in plain terms: big data and platform engineer, SAP CAP developer, product owner SAP integration, product manager SAFe, IT operations manager, network security engineer, senior network and applications specialist. Two of these are ERP integration, two are network operations, one is security. That is the professionalisation of one's own IT, not the building of a digital business. At Emil Frey there are two positions: a network security engineer and a system engineer focused on M365. Both operations, none of it build-up.

ONE NUMBER FROM THE LINE BEHIND

AMAG posts 21 apprenticeship places and seven IT positions. A ratio of three to one in favour of the classical skilled-trades base. This is not a criticism. It is a statement about priority, and it is unambiguous.

SECTION 06

06 The structural bottleneck

Organisations do not change by posting new capabilities. They change only when they leave their existing talent circuit. The decisive variable is therefore not the number of digital positions but the origin of the people who fill them.

The bottleneck does not sit at the question of whether a dealership wants to digitalise. The postings show that several are trying. Nor does it sit at the budget. An organisation whose personnel exchange takes place four fifths within its own industry does not have a recruitment problem but a visibility problem in the wrong market. It can describe the position correctly, pay appropriately and still fail to fill it, because the candidates never see it. And when it is filled, it is filled by someone from the same circle.

WHAT THIS MEANS FOR THE MANUFACTURERS

The manufacturers are now catching up on what did not happen in the decade to 2020. A quarter to a third of their postings are digital. Whether capability results from this will be decided not by the number of positions but by whether the talent circuit opens. So far it is closed.

WHAT THIS MEANS FOR THE DEALER TRADE

The dealer trade stands at a fifth to a tenth of the manufacturer level, and its movement has not begun. The spread within the trade, a factor of 2.7 between the weakest and the strongest group, nevertheless shows that this is not an industry inevitability. It is a decision, and some have already taken it differently.

This study measures six manufacturers and four dealer groups. It says nothing about whether car retail lags other retail formats or whether all bricks-and-mortar retail shows the same pattern. That question is answerable with the same instrument and would be the next step.

SECTION 07

07 Scope and measurement limits

TWO LAYERS, NEVER NETTED

The panel layer and the live layer are never netted against one another. At a large corporation the panel layer captures roughly two per mille of the workforce, while at a dealer group the live layer captures roughly forty per cent. A share drawn from a thin sample is systematically higher than one drawn from a dense reading, so historical and present-day values stand side by side in this document and never inside one calculation.

OBSERVATION DEPTH IS A SAMPLE

For Volkswagen the panel layer captures 1,164 people in 2020 against roughly 670,000 real employees. Anyone who lands in such a sample is above-averagely likely to maintain a professional profile and to work in an office role, so the measured software share of the early years runs high rather than low and the observed increase flat rather than steep. Where the curves rise, the result is conservative.

THE INDEX IS A RATIO, NOT A COMPOSITION

The capability buckets overlap, because one person counts in several, and their shares sum to roughly 330 per cent across all buckets. Admissible are the movement within one bucket over time and the ratio to the panel mean, and that is exactly how they are used here.

THE EXCHANGE IS A TOP-TEN LIST

The ten strongest inflows and outflows per company were analysed. They show the centre of the personnel flow, not its edges. The supported statement is that software does not appear at the centre of the exchange, not that no exchange exists at all. Assignment to the automotive industry was made on company names; it is traceable, and it is an assignment rather than a measurement.

Four further measurement constraints, the treatment of the two corporations missing from the present-day layer, and the null results from approaches that were tested and rejected are documented in the method protocol and disclosed on request.

An industry that recruits its software capability from itself can rename positions without any new capability entering the system.

SECTION 08

08 Running this against a named competitor set

The construction behind this study is not specific to the automotive industry. It takes a list of names and returns the same three measurements against them.

THE CAPABILITY INDEX	Software, data or any of the panel's capability areas per company and year, divided by the panel mean of the same year, from 2000 onward and against a stated minimum observation base. That is the series in section 03, built for any company or unit in the panel.
THE TALENT CIRCUIT	Where a company's people come from and where they go, ranked by flow, with each counterparty resolved and classified. That is the measurement in section 02, and it is the one that separates a stated capability from an acquired one.
THE PRESENT-DAY DEMAND	Open positions collected and classified under a rule set fixed in advance, with contract form recorded as an independent dimension. That is section 05, and it shows what an organisation is currently trying to become.

Given a list of manufacturers, suppliers, dealer groups or any competitor set, the same three measurements run against it, on the same taxonomy, with the classification rules fixed before the result is examined. The cohort in this paper was built exactly that way, and nothing about the construction is specific to the names in it.

What comes back

A LABELLED FEATURE FILE	One row per company and year, with the capability index against the panel mean, the observation base behind each value, the ranked exchange counterparties with their classification, and the present-day posting mix. Every field is named and every threshold is stated in the file itself.
A WRITTEN ANALYSIS	The same measurements in the form of this paper: a finding, the evidence behind it, the scale validation, and the scope section that states where the reading holds. Suited to a board, an investment committee or a diligence file.
A REPEAT READING	The present-day layer is a snapshot by construction, so the posting classification can be re-run at a fixed interval against the same rule set. What that produces over four quarters is a movement rather than a position, which is the measurement that answers whether a talent circuit is actually opening.

Data basis: Vivameda panel layer, capability layer 1990 to 2020, analysis window 2000 to 2020, minimum base of 100 observed people per company-year, roughly 306,000 companies with a capability layer. Vivameda live layer as of August 2026, 2,325 classified job postings across seven companies, identical rule set fixed in advance. Scale validation against nine reference companies with expectations recorded in advance. All classification rules and queries are reproducible and are disclosed on request.

LONGITUDINAL WORKFORCE ANALYSIS

The 1998 Head Start

A corporation holding its entire visible workforce in one entity in 1998 stood in the same place in 2018. Structure is not a decision taken when the need arises.

HEADLINE

Corporations already distributed in 1998 reduced concentration almost ten times as much as concentrated ones, at an identical rate of entity growth

EVIDENCE

41 multinational corporations, 1994 to 2020. Permutation test across 100,000 draws, $p = 0.0015$. The finding holds in all seven twenty-year windows

WHAT IT CHANGES

The base year predicts the end year better than anything a corporation decided or announced in between

SECTION 01

01 The result

Since 2020 every consultancy has been selling resilience. De-risking, China plus one, nearshoring, supply chain diversification. The shared premise of these programmes is that structure is a decision you make once the need arises. This study tests that premise on 41 multinational corporations across 1994 to 2020, and it runs against it.

–17.4%

relative reduction in workforce concentration among already distributed corporations, 1998 to 2018

–1.8%

the same measurement among concentrated corporations, at an identical rate of entity growth

$p = 0.0015$

permutation test, 100,000 random reassignments, one-sided

Corporations did not distribute their workforce because they grew. They distributed it because they were already distributed. Two decades of globalisation, financial crisis and supply chain realignment changed that structure almost not at all.

Workforce structure follows a path. It does not come into being at the moment it is needed.

What is measured is the share of a corporate family's observed total workforce held by its largest entity. Multinationals appear in the panel not as one row but as a family of legally separate entities: Bosch, for instance, comprises 344 entities across 43 countries, with observations from 1963 onward. Each family is compared against itself between two points in time, which removes every compositional effect that could distort a cohort comparison. Section 06 sets out what this measure is and is not.

SECTION 02

02 The evidence

41 corporations were drawn from 848 anchor firms with at least 2,000 observed employees between 1995 and 2005. Armed forces, government bodies, purely domestic retail chains and every name whose attribution would have been too imprecise were excluded. The cohort is listed by name in Appendix A.

Three classes were formed from the base year alone, on thresholds fixed before analysis: distributed (1998 concentration below 80 per cent, 10 corporations), partly distributed (80 to 95 per cent, 9 corporations) and concentrated (95 per cent and above, 22 corporations). The classification uses 1998 data exclusively and no information about what followed.

CLASS	N	CONCENTRATION 1998	MEDIAN CHANGE	RELATIVE	STRUCTURAL GROWTH
A distributed	10	62.5 %	-9.9 pp	-17.4 %	0.63
B partly distributed	9	88.2 %	-3.6 pp	-4.1 %	0.52
C concentrated	22	98.2 %	-1.7 pp	-1.8 %	0.56

Concentration is the share of the largest entity in the observed workforce of the corporate family. Structural growth is the mean logarithmic increase in entity count.

Two obvious alternative explanations are ruled out by the data. The first is differential growth: all three classes grew structurally at the same rate, 0.63, 0.52 and 0.56. The second is a ceiling effect, the objection that a corporation at 99 per cent has less room to fall than one at 40 per cent. The correlation between base level and change across all 41 corporations is $r = -0.01$, effectively zero. The effect attaches to class membership, not to available headroom.




FIGURE 1 Relative change in concentration per corporation, 1998 to 2018, by base class. Each point is one corporation, the vertical mark is the median. The classification derives exclusively from 1998.

The difference holds statistically: a permutation test across 100,000 random reassignments yields $p = 0.021$ for the absolute and $p = 0.0015$ for the relative measurement.

SECTION 03

03 It holds in every window

A result that holds only for one particular pair of years is worthless. The measurement was repeated across seven consecutive twenty-year windows. It holds in every one, with a ratio between 4.6 and 9.8 in favour of the distributed class. The originally chosen window, 1998 to 2018, is among the weaker ones.




FIGURE 2 Median change in concentration across seven twenty-year windows, by base class. The cohort comprises the same 41 corporations in every window.

The individual corporations

Seven of the ten distributed corporations distributed further, three concentrated. In the concentrated class, 20 of 22 moved toward distribution, though only marginally in each case, because a base value of 98 per cent leaves no room.




FIGURE 3 Change in concentration per corporation, 1998 to 2018, sorted. Colour indicates the base class.

The three counter-cases

Infosys (plus 26.2 points), Bayer (plus 11.3) and Schneider Electric (plus 10.4) moved against the finding, and all three are explicable. Infosys was not a global corporation in 1998 but an Indian services firm still forming; concentration rose because a centre came into existence at all. Bayer integrated a major acquisition from 2018. Schneider Electric built a central group structure over the same period where grown national subsidiaries had stood before.

READING THE COUNTER-CASES

These cases sharpen the finding rather than weakening it. Concentration is neither a law of nature nor a one-way street. It is the accumulated result of decisions taken over decades, and it moves in either direction when a structural occasion arises.

SECTION 04

04 A 2026 snapshot

The historical panel layer ends in 2020. To see what the result looks like today, an independent sample was drawn for one corporation from a second source: 100 Bosch entities with current workforce figures, as of June 2026.




FIGURE 4 Observed workforce of the ten largest countries in a partial sample of 100 Bosch entities, 2026. Second source with its own collection methodology, not comparable with the panel.

India carries nearly four times the German workforce in this sample, spread across seven entities rather than 23. A single entity, Bosch Global Software Technologies, employs more people there than all German entities in the sample combined.

This figure is an illustration, not a data point of the study. The sample covers 100 of roughly 390 findable entities, the parent company is not included, and the workforce measurement follows a different methodology from the panel, so no concentration value can be derived from it. The country distribution within the sample is unaffected by this and indicates the direction in which the movement measured in the panel points.

SECTION 05

05 What this changes for assessing companies

Workforce structure does not come into being at the moment it is needed. It comes into being over the decades before, mostly as a side effect of decisions that had nothing to do with structure as a goal.

For anyone assessing a company, that shifts the question. Not: which structural measures were adopted. But: what structure was already there, and in which direction has it moved over twenty years. The first question is answered by every press release. The second is answered only by the past, and unlike the first it is verifiable.

That verification is what the panel supplies. Entity-level workforce distribution, per company and per year, on a consistent basis across decades and across peers, for any corporate family in the panel. The same construction that produced this cohort produces a cohort for any industry or any competitor set.

What an organisation can be in a crisis, it decided decades earlier, without recognising it as a decision.

The environment does not teach an organisation what to become. It amplifies what it already is. Whether that amplification proves an advantage or a liability in the event is decided by factors this study does not measure, and section 06 names them.

SECTION 06

06 Scope of application

WHAT THE MEASURE IS

What is measured is organisational and legal workforce concentration. It is not a measure of operational resilience, single-site failure tolerance, supply chain robustness, redundancy of production capacity, decision-making decentralisation or supplier and technology diversity. A corporation can run dozens of legally separate entities whose operational dependencies still run through one centre. Workforce distribution is plausibly correlated with several of those variables, because people sit where activity takes place, but a statement about robustness requires site, supplier or production data the panel does not contain.

WHY THE PAIRED COMPARISON IS THE CORRECT PROCEDURE

Each family is compared against itself between two points in time, which removes every compositional effect that could distort a cohort comparison. Controlling within constant entity counts removes the effect, and that is the expected result: corporations distribute by adding entities, so holding entity count constant controls away the mechanism through which the phenomenon expresses itself. Both procedures are disclosed, so the choice is checkable rather than asserted.

THE PRECISION OF THE ESTIMATE

A bootstrap across 20,000 draws yields a 95 per cent interval for the class A median of -16.2 to +10.4 points. The supported statement is that the classes differ significantly, at $p = 0.0015$ across 100,000 permutations and in all

seven twenty-year windows, not that the decline amounts to exactly 9.9 points.

COVERAGE AND MATCHING RUN AGAINST THE FINDING

Entities are assigned to corporations by name match rather than ownership, so subsidiaries under a different brand are systematically missing and the concentration measure understates actual distribution. Roughly 88 per cent of the observed workforce falls in the United States, which under-records relocations to Asia and Eastern Europe. Both biases work against the finding, and the US share is stable enough, 87.6 to 88.7 per cent between 2001 and 2020, to serve as a reference line: the corporations under study move in the opposite direction over the same period.

APPENDIX A

The 41 corporations

Concentration as a percentage of observed corporate workforce held by the largest entity. Classification derives exclusively from the 1998 base year.

CORPORATION	ENTITIES 98	ENTITIES 18	CONC. 98	CONC. 18	CHANGE	CLASS
Accenture	28	84	99.57	89.26	−10.31	C
Alstom	12	18	56.48	46.74	−9.74	A
AstraZeneca	4	9	96.94	98.22	+1.28	C
BASF	13	15	98.20	98.74	+0.54	C
Bayer	26	46	28.60	39.94	+11.34	A
Boeing	22	26	97.39	94.24	−3.15	C
Bosch	90	189	40.88	26.93	−13.95	A
Caterpillar	30	62	84.83	78.76	−6.07	B
Deloitte	65	143	96.26	78.69	−17.57	C
Eli Lilly	9	13	98.39	97.40	−0.99	C
Ericsson	15	25	99.04	95.96	−3.08	C
Goldman Sachs	5	10	95.76	93.80	−1.96	C
HSBC	12	22	92.16	78.52	−13.64	B
Honeywell	28	49	93.30	96.11	+2.81	B
Infosys	37	197	61.03	87.18	+26.15	A
JPMorgan	6	13	99.84	99.72	−0.12	C
Johnson & Johnson	19	34	85.45	70.78	−14.67	B
Johnson Controls	14	22	97.84	95.85	−1.99	C

CORPORATION	ENTITIES 98	ENTITIES 18	CONC. 98	CONC. 18	CHANGE	CLASS
KPMG	82	148	76.37	49.15	-27.22	A
Lockheed Martin	14	21	99.76	98.56	-1.20	C
Medtronic	16	29	97.79	95.49	-2.30	C
Merck	34	46	96.76	92.36	-4.40	C
Michelin	14	32	92.75	92.41	-0.34	B
Morgan Stanley	7	18	99.70	99.46	-0.24	C
Motorola	23	17	82.49	64.25	-18.24	B
Nestlé	36	100	70.35	51.85	-18.50	A
Nokia	10	20	89.26	92.42	+3.16	B
Nortel	14	7	99.28	92.74	-6.54	C
Northrop Grumman	27	50	99.07	98.03	-1.04	C
Novartis	24	38	74.25	58.88	-15.37	A
PepsiCo	6	18	99.94	99.71	-0.23	C
Pfizer	8	16	99.82	99.63	-0.19	C
Philips	58	95	85.47	81.88	-3.59	B
Procter & Gamble	11	13	95.28	88.30	-6.98	C
Raytheon	35	41	80.00	84.10	+4.10	B
Sanofi	16	23	56.02	46.17	-9.85	A
Schneider Electric	14	27	72.01	82.36	+10.35	A
Siemens	57	76	71.26	61.24	-10.02	A
Unilever	10	18	99.42	97.90	-1.52	C
UnitedHealth	3	11	99.56	99.47	-0.09	C
Xerox	56	98	95.38	80.81	-14.57	C

APPENDIX B

Method protocol

DATA BASIS	Vivameda Workforce Panel, table company_year_snapshot, filtered to observed company-years with positive workforce. The panel spans 1950 to 2026; the historical layer used throughout this study runs to 2020.
ANCHOR SELECTION	848 companies with a maximum observed workforce of at least 2,000 in the period 1995 to 2005. From these, 41 corporate families after the exclusions set out in section 02.
FAMILY CONSTRUCTION	Name matching with a block list for known namesakes (Stellenbosch University, Jeroen Bosch Ziekenhuis, Katz Abosch, Gemeente 's-Hertogenbosch and others). Families with fewer than three entities in either measurement year are excluded.
CONCENTRATION MEASURE	Share of the largest entity in the sum of observed workforce across all entities of the same family in the same year.
CLASSIFICATION	Three classes by base-year value, thresholds at 80 and 95 per cent, fixed before analysis.
SIGNIFICANCE	One-sided permutation test across 100,000 random reassignments between classes A and C, test statistic is the difference of medians. Bootstrap confidence intervals across 20,000 draws.
ROBUSTNESS	Repetition across seven windows (1994/2014 to 2000/2020), correlation check for ceiling effect, total-panel reference line, banded check as counter-test.
FULL PROTOCOL	The complete protocol, including the alternative measurement approaches that were tested and rejected before this construction was adopted, is documented and disclosed on request.

Data basis: Vivameda Longitudinal Workforce Panel. 4 million companies, 48 million observed company-years, 1.88 billion skill rows, 46.5 million capability buckets, 100+ countries, 1950 to 2026. All classification rules and queries are reproducible and are disclosed on request. This analysis is provided for information and does not constitute investment, legal or tax advice.

WORKFORCE COMPOSITION BEFORE ORGANISATIONAL FAILURE

The Portable Asset

Every AI company that contracted was already thinning its technical core three years earlier. Every one that survived was thickening it.

CORE FINDING

Four contracted AI companies lost engineering share 2017 to 2020. Two survivors gained it. Six of six in the expected direction

WHY IT MATTERS

No annual report disclosed it, and two of the six no longer exist. The composition survives only in the panel

DESIGN

Cohort and classification fixed in writing 10 August 2026, before measurement. Identical window for every company

SECTION 01

01 The signal is in the composition, not the headcount

An acquirer looking at an AI company is asking one question: is the capability durable. Every available source answers a different question instead.

Financial statements give a headcount. They do not say how many of those people were technical, and no reporting standard requires that breakdown. For privately held companies no figure is published at all. A data room shows what management chose to assemble. A current snapshot shows who is there today, which is by construction the population that has not yet left.

Workforce composition is the variable that answers it, and it is observable at company level, comparable across peers, available for firms that have since disappeared, and deep enough in time to test whether a pattern actually preceded an outcome. That combination is what every study in this portfolio exploits. Study II used it to separate surviving retailers from failing ones. This study points the same construction at AI companies.

Growth data ranks companies. Workforce structure explains the ranking, and it is visible earlier.

SECTION 02

02 The result

Six AI companies, classified in 2020 on the panel's skill taxonomy and split by what subsequently happened to them. For each, the engineering share of the observed workforce in 2017 and in 2020, measured over an identical window for every company, on one role taxonomy.

4 of 4

companies that later contracted were thinning their technical core by 2020

2 of 2

companies that survived were thickening theirs over the same window

3 yrs

between the composition signal and the contraction it preceded

COMPANY, ENGINEERING SHARE OF OBSERVED WORKFORCE	2017	2020	CHANGE
CONTRACTED BETWEEN 2022 AND 2026			
Primer	61.3 %	46.8 %	-14.5 pp
H2O.ai	56.8 %	53.5 %	-3.3 pp
Clarifai	41.3 %	38.8 %	-2.5 pp
Petuum	66.7 %	65.3 %	-1.4 pp
PERSISTED OVER THE SAME PERIOD			
WeRide	52.2 %	74.1 %	+22.0 pp
Quantiphi	61.1 %	62.9 %	+1.8 pp




FIGURE 1 Each bar spans the company's 2017 and 2020 values on a scale of 0 to 100 per cent engineering share of the observed workforce. Copper runs downward, blue upward. All rows carry observed evidence and role coverage above 78 per cent.

Primer lost 14.5 points of engineering share and Clarifai lost 2.5, both in the years before the contraction that followed. Petuum, which the panel shows holding a technical core of two thirds through 2020, is now a shell. No annual report disclosed any of this. Two of these companies are privately held and never published a workforce figure of any kind.

SCOPE OF THE CLAIM

Six companies is a small cohort and the separation is clean rather than statistically decisive: on an exact test the probability of six of six falling this way by chance is about 0.07. This is a pattern consistent across every company measured, on an identical window, with the classification fixed before the outcome was examined. It is direction and precondition, not a forecast for any individual company.

SECTION 03

03 Why this cannot be reconstructed anywhere else

The question is: how was Petuum composed in 2017. It is not a curiosity. It is the question one would have had to ask in 2019 to see something three years before the company emptied.

SOURCE	WHAT IT ANSWERS	VERDICT
Annual report	Nothing. Privately held, no workforce figure published in any year	not available
Professional network today	Profiles show the state of 2026, not of 2017. Titles are overwritten, not versioned	not available
Web archive	Snapshots of pages. No structured workforce composition	not available
Vivamedata panel	Observed role composition per year across nine buckets, 2017 to 2020, on the same taxonomy as any peer	complete

None of this can be collected after the fact. Profiles are overwritten rather than versioned, so when someone updates a title the old one is gone. Company pages disappear with the company. No reporting standard requires a breakdown of workforce by function. The panel recorded these states continuously rather than reconstructing them later, and brought them onto a uniform taxonomy. Only that standardisation makes a company of 2017 comparable with one of 2020, and a defunct one with a surviving one.

SECTION 04

04 The panel does not stop when the company does

Composition is one half. The other is that the people remain observable after the entity has gone. Of the population carried at Petuum between 2016 and 2020, seventeen subsequently held a position at Google, each with both positions and both date ranges attached.

PERIOD	EMPLOYER	POSITION
2017 to 2019	Petuum	Software Engineer
2022 onward	Google	Software Engineer
2018 to 2021	Petuum	Software Engineer II, Machine Learning
2021 onward	Google	Software Engineer

Two of the seventeen, as carried in the panel.

The same construction resolves any named destination: frontier laboratories, large technology employers, competitors, or employers outside the sector entirely. For a buyer assessing whether a target's technical core has already left, that is the difference between an assertion and an evidenced count.

TWO LAYERS, NEVER NETTED

Composition and subsequent presence are measured on separate layers whose observation depths differ by orders of magnitude, so no figure from one is divided by, subtracted from or spliced onto a figure from the other. Section 02 and section 04 stand side by side in this document and never inside one calculation.

SECTION 05

05 The fifth question of one instrument

STUDY	WHAT IT DETECTS	RELEVANT DECISIONS
I The Unicorn Signature	Pre-hypergrowth patterns	Sourcing, screening
II The Defensible Shelf	Structural business vulnerability	Diligence, credit risk
III The Closed Circuit	Capability and talent-network lock-in	Transformation, recruiting
IV The 1998 Head Start	Long-term organisational readiness	Resilience, operating model
V The Portable Asset	Capability decay before failure	Acquisition, portfolio review

Read together the five map the same variable across the full life of an organisation. Study I asks which companies are structurally ready to scale before the growth curve is visible. Study II asks which are structurally exposed before the market removes them, using the technical share construction applied here. Study III asks whether an industry that says it is transforming is actually transforming. Study IV asks what an organisation already was, decades before the question mattered. Study V asks what a company was made of in the years before it emptied, and where those people are now.

Data basis: Vivameda Longitudinal Workforce Panel. 4 million companies, 48 million observed company-years, 1.88 billion skill rows, 46.5 million capability buckets, 100+ countries, 1950 to 2026. Composition measured on the panel's role layer, observed role classification per company-year, identical 2017 to 2020 window for every company in the cohort. Subsequent presence measured on a separate layer from dated employment histories; the two layers are reported separately and are never netted. Cohort and classification fixed 10 August 2026 before measurement. Panel figures are observed workforce records on a uniform taxonomy, not reported corporate headcount: they show structure and direction, not absolute employment. All classification rules and queries are reproducible and are disclosed on request. This analysis is provided for information and does not constitute investment, legal or tax advice.

CLOSING SECTION

Applying the research to your own universe

Each study was produced by pointing one instrument at a defined set of companies. The set is the only thing that changes: given a company list, a market definition, a portfolio or an acquisition universe, the same reconstruction runs against it.

A NAMED COMPANY LIST

A CRM export, target list or research universe is matched into the panel and returned with role and capability composition per company and year, trajectory measures, pattern flags and model scores. This is the construction behind Study I: a defined universe, ranked, with the reasoning attached to each row.

A MARKET OR SEGMENT

Where the set is defined by a segment rather than by names, the cohort is built first, on a written rule with a stated minimum observation threshold, and only then measured. Study II built its 225-company cohort that way and split it on one pre-declared criterion before looking at outcomes.

A PORTFOLIO OR CREDIT BOOK

Holdings are measured against their own history and against peers on one taxonomy: workforce trajectory, technical and capability share, direction of movement over a decade. This supports structural review and ranked attention, not an event forecast for a single name.

AN ACQUISITION UNIVERSE

The panel answers questions no data room contains: what a target was composed of before current management, whether a stated capability appears in the workforce, and where its people come from. Study III found an industry recruiting its declared capability entirely from itself. Study V measured what a technical core looked like three years before the company emptied, and named where those people went.

What is needed, and what holds in every case

To start: a list of company names with any available identifiers, plus the question the list is meant to answer. What comes back is either a labelled feature file, one row per company with measured fields, flags and explanation columns, or a written analysis in the form of these studies. Delivery formats are set out at the end of Study I.

- Classification rules are fixed before the outcome is examined, and are disclosed on request.
- Historical and present-day layers are reported separately and never netted, because their observation depth differs by orders of magnitude.
- Panel figures are observed workforce records on a uniform taxonomy, not reported headcount: structure and direction, not absolute employment.
- Classification rules, cohort selection and the full method protocol, including approaches that were tested and rejected, are documented and disclosed on request.

The question is not what a company says it is becoming. It is what its workforce already shows it to be.

CONTACT

Vivameda Ltd, Cyprus
oli@vivameda.com · vivameda.com

Data basis throughout: Vivameda Longitudinal Workforce Panel. 4 million companies, 48 million observed company-years, 1.88 billion skill rows, 100+ countries, 1950 to 2026. Individual studies state their own windows, cohort rules and thresholds. Provided for information; not investment, legal or tax advice.