Did Fundamentals Justify the AI Repricing of Enterprise Software?

Evidence from 2023–2026

Abhishek Sehgal

Eye Dream Pty Ltd (trading as Finsoeasy) • FSE Research & Investments Pty Ltd

ORCID: 0009-0000-9424-4695

This version: June 18, 2026

Abstract

Research question. The post-2022 AI narrative pushed enterprise-software multiples down on the fear that generative AI would commoditise SaaS workflows. We ask whether the repricing reflected deterioration in business value or a blunt AI-disruption narrative applied across firms whose fundamentals were in fact diverging.

Method. We rebuild the opportunity set from financial statements, not from price action. After carving out 12 payments/fintech firms whose float economics differ from SaaS, the core sample contains 335 firms with usable 2023 and latest-2026 endpoints. The analysis then follows one chain: (1) test business reality, (2) estimate intrinsic value independently of price, (3) read current price in reverse through implied growth, (4) attach signed AI exposure, and (5) convert only the surviving evidence into a portfolio watchlist.

Results. The evidence indicates a cohort-specific mispricing of growth expectations rather than a uniform market error. *(1) Expectations gap.* By reverse DCF, the repricing mis-set implied growth with opposite signs across cohorts: AI-native firms are priced for materially higher growth than they delivered (implied 43.6% versus realised 25.5%), whereas profitable workflow incumbents are priced for lower growth than they delivered (implied 8.6% versus realised 13.8%). *(2) Conversion test.* Firms with the highest AI salience did not convert it into superior economics: their gross-margin change is significantly worse than matched controls (difference-in-differences -2.7 pp, $t = -2.87$). *(3) Fundamentals and level value.* Over the period the 167 FCF-positive firms improved (median SBC-adjusted FCF margin 5.5%→11.6%, ROIC 5.7%→14.0%) while the 168 pre-FCF firms deteriorated (FCF margin -27.8%→-33.6%, median runway 0.97 years). On level value, broad SaaS is not cheap (median margin of safety -17.7% FCF-positive, -84.4% pre-FCF); the only cohort with a positive median margin of safety is the exposed-incumbent group (19 firms, median MoS 2.9%, 57.9% undervalued).

Implications. The output is a small, gated, reproducible watchlist of workflow incumbents whose prices imply impairment that the financial statements do not yet show. It is an ex-ante valuation gap, not demonstrated realised alpha: a point-in-time, sector-neutral implementation of the signal earns no significant factor-adjusted return over the sample (FF5+momentum $\alpha = -14.0\%$, $t = -0.97$), so the gap had not yet been harvested. The same signed-exposure framework can be reused for future narrative shocks.

Keywords: intrinsic value; enterprise software; AI disruption; economic moats; discounted cash flow; reverse DCF; portfolio construction; narrative-shock repricing.

JEL classification: G11, G12, G14, G17, M41, O33.

Summary of findings. Over 2023–2026 the enterprise-software repricing is consistent with a cohort-specific error in growth expectations rather than a uniform market mistake. By reverse DCF the market priced AI-native firms for higher growth than they realised (implied 43.6% vs realised 25.5%) and profitable workflow incumbents for lower growth than they realised (implied 8.6% vs realised 13.8%); firms with the highest AI salience did not convert it into better economics; and the under-expected incumbent cohort is also the only group that screens cheap on level value.

Contents

Reader's Map **5**

Portfolio-Manager Decision Sheet . . . 5

1 Introduction **6**

1.1 Scope and claims . . . 6

1.2 Contribution . . . 6

1.3 Roadmap . . . 7

2 Related Work and Hypotheses **7**

2.1 Intrinsic value and the price–value gap . . . 7

2.2 Intangible capital, R&D, and the accounting of software . . . 8

2.3 Artificial intelligence and firm value . . . 8

2.4 Measuring exposure from disclosure text . . . 8

2.5 Asset pricing controls and multiple-testing discipline . . . 8

2.6 Hypotheses . . . 8

2.7 Summary of findings . . . 9

3 Universe, Data and Selection **10**

3.1 Purpose . . . 10

3.2 Data sources . . . 10

3.3 Selection funnel . . . 10

3.4 Cohorts . . . 11

3.5 The two-world split: definition . . . 12

3.6 Point-in-time discipline and look-ahead mitigation . . . 12

3.7 Snapshot and data-tier disclosure . . . 12

4 First-Principles Metric Palette **13**

4.1 The six factor families . . . 14

4.2 Treatment of stock-based compensation . . . 14

4.3 The diagnostic and the investment question . . . 14

5 Business Reality: 2023 to 2026 **14**

5.1 World 1: FCF-positive firms improved . . . 14
5.2 World 2: pre-FCF firms deteriorated . . . 15
5.3 Trajectory by cohort . . . 15

6 Intrinsic Value versus Price **16**
6.1 Routed valuation architecture . . . 16
6.1.1 Route A: FCF-positive firms . . . 16
6.1.2 Route B: pre-FCF firms . . . 17
6.1.3 Fair value and margin of safety . . . 17
6.2 Reading the market in reverse: the central result . . . 17
6.3 Broad SaaS is not cheap . . . 18
6.4 Price-space corroboration: the localized margin of safety . . . 18
6.5 Did the gap pay in-sample? A realized-return reality check . . . 19
6.6 The justification quadrant . . . 20
6.7 The over-corrected screen . . . 20
6.8 Pre-FCF: a few survivors, many traps . . . 21

7 AI Exposure, Moats and R&D Productivity **22**
7.1 Signed exposure . . . 22
7.2 Talk is not conversion: an AI-pivot event study . . . 22
7.3 R&D productivity as a moat proxy . . . 23

8 Portfolio Construction **24**
8.1 Gates and score . . . 24
8.2 First-principles overlay from filings . . . 25
8.3 Ranked candidates . . . 26
8.4 Valuation-gap profile and naive benchmarks . . . 27
8.5 Industry concentration . . . 28
8.6 How a portfolio reader should use the output . . . 28
8.7 An out-of-sample factor-spanning test: does the signal pay? . . . 29

9 A Generalized Narrative-Shock Framework **30**

10 Robustness, Falsification and Limitations **31**
10.1 Robustness . . . 31
10.2 Identification: the AI narrative versus the rate shock . . . 32
10.3 Falsification . . . 32
10.4 Limitations . . . 33

11 Replicability **33**

12 Conclusion **34**

Acknowledgments **35**

Data Availability Statement **35**

How to Cite **35**

Disclosure Statement **36**

A Reproducibility and Audit **38**

B Statistical and Valuation Methodology Details **42**

B.1 Cost of capital (CAPM / WACC) 42

B.2 Owner-earnings DCF (Route A) 43

B.3 Path-to-profitability DCF (Route B) 43

B.4 Fair value and margin of safety 43

B.5 Reverse DCF (implied growth) 43

B.6 R&D capitalization and productivity 44

B.7 Event study: difference-in-differences 44

B.8 Quartile sorts 44

C Data Sources and Quality **44**

C.1 Primary fundamentals and prices: Sharadar 45

C.2 Factors and risk-free rate: Kenneth French 45

C.3 Return panel for the spanning test: public adjusted closes 45

C.4 Qualitative fields: SEC EDGAR 45

C.5 AI-text panel 45

C.6 Snapshot integrity 45

D EDGAR Extractor Methodology **46**

E Alternative Specifications **46**

F Glossary of Terms **47**

Reader's Map

Problem. The market repriced SaaS as if AI created one broad impairment risk. A value investor needs to know whether that repricing matched actual business deterioration.
Test. We start from cash-flow conversion, margins, ROIC, balance-sheet runway, dilution, R&D productivity, disclosure, DCF value and reverse-DCF implied growth, and only then ask what the multiple is doing.
Answer. The repricing mis-set *expectations*, not just multiples. Reverse DCF shows the market over-expects AI-native growth and under-expects profitable workflow incumbents; and the firms that talked most about AI did not convert it into better economics. Broad SaaS is not a bargain; the opportunity is the narrow, under-expected incumbent cohort, which also screens cheap on level value.
Output. What survives the gates is a small watchlist and a method that carries to the next narrative shock. Neither is a backtested claim of realised alpha, and the paper says so plainly.

The paper should be read as a sequence of gates. A firm must first belong in the SaaS universe, then show business resilience, then trade below independently estimated value, then carry a signed AI exposure that widens rather than erodes the moat, and finally survive practical portfolio gates. The watchlist is what remains after the gates, not what the author wanted to find.

Portfolio-Manager Decision Sheet

The sheet below restates the research design as the steps a portfolio manager would repeat before underwriting a position. It is a checklist for doing the work, not a list of names to buy.

Decision sheet: how to use the evidence without turning it into a story trade.

Question	Evidence used in this paper	Decision implication
Did the business actually weaken?	2023–2026 changes in revenue, gross margin, SBC-adjusted FCF margin, ROIC, runway and dilution.	If fundamentals weakened, do not call the selloff an over-reaction.
Can value be estimated from cash today?	FCF-positive firms use owner-earnings DCF; pre-FCF firms use path-to-profitability DCF with survival and dilution gates.	Do not use one valuation method for both worlds.
Is price below independently estimated value?	Margin of safety, reverse DCF, and value-versus-price quadrants.	Only research names where the valuation gap survives multiple lenses.
Does AI widen or erode the moat?	Signed exposure, filing/call stance, R&D productivity, switching costs and workflow embeddedness.	Avoid treating high AI salience as bullish by itself.
Can the thesis be monitored quarterly?	FCF margin, ROIC–WACC spread, realised versus implied growth, dilution, disclosed retention and AI-risk language.	Keep the watchlist alive only while the business evidence remains consistent.

Decision tree used by the portfolio screen.

Step	Gate	Output if the gate fails
1	Core SaaS universe after payments/fintech carve-out	Exclude from headline SaaS test
2	Latest SBC-adjusted FCF positive?	Route to pre-FCF survival/path-to-profit test
3	Positive base-case margin of safety?	Do not enter core long-research list
4	Exposed incumbent / workflow-embedded economics?	Treat as statistical value, not core thesis
5	Improving cash conversion and moat-quality gate?	Watch only; no core underwriting
6	Non-negative text stance, low dilution and investable size?	Exclude from core watchlist
7	Conservative valuation band positive?	Tier 1 if yes; Tier 2 watchlist if base case only

1 Introduction

Enterprise software sold off after 2022 because the market feared generative AI would commoditise SaaS: if a model can write the workflow, why pay for the seat? Multiples compressed across the sector. A value investor, however, does not ask whether prices fell; prices are an output. The investor asks whether *cash flows, margins, returns on capital and balance sheets* deteriorated enough to justify the move, and where they did not.

This paper tests the repricing in the order a value investor would underwrite it. First, did the businesses actually get worse? Second, if fundamentals did not deteriorate, did intrinsic value still fall because discount rates, growth or margins changed? Third, if value held up, did price fall far enough to create a margin of safety? Fourth, is the apparent bargain a real moat-widener or merely a cheap stock exposed to substitution? The paper is coherent only if the answer survives all four questions. A cheap multiple alone is not enough.

The conclusion is narrow. The repricing was largely consistent with fundamentals for cash-burning SaaS and still-rich AI narratives, but not across the board. The clearest over-discount appears in a small set of profitable, improving, workflow-embedded incumbents: firms where AI looks more like a feature and productivity layer than a product substitute. That is the "moat test." AI is not a sector call; it is a sorting device.

Finding

The paper's thesis in one line: the AI SaaS selloff was *mostly justified in aggregate but locally excessive for profitable workflow incumbents.* Everything else in the paper is a test of that sentence.

1.1 Scope and claims

This is a research paper, not a strategy memo, written in the first-principles value tradition of Greenwald et al. [1], Damodaran [2] and the expectations-investing logic of Rappaport and Mauboussin [3]. My prior going in was that the selloff looked overdone; the data agreed only in part, and where it did not I have said so rather than dress the result up. No name reaches the watchlist without clearing the valuation, moat, survival and risk gates that produced it, and every headline number resolves from a frozen, validated data artifact rather than the keyboard (Appendix A).

1.2 Contribution

Two empirical results carry the paper, and we lead with them because they are sharper and more defensible than any single cheapness statistic. **First, an expectations gap.** Reading price in reverse through implied growth shows the AI repricing mis-set expectations cohort by cohort, with a consistent and economically large sign. AI-native firms are priced for growth

far above what they delivered (median implied 43.6% versus realised 25.5%), while profitable workflow incumbents are priced for growth *below* what they delivered (median implied 8.6% versus realised 13.8%; Table 8). The market did not derate software uniformly; it over-paid for the AI narrative and under-paid the cash-flow reality, and these two errors point in opposite directions. **Second, a conversion test.** An AI-pivot event study shows that the firms that talked most about AI did not convert that talk into superior economics: their gross-margin change was significantly *worse* than matched controls (difference-in-differences -2.7 pp, $t = -2.87$; Table 15). Together these say the mispricing is about expectations, not realised fundamentals, and that AI salience is a cost signal before it is a moat signal.

The localized positive margin of safety we find in the incumbent cohort (Section 6) is the price-space *confirmation* of the same expectations gap, not a standalone headline; it is a modest median and we treat it as corroboration. Supporting these two results is the rest of the apparatus: (i) a full-population, fundamentals-first map of post-AI enterprise software; (ii) a two-route valuation architecture that treats cash-generative and cash-burning firms differently; (iii) a signed-exposure method that separates moat-wideners from moat-losers; (iv) a narrow, gated portfolio watchlist; and (v) a method general enough to carry to later shocks such as quantum, robotics or regulation. Relative to text-based AI-exposure measurement [4, 5] and the AI-adoption literature [6, 7], the methodological contribution is to make exposure accountable to realised cash economics and intrinsic value rather than to language alone.

1.3 Roadmap

The order of the paper mirrors how an underwriter actually works the question. Section 2 turns the question into five falsifiable claims. Section 3 defines the sample so the opportunity set is auditable. Section 4 defines the financial statement palette. Section 5 asks whether the businesses got worse. Section 6 asks whether prices fell more than intrinsic value. Section 7 asks whether the apparent gap is a moat-widening AI exposure or a value trap. Section 8 converts only the surviving evidence into a watchlist. Section 9 abstracts the method, Section 10 states what would falsify it, and Section 12 synthesises the findings.

2 Related Work and Hypotheses

This paper sits at the intersection of four literatures: intrinsic-value investing, the accounting and asset pricing of intangible capital, the economics of artificial intelligence, and the measurement of firm-level exposure from disclosure text. We synthesise the strands that bear on the repricing question and then state five testable, falsifiable hypotheses.

2.1 Intrinsic value and the price–value gap

The premise that price and value are distinct, and that the gap is the object of analysis, is the foundation of value investing [1] and of practitioner valuation [2]. Rappaport and Mauboussin [3] reframe the same idea as *expectations investing*: rather than forecasting cash flows and comparing to price, one inverts the valuation to recover the expectations the current price embeds, then asks whether those expectations are too high or too low. We use both directions, a forward discounted cash flow (DCF) to estimate value and a reverse DCF to read the market's implied growth, because a narrative shock is precisely an event in which expectations, not realised fundamentals, move first.

2.2 Intangible capital, R&D, and the accounting of software

Software firms expense rather than capitalise most of what builds their moat: R&D, sales-led customer acquisition, and the data and integrations that create switching costs. A long literature shows that expensing intangibles distorts reported profitability and can lead to systematic mispricing. Lev and Sougiannis [8] construct R&D capital and link it to subsequent returns; Chan et al. [9] document that the market appears to undervalue R&D-intensive firms. Peters and Taylor [10] build firm-level intangible capital stocks from R&D and a portion of SG&A, and Eisfeldt and Papanikolaou [11] show that organisation capital carries a return premium. Novy-Marx [12] establishes gross profitability as a robust quality signal, and in earlier work we document that R&D intensity itself earns a long-run return premium in the broad equity cross-section [13]. We draw on this work to capitalise R&D by perpetual inventory and to measure R&D *productivity*, gross profit created per cumulative R&D dollar, as a durability proxy; the question here is narrower, namely whether that R&D converts into gross profit for a shock-exposed industry rather than whether intensity is priced market-wide.

2.3 Artificial intelligence and firm value

The macro and labour literature frames AI as automation that can either substitute for or complement existing factors [7], with heterogeneous firm-level effects depending on task structure and adoption capacity [6]. Most relevant here, Eisfeldt et al. [14] show that generative-AI exposure is priced in the cross-section of equities and that the sign of the effect depends on whether AI complements or substitutes a firm's workforce and product. This is the economic basis for our central construct, a *signed* exposure, rather than a magnitude-only "AI score."

2.4 Measuring exposure from disclosure text

Firm-level exposure to a diffuse shock can be measured from the language firms use. Hassan et al. [4] quantify political risk from earnings-call transcripts; Loughran and McDonald [5] provide the finance-specific sentiment lexicons that discipline such measurement; Li [15] shows that the information content of forward-looking statements in filings can be recovered with simple machine classifiers; and Hoberg and Phillips [16] build text-based industry classifications from as-filed disclosure. We adapt these methods to measure AI salience and stance, but, learning from the critique that talk is cheap, we validate the text signal against realised cash economics in an event study rather than treating language as the outcome.

2.5 Asset pricing controls and multiple-testing discipline

Any claim of mispricing must survive standard factor controls and the file-drawer problem. We adopt the factor vocabulary of Fama and French [17, 18] and Carhart [19], the cross-sectional methodology of Fama and MacBeth [20], the quality-minus-junk perspective of Asness et al. [21], the implementation-cost framing of Novy-Marx and Velikov [22], and the sentiment/limits-to-arbitrage context of Stambaugh et al. [23]. We treat the central return claim as associational and explicitly killable, in the spirit of the multiple-testing correction of Harvey et al. [24] and the false-discovery-rate control of Benjamini and Hochberg [25], and we pre-state falsification conditions (Section 10).

2.6 Hypotheses

We structure the analysis around five hypotheses. Each is paired with the section that tests it and with the evidence that would falsify it, so the paper can be read as a set of refutable

claims rather than a narrative.

1. **H1 (Two worlds, not one).** The post-AI fundamental path is bimodal: FCF-positive firms did not broadly deteriorate in cash economics over 2023–2026, while pre-FCF firms did. *Rationale.* "AI hurt SaaS" treats software as one object; if the cash statements split into two regimes, the sector-level claim is the wrong unit of analysis. *Falsified if* FCF-positive firms show broad margin/ROIC impairment, or pre-FCF firms show broad improvement. (Section 5.)
2. **H2 (Localized, not broad, over-reaction).** Broad SaaS is not cheap on intrinsic value; any positive margin of safety is concentrated in a narrow, identifiable cohort rather than spread across the sector. *Rationale.* A disciplined value test should resist the temptation to declare the whole sector mispriced. *Falsified if* the median firm is undervalued, or if no cohort carries a positive median margin of safety. (Section 6.)
3. **H3 (Under-expected incumbents).** For the exposed-incumbent cohort, market-implied growth recovered by reverse DCF is below realised growth: the market is paying for less than the firms delivered. *Rationale.* This is the expectations-investing signature of an over-reaction. *Falsified if* implied growth meets or exceeds realised growth for incumbents. (Section 6.)
4. **H4 (Talk is not conversion).** Firms with the highest AI-narrative salience did not convert it into superior subsequent cash economics relative to controls. *Rationale.* If a cheap screen merely rewards firms that talk about AI, it is a narrative trade, not a value one. *Falsified if* high-salience "pivoters" show significantly better fundamental deltas. (Section 7.)
5. **H5 (Quality is priced, not free).** R&D productivity identifies durable, high-quality firms but not cheap ones: higher-productivity quartiles do not carry higher margins of safety. *Rationale.* Separating quality from cheapness prevents conflating a good business with a good investment. *Falsified if* the highest-productivity quartile is also the cheapest. (Section 7.)

2.7 Summary of findings

Table 1 previews the results so the paper can be read as a set of resolved claims. Every figure in it resolves from the validated snapshot (Appendix A); the two lead results are H3 (the signed expectations gap) and H4 (the conversion test).

Table 1. Summary of findings. Each hypothesis is paired with its prediction, the realized result with its key statistic, and a verdict. "Supported" is an in-sample association, not a claim of realized alpha; the realized-return row is reported separately to make that distinction explicit.

Hypothesis	Prediction	Result (key statistic)	Verdict
H3 *(lead)*	Incumbents priced for less growth than realized	Reverse DCF: incumbents implied 8.6% vs realized 13.8%; AI-native implied 43.6% vs realized 25.5% (opposite sign)	Supported
H4 *(lead)*	AI talkers do not convert talk to economics	Event study DiD on gross margin -2.7 pp ($t = -2.87$); FCF-margin null but underpowered (MDE 12.5 pp)	Supported
H1	Two worlds: FCF-positive stable, pre-FCF deteriorates	FCF-positive FCF margin 5.5%→11.6%; pre-FCF -27.8%→-33.6%	Supported
H2	Over-reaction localized, not broad	Median MoS -17.7% (FCF-positive); only incumbents positive (2.9%, 19 firms)	Supported (modest)
H5	R&D productivity is quality, not cheapness	Top productivity quartile is the most expensive, not the cheapest	Supported
Realized return	Does the valuation gap pay in-sample?	Under-expected cohort de-rated (median -41.1%); systematic sleeve earns no significant FF5+MOM alpha (-14.0%, $t = -0.97$, 31 mo)	Ex ante only

3 Universe, Data and Selection

3.1 Purpose

Before the paper can say whether the market was wrong, it must define what "software" means. This section narrows a broad software screen into the actual SaaS opportunity set used for valuation. The reader must be able to see which companies entered, which were excluded, and why.

3.2 Data sources

Fundamentals are point-in-time from the Sharadar Core US Fundamentals feed; prices and split/dividend-adjusted returns are from Sharadar Equity Prices; market factors and the risk-free rate are from the Kenneth French data library; and qualitative fields are extracted from as-filed annual reports (SEC EDGAR `10-K`, `20-F` and `40-F`). All foreign-currency line items are converted to USD before ratios are formed. The study window is the 2023 baseline against the latest available 2026 quarter (Sep 2023–Mar 2026), a median observed span of 2.25 years per firm. We use a single accessible-data tier (Tier-1 (Sharadar + SEC EDGAR + Ken French)) rather than the CRSP/Compustat academic standard; the trade-offs of that choice are stated in Appendix C.

3.3 Selection funnel

Starting from the Sharadar application- and infrastructure-software universe, we apply a survivorship-safe inclusion rule, a data-availability filter (a firm must have both a 2023 baseline

and a latest-2026 endpoint), and a business-model carve-out for payments/fintech firms whose float and spread economics distort cash-flow and valuation ratios. 347 firms clear the endpoint filter; 12 are carved out as payments/fintech; the final core SaaS analysis set is 335 firms. Table 2 and Figure 1 show the funnel and its terminal split into FCF-positive and pre-FCF firms.

Table 2. Selection funnel from endpoint-eligible firms to the core SaaS analysis set. Counts are computed directly from the frozen analysis file; upstream Sharadar software screening and the positive-gross-profit / endpoint-availability filters precede row 1.

Step	Description	Count	Reason
1	Firms with 2023 + latest-2026 endpoints	347	valuation requires both endpoints
2	Less payments/fintech carve-out	−12	float/spread economics differ
3	**Core SaaS analysis set**	**335**	paper universe
4	of which FCF-positive	167	already cash-generative
5	of which pre-FCF / cash-burning	168	value depends on survival

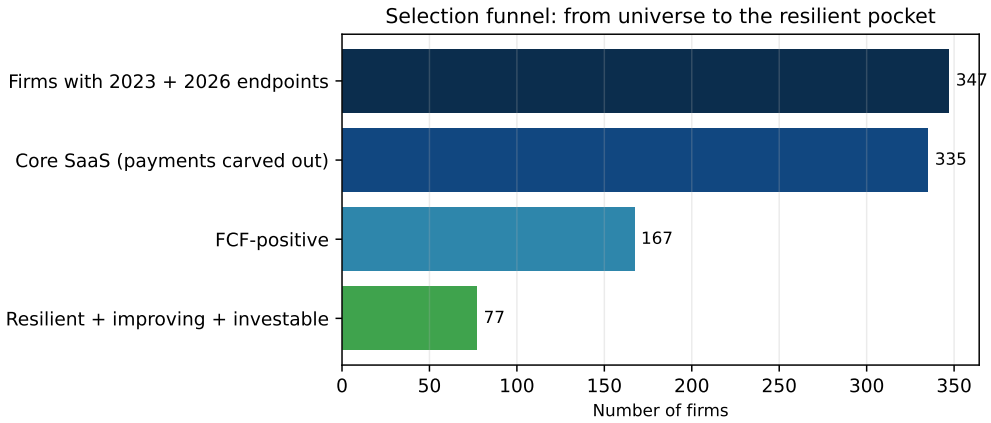

Figure 1. Selection funnel. From endpoint-eligible firms, through the payments carve-out, to the core SaaS set, the FCF-positive subset, and the resilient-and-improving investable pocket that anchors the portfolio in Section 8.

3.4 Cohorts

We label each firm by its economic relationship to the AI narrative. The labels are rule-based with documented overrides, and they carry an *exposure sign* used later in valuation and portfolio gates (Table 3). The cohorts are deliberately economic, not thematic: an "exposed incumbent" is a seat- or workflow-priced system whose customers face switching costs, not merely a firm that mentions AI.

Table 3. Cohort definitions and the sign of AI exposure. Representative types are illustrative of the cohort, not endorsements.

Cohort	Economic meaning	AI exposure sign	Representative type
AI-native / enabler	sells AI or AI infrastructure	priced as beneficiary	model / data-platform names
Exposed incumbent	seat/workflow software, narratively threatened	ambiguous; depends on switching cost	enterprise workflow SaaS
Exposed other	application software with AI-relevant workflow exposure	mixed	vertical application names
Infrastructure / neutral	security, observability, data plumbing	lower substitution risk	infra / security names
Payments / fintech	not core SaaS (excluded from headline)	n/a	float/spread economics

3.5 The two-world split: definition

The single most important partition in the paper is whether a firm already converts its business into owner cash flow. We define a firm as *FCF-positive* when its latest SBC-adjusted free cash flow is positive,

$$\text{vgroup}_i = \begin{cases} \text{FCF-positive} & \text{if } \text{FCF}^{\text{SBC-adj}}_{i,\,\text{latest}} > 0, \\ \text{pre-FCF} & \text{otherwise,} \end{cases} \tag{1}$$

where $\text{FCF}^{\text{SBC-adj}}$ is operating cash flow less capital expenditure less stock-based compensation (SBC treated as a cash-equivalent cost, Section 4). The split is economic, not stylistic: it determines which valuation route a firm receives (Section 6) because the two groups require different evidence: cash-generative firms can be anchored in owner earnings, while cash-burning firms must first clear a survival and path-to-profit test. The 167 FCF-positive and 168 pre-FCF firms are analysed side-by-side throughout.

3.6 Point-in-time discipline and look-ahead mitigation

The baseline and latest readings use as-reported, point-in-time fundamentals: the 2023 endpoint uses data that was public as of the 2023 fiscal observation, and the latest endpoint uses the most recent available quarter, so no firm is evaluated on figures it had not yet filed. Prices are split- and dividend-adjusted to the same endpoints. Because the design is a two-endpoint *change* study (2023 baseline versus latest 2026 reading) rather than a monthly tradeable backtest, look-ahead enters only through the endpoints themselves; we therefore (i) require both endpoints to exist before a firm is admitted, (ii) measure price change over the identical window as the fundamental change, and (iii) treat the resulting portfolio as an *ex-ante* watchlist whose realised performance is explicitly out of sample (Section 10). The reverse-DCF and event-study tests compare contemporaneous quantities and so are not sensitive to formation timing.

3.7 Snapshot and data-tier disclosure

Every reported number is pinned to a frozen snapshot of three artifacts: the firm-level valuation run, the first-principles overlay, and the AI-text panel, identified by a content hash so the build cannot silently drift (snapshot `8a7ec05c7873`, built 17 June 2026; Appendix C). Table 4 records the exact formation and valuation parameters used to generate this manuscript;

data-derived entries (window, span, conversion k, target margin) are computed from the snapshot, while configuration entries (discount-rate inputs, glide length, depreciation, winsorization, screens) are the documented engine inputs.

Table 4. Formation and valuation parameters (snapshot `8a7ec05c7873`). Data-derived entries are computed from the frozen artifacts; configuration entries are documented engine inputs.

Parameter	Value / definition
Sample and window	
Study window	Sep 2023 baseline → Mar 2026 latest
Median observation span	2.25 years
Inclusion filter	positive gross profit; 2023 and latest-2026 endpoints
Business-model carve-out	payments / fintech excluded from headline
Route A: FCF-positive valuation	
Core model	owner-earnings DCF on SBC-adjusted FCF
Discount rate	per-firm CAPM WACC (Eq. 3)
Beta window	5 years of monthly returns vs. market
Risk-free rate r_f	4.30% (Ken French trailing 12m)
Equity risk premium	5.0%
Tax rate τ	21%
Explicit horizon / terminal growth	10 years / 3%
Growth fade and clamps	stage-1 growth fades to terminal; revenue growth clipped -5% to 30%
Triangulation	median of DCF, normalized DCF, peer EV/Sales
Route B: pre-FCF valuation	
Core model	path-to-profitability DCF with survival gate
Discount rate	per-firm CAPM WACC +2pp exec. premium (cap 20%; 12% fallback)
Margin glide	7 years to target
Steady-state conversion k	0.19 (median)
Target FCF margin	$k\times$ gross margin (median 10.6%)
Valuation clamps	target margin 5–40%; annual dilution cap 25%; displayed MoS -95–500%
Moat and event study	
R&D capitalization	perpetual inventory, 20%/yr depreciation
Event-study treatment	top 33% of 2023–24 AI-call salience
Delta winsorization	2.5/97.5 percentiles
Significance test	Welch (unequal-variance) t
Screens	
Resilient market-cap floor	\$500M
Investable over-corrected floor	\$1bn
Margin of safety	fair value/price − 1

4 First-Principles Metric Palette

With the sample fixed, the next question is what counts as "business reality." The paper is built from six families of factors. Together they answer whether a business is still growing, whether scale improves its economics, whether growth becomes owner cash flow, whether it earns above its cost of capital, whether it can survive the correction, and whether AI is likely to widen or erode its moat.

4.1 The six factor families

1. **Growth and durability:** revenue CAGR, gross-profit growth, deferred-revenue growth where disclosed. *Is demand still expanding, or did the category structurally weaken?*
2. **Profitability and operating leverage:** gross, operating and net margins. *Is scale improving economics, or is AI compressing margins?*
3. **Cash-flow quality:** FCF margin, SBC-adjusted FCF margin, and stock-based compensation (SBC) as a share of revenue. *Is reported growth turning into owner cash flow?*
4. **Returns on capital:** ROIC and ROIC minus the weighted average cost of capital (WACC). *Is the business creating value above its cost of capital?*
5. **Balance sheet and survival:** net cash / net debt, runway for pre-FCF firms, and annualised dilution. *Can the company survive long enough for the thesis to matter?*
6. **Moat and reinvention:** R&D intensity, R&D productivity, capitalised R&D stock, the Rule-of-40, and signed AI exposure. *Is AI likely to widen the moat or erode it?*

4.2 Treatment of stock-based compensation

We treat SBC as a real economic cost throughout: the FCF measure that drives valuation is *SBC-adjusted*, and dilution is tracked separately so that per-share value reflects shareholder leakage. This is deliberately conservative; it is also where many cash-burning SaaS valuations break.

4.3 The diagnostic and the investment question

We define the first-principles diagnostic as a function of the six families,

$$\text{Business Reality}_i = f(\text{Growth}_i,\ \text{Profitability}_i,\ \text{CashFlow}_i,\ \text{Returns}_i,\ \text{BalanceSheet}_i,\ \text{Moat}_i),$$

and the investment question as the gap between how value moved and how price moved,

$$\text{Mispricing}_i \ = \ \Delta\text{Intrinsic Value}_i \ - \ \Delta\text{Price}_i. \tag{2}$$

Everything that follows is an attempt to measure the right-hand side of Equation (2) honestly, firm by firm.

5 Business Reality: 2023 to 2026

This is the first major evidence section because it comes before valuation. If fundamentals deteriorated, the selloff may be justified and there is no need to force a mispricing story. The literal question is: did SaaS businesses actually get worse after AI entered the narrative? The answer is not "yes" or "no." The universe split into two worlds.

5.1 World 1: FCF-positive firms improved

Among the 167 FCF-positive firms, median gross margin held (72.2% to 73.5%), median SBC-adjusted FCF margin rose from 5.5% to 11.6%, and median ROIC roughly doubled from 5.7% to 14.0%, while revenue still compounded at a median 11.5%. These businesses looked more efficient, not broken; yet their median price fell -6.5% over the same window. The improvement is broad, not driven by a few names: across the whole SaaS population, 68.7% of firms improved operating margin (median +6.1 pp), 62.4% improved net margin, 60.6% improved SBC-adjusted FCF margin, and 57.9% improved ROIC.

5.2 World 2: pre-FCF firms deteriorated

Among the 168 pre-FCF firms, median gross margin was flat (55.3% to 55.2%) but median SBC-adjusted FCF margin worsened from -27.8% to -33.6% and median ROIC fell from -24.6% to -33.5%. Growth did not offset burn (median revenue CAGR 11.3%), median runway is only 0.97 years, 63.1% have under two years of runway, and median annual dilution is 17.9%. For this world, the market's caution is largely justified by the accounts.

Table 5. Fundamental trajectory, 2023 to latest 2026 (group medians). Margins, ROIC and revenue CAGR are percentages; Δprice is the median price change over the window.

Group	Gross margin (2023→2026)	SBC-adj FCF margin (2023→2026)	ROIC (2023→2026)	Rev. CAGR	Median MoS	Median Δprice
FCF-positive (n=167)	72.2 → 73.5	5.5 → 11.6	5.7 → 14.0	11.5	-17.7	-6.5
Pre-FCF (n=168)	55.3 → 55.2	-27.8 → -33.6	-24.6 → -33.5	11.3	-84.4	-9.1

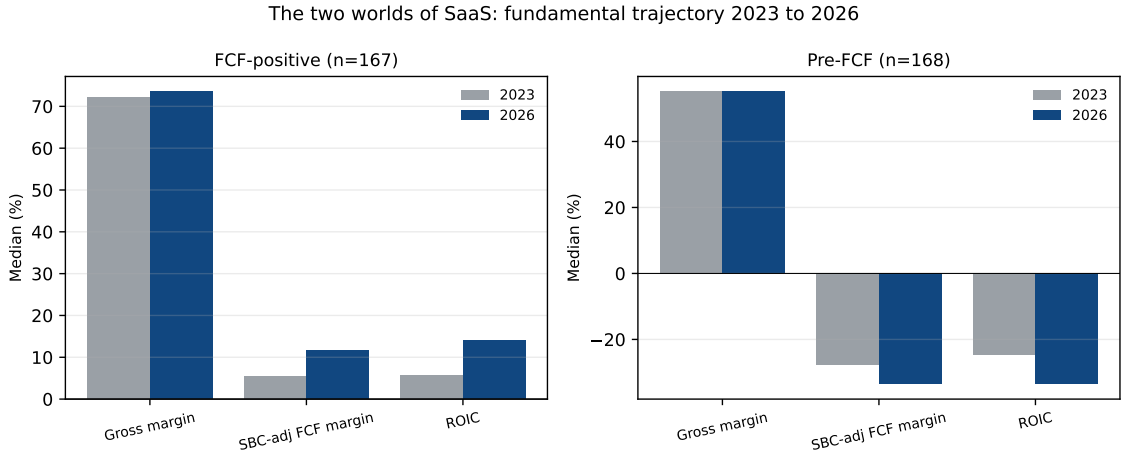

Figure 2. The two worlds of SaaS. Left: FCF-positive firms show rising SBC-adjusted FCF margin and ROIC with stable gross margin. Right: pre-FCF firms show deteriorating cash margin and ROIC. Bars are group medians for 2023 and the latest 2026 reading.

5.3 Trajectory by cohort

The two-world split also has cohort structure (Table 6). Among FCF-positive firms, exposed incumbents are the most profitable cohort, with median operating margin and ROIC well above the group, while AI-native names show the highest growth but the thinnest cash margins. Among pre-FCF firms (Table 7) almost every cohort is deeply cash-negative with heavy dilution; the pre-FCF world is a survival problem, not an AI problem.

Table 6. FCF-positive firms: latest-period profile by cohort (medians).

Cohort	n	Rev. CAGR (%)	Gross margin (%)	Op. margin (%)	SBC-FCF margin (%)	ROIC (%)	Rule-40 (%)
AI-native	7	25.5	79.9	-1.9	5.4	2.0	31.2
Exposed incumbent	19	13.8	79.6	17.1	16.0	32.2	28.7
Exposed other	81	9.7	72.7	9.1	12.1	13.1	23.9
Infrastructure	53	8.0	69.6	10.1	10.6	11.7	21.4
Neutral infra	7	13.3	77.4	5.5	12.1	19.1	29.2

Table 7. Pre-FCF firms: latest-period profile by cohort (medians). Margins negative across the board.

Cohort	n	Rev. CAGR (%)	Gross margin (%)	Op. margin (%)	SBC-FCF margin (%)	ROIC (%)	Rule-40 (%)
AI-native	6	28.0	70.2	-31.2	-28.1	-39.1	-3.7
Exposed incumbent	3	22.9	84.0	-16.8	-11.0	-38.0	14.5
Exposed other	101	3.4	50.1	-28.3	-29.2	-32.5	-19.4
Infrastructure	55	12.9	57.9	-44.5	-43.0	-39.8	-26.7
Neutral infra	3	23.6	77.5	-16.0	-3.6	-12.2	22.2

Finding

The profitable half of SaaS did not show broad AI-driven fundamental impairment over the observed window; its prices fell anyway. The cash-burning half got weaker, and its caution is earned. "AI hurt SaaS" is the wrong unit of analysis; the correct unit is the firm's cash economics.

6 Intrinsic Value versus Price

Section 5 showed that the profitable half of SaaS did not look broken while the cash-burning half did. That is still not enough. A falling price is only an opportunity if intrinsic value did not fall with it. So we estimate value *independently of price*, then ask whether the price move was larger than the value move.

6.1 Routed valuation architecture

Because the two worlds have different economics, a single valuation recipe would misstate one of them. We route valuation by business type.

6.1.1 Route A: FCF-positive firms

We use an owner-earnings DCF on SBC-adjusted FCF, discounted at a per-firm WACC, and triangulate against a peer EV/Sales relationship conditioned on the Rule-of-40 [2]. The cost of equity uses a CAPM beta from five years of monthly returns against the market factor:

$$\mathrm{WACC}_i = w_i^e(r_f + \beta_i\,\mathrm{ERP}) + w_i^d\, r_i^d(1-\tau), \tag{3}$$

with $r_f = 4.30\%$ the trailing risk-free rate (Ken French, annualised), $\mathrm{ERP} = 5.0\%$ the equity risk premium and $\tau = 21\%$ the tax rate. The beta is estimated from 5 years of monthly returns against the market factor. We estimate β for 309 of the SaaS firms; the median CAPM beta is 1.14 and the median WACC is 9.5% (interquartile range 7.6%–11.0%), shown in Figure 3. The economic logic is simple: if a firm already generates cash, value should be anchored in cash generation. Full mechanics are in Appendix B.

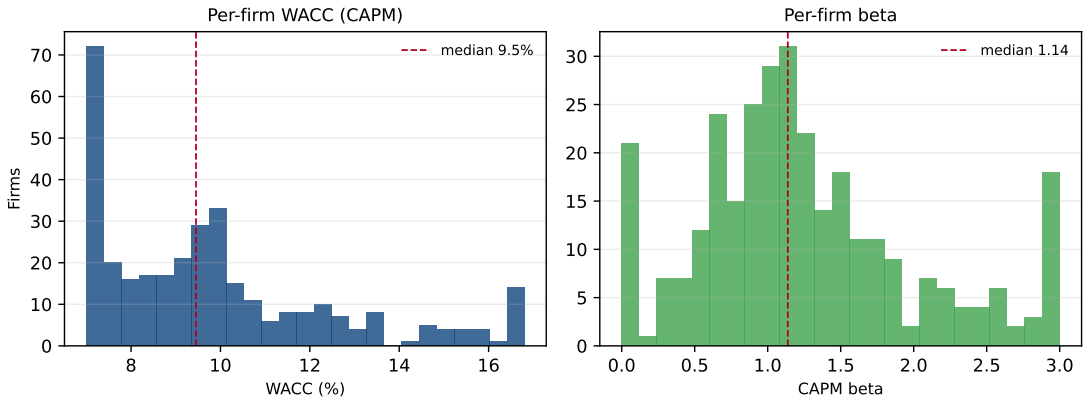

Figure 3. Per-firm cost of capital. Left: distribution of CAPM-based WACC (median 9.5%). Right: distribution of CAPM beta (median 1.14). Replacing a flat discount rate with per-firm WACC avoids systematically over- or under-valuing high-beta names.

6.1.2 Route B: pre-FCF firms

We use a path-to-profitability DCF: revenue is grown and margins converge over a 7-year glide toward a target FCF margin equal to a steady-state conversion $k = 0.19$ times gross margin (median target 10.6%), discounted at the firm's CAPM WACC plus a two-point execution-risk premium (capped at 20%; a 12% flat fallback where beta is unavailable), with dilution drag and a survival/runway gate. A contamination guard excludes implausibly high FCF yields from the FCF-based route. If a firm does not yet generate cash, value depends on whether it can survive and plausibly convert gross profit into owner earnings. Full mechanics are in Appendix B.

6.1.3 Fair value and margin of safety

For each firm we take the triangulated point estimate and a bear–base–bull band,

$$\text{Fair Value}_i = \text{median}\big(\text{DCF}_i,\ \text{Normalized-DCF}_i,\ \text{Peer-Multiple}_i\big), \qquad \text{MoS}_i = \frac{\text{Fair Value}_i}{\text{Price}_i} - 1. \tag{4}$$

6.2 Reading the market in reverse: the central result

The paper's primary evidence is not a margin-of-safety count; it is what price is *implying*. We invert the DCF, after Rappaport and Mauboussin [3]: the growth rate that sets fair value equal to the current price is the *market-implied growth*, and comparing it to realised growth shows what the market is actually paying for (Table 8, Figure 4). Two facts matter. At the aggregate level SaaS is fairly expected: across FCF-positive firms implied growth (12.4%) sits close to realised (11.5%), so there is no broad mispricing to harvest. But the aggregate hides a sign-flipping dispersion across cohorts. AI-native firms are priced for far more growth than they delivered, median implied 43.6% against realised 25.5%, an over-expectation of roughly the order of the realised number itself. **Exposed workflow incumbents are priced for the opposite: less growth than they realised**, median implied 8.6% against realised 13.8%. The market under-expects exactly the cohort whose cash economics held up and over-expects the cohort selling the narrative. This expectations gap, signed and cohort-consistent, is the cleanest statement of the over-reaction and the result the rest of the section corroborates.

Table 8. Reverse DCF: market-implied growth vs realised revenue CAGR, FCF-positive firms by cohort. A positive "realised − implied" means the market is paying for less growth than the firm delivered.

Cohort	Firms	Implied growth (%)	Realised CAGR (%)	Realised − implied (%)
AI-native	7	43.6	25.5	-21.7
Exposed incumbent	19	8.6	13.8	5.1
Exposed other	81	12.6	9.7	-2.6
Infrastructure	53	10.8	8.0	-3.7
Neutral infra	7	21.6	13.3	-8.3

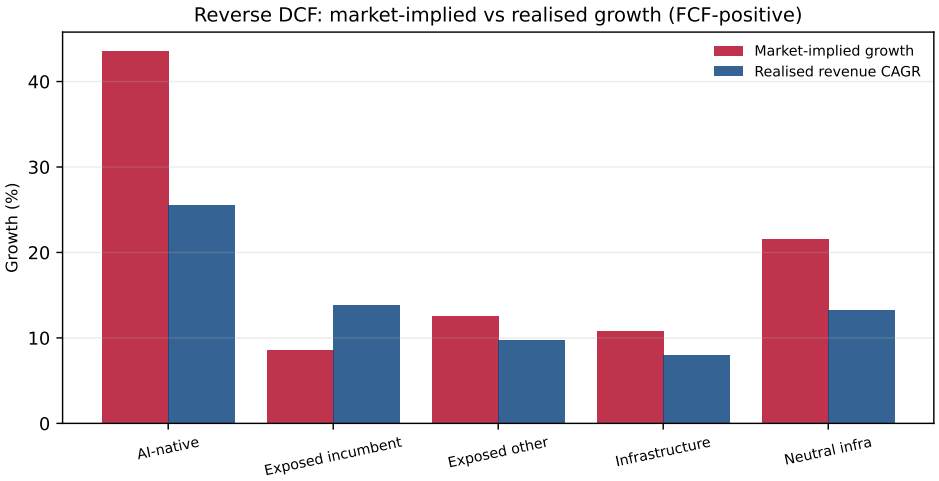

Figure 4. Reverse DCF by cohort. AI-native firms carry rich implied growth; exposed incumbents carry implied growth below what they realised.

6.3 Broad SaaS is not cheap

The first valuation result is a discipline check against wishful thinking. At the median, neither world is cheap: median MoS is -17.7% for FCF-positive firms (38.9% undervalued) and -84.4% for pre-FCF firms (20.2% undervalued). Even restricting to resilient FCF-positive firms, improving cash margin and Rule-of-40 above 20%, the median MoS is still -26.2% with only 29.9% undervalued (77 firms). The evidence does not support "buy SaaS broadly."

6.4 Price-space corroboration: the localized margin of safety

The expectations gap should also show up in level prices, and it does, in the same cohort. Among FCF-positive firms by cohort (Table 9, Figure 5), the only group with a *positive* median margin of safety is the exposed-incumbent group: 19 firms, median price change -41.1%, median MoS 2.9%, and 57.9% undervalued. We are deliberately modest about this number: a median margin of safety of 2.9% is thin, and on its own it would be a weak hook. Its value is that it independently confirms the reverse-DCF result: the cohort the market under-expects on growth is the same cohort that screens cheap on level value, while AI-native firms remain expensive on both lenses. Two methods, one answer.

Table 9. Margin of safety by cohort, FCF-positive firms only. The exposed-incumbent cohort is the single group with a positive median margin of safety.

Cohort	Firms	Median Δprice (%)	Median MoS (%)	% undervalued
AI-native	7	4.3	-81.2	0.0
Exposed incumbent	19	-41.1	2.9	57.9
Exposed other	81	-16.6	-19.4	35.8
Infrastructure	53	12.0	-21.6	43.4
Neutral infra	7	18.0	-62.0	28.6

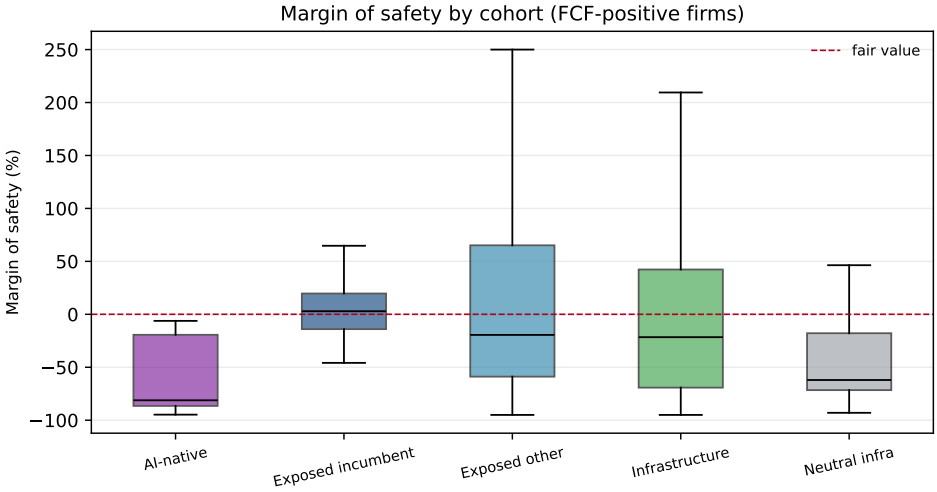

Figure 5. Margin of safety distribution by cohort (FCF-positive firms). The dashed line is fair value (MoS = 0). Only the exposed-incumbent box sits with its median above the line.

6.5 Did the gap pay in-sample? A realized-return reality check

The natural referee question is whether this ex-ante gap was already a realized winner over the sample. It was not, and saying so plainly is what keeps the paper honest. Over the realized 2023–2026 window the cohort the market *under*-expected on growth is the cohort whose price fell hardest: exposed incumbents posted a median price change of -41.1% and only 15.8% of them rose, while the *over*-expected AI-native cohort held up (median 4.3%, 57.1% rose; Table 9). That is, the discount on resilient incumbents widened across the window rather than closing. The signal is therefore **contrarian and forward-looking**: it identifies a price–value gap that is open *at the end* of the sample, not a return that was harvested *during* it. This is exactly why the paper's central claim is an ex-ante valuation gap and not demonstrated alpha.

We do not present the frozen cross-section as a backtest. Realized window return (Δprice) and the margin of safety share the latest price in their construction, so an in-sample regression of one on the other would be mechanically confounded rather than predictive. The *proper* test is a monthly, point-in-time portfolio of the signal regressed on a factor model; we run exactly that in Section 8.7 and it returns no significant factor-adjusted alpha, which is consistent with the contrarian, ex-ante reading here. The honest read is unambiguous: the expectations gap is real and signed, and it had not paid, in level prices or in factor-adjusted returns, by the snapshot date.

6.6 The justification quadrant

Equation (2) is best read visually. Figure 6 plots the change in fundamental value against the change in price for every SaaS firm, and Table 10 counts the quadrants. Among FCF-positive firms, 45 names are over-corrected (value up, price down) versus 46 justified corrections, 37 earned re-ratings and 34 still-rich. The over-corrected names cluster in the upper-left and are disproportionately exposed incumbents.

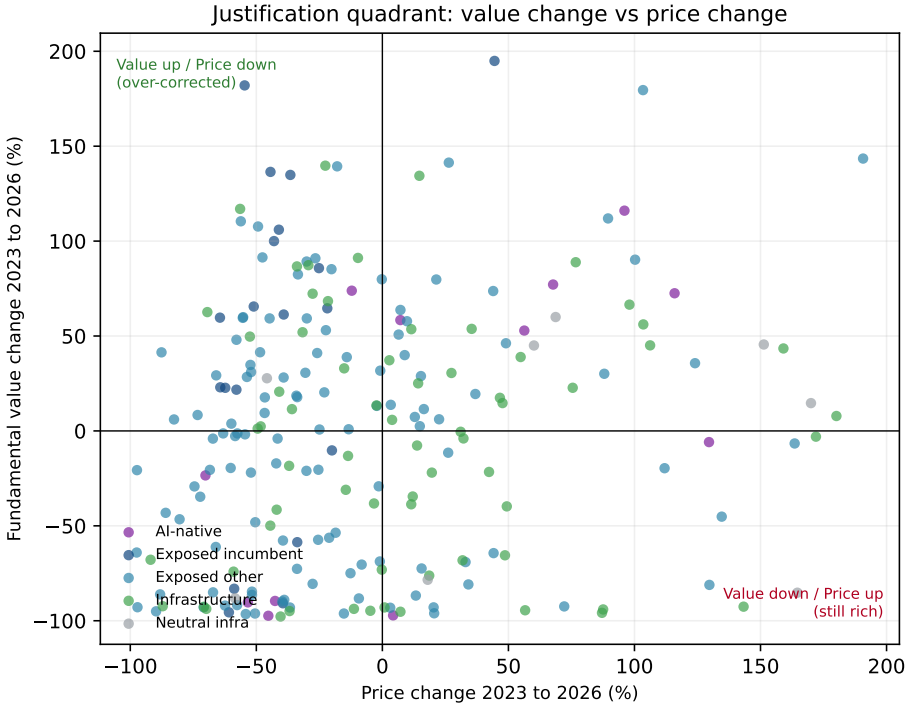

Figure 6. Justification quadrant. Horizontal axis: price change 2023 to 2026. Vertical axis: fundamental value change. Points colored by cohort; axes clipped for readability. Upper-left is the over-corrected opportunity; lower-right is the still-rich risk.

Table 10. Justification-quadrant counts by group (firms with a classified quadrant; unclassified omitted).

Quadrant	FCF-positive	Pre-FCF	All
Over-corrected (value up, price down)	45	46	91
Justified correction (both down)	46	27	73
Earned re-rating (both up)	37	41	78
Still rich (value down, price up)	34	13	47

6.7 The over-corrected screen

Table 11 lists the 17 investable (market cap $\geq$ \$1bn) names that are over-corrected and undervalued on triangulated fair value after bounding away model-sensitive tails (positive realised growth and MoS no greater than 200%). The list is dominated by profitable incumbents and infrastructure names, not by AI-native or pre-FCF firms. It is the raw screen before the moat, survival and exposure gates of Section 8; it is not the portfolio.

Table 11. Investable over-corrected and undervalued FCF-positive names (market cap ≥ $1bn), ranked by margin of safety. Raw screen, before portfolio gates.

Ticker	Cohort	Mkt cap ($bn)	Rev. CAGR (%)	SBC-FCF margin (%)	ROIC (%)	EV/FCF	MoS (%)
MNDY	Exposed incumbent	3.7	29	10	-46	22	98
SPSC	Infrastructure	2.1	17	14	114	18	83
GTM	Exposed other	1.8	1	21	5	12	68
NICE	Exposed incumbent	5.9	11	12	58	15	52
PCTY	Exposed incumbent	5.9	14	16	53	21	46
DUOL	Exposed other	4.8	38	24	59	14	42
GDDY	Infrastructure	11.5	8	26	38	11	41
RAMP	Infrastructure	2.3	12	10	81	23	25
QLYS	Neutral infra	3.2	10	31	61	14	22
ALRM	Exposed other	2.3	8	13	15	18	22
PAYC	Exposed incumbent	6.6	10	16	33	21	20
CHKP	Infrastructure	11.9	6	35	34	13	18
BLKB	Exposed other	1.8	1	14	60	16	15
DOCU	Exposed incumbent	9.0	8	15	32	17	11
BL	Exposed other	1.8	9	6	6	56	8
WDAY	Exposed incumbent	31.7	15	14	14	25	8
DBX	Exposed incumbent	6.7	0	27	21	14	2

6.8 Pre-FCF: a few survivors, many traps

The pre-FCF world is mostly uninvestable on these rules, but not uniformly. Reverse DCF on normalised earnings implies 14.6% growth against realised 11.3%, demanding expectations for firms that are still burning. Table 12 shows the handful of credible path-to-profit names after investability and sanity gates (market cap at least $500M, 5–60% realised revenue CAGR, under 25% annual dilution, ≥ 2 years runway, and positive MoS), while Table 13 shows the opposite end: decelerating, burning, sub-two-year-runway value traps.

Table 12. Credible pre-FCF path-to-profit candidates after investability and sanity gates: market cap ≥ $500M, 5–60% realised revenue CAGR, ≤ 25% annual dilution, ≥ 2 years runway, and positive path-to-profit margin of safety.

Ticker	Cohort	Mkt cap ($bn)	Rev. CAGR (%)	SBC-FCF margin (%)	Target margin (%)	Runway (yrs)	MoS (%)
ODD	Infrastructure	0.55	18	-10	14	3.7	226
LSPD	Exposed other	1.17	17	-5	8	7.4	87
ASAN	Exposed incumbent	1.54	10	-12	17	2.0	60
PLUS	Exposed other	2.32	6	-5	5	3.1	20
TEAM	Exposed incumbent	22.55	23	-6	16	3.2	16
RSKD	Exposed other	0.69	8	-3	10	14.3	2

Table 13. Pre-FCF value traps: decelerating revenue, deeply negative cash margin, under two years of runway.

Ticker	Cohort	Mkt cap ($bn)	Rev. CAGR (%)	Runway (yrs)	Dilution (%/yr)
LGCL	Exposed other	0.01	-16	0.0	19
QH	Exposed other	0.00	-17	0.0	302
FRGT	Exposed other	0.00	-12	0.1	1842
CISO	Infrastructure	0.01	-30	0.1	79
SURG	Exposed other	0.01	-30	0.1	12
SMSI	Exposed other	0.02	-32	0.1	56
ANY	Exposed other	0.01	-29	0.1	47
MYSE	Exposed other	0.01	-9	0.2	26
ARBE	Infrastructure	0.15	-1	0.2	22
CYN	Exposed other	0.02	-53	0.2	3222

Finding

The over-reaction is an *expectations* error, not a sector-wide discount. Reverse DCF shows the market over-expects AI-native growth (implied 43.6% vs realised 25.5%) and under-expects exposed-incumbent growth (implied 8.6% vs realised 13.8%). Level prices corroborate it: the same incumbent cohort is the only one with a positive, if modest, median margin of safety. Broad SaaS is not cheap; the mispricing is localized and signed.

7 AI Exposure, Moats and R&D Productivity

Section 6 located the valuation gap. Whether that gap is investable is a separate question. A cheap stock facing substitution is a trap; a cheap workflow incumbent using AI to deepen the product may be an opportunity. The difference is the sign of AI exposure.

7.1 Signed exposure

Exposure is not just "high" or "low": it has a sign. Positive exposure means AI enhances the product, workflow or moat. Negative exposure means AI substitutes or commoditises the product. Neutral exposure means AI is less central to the core economics. The sign depends on switching costs, data and workflow embeddedness, pricing power, and integration depth (Table 14).

Table 14. Exposure sign matrix. The same technology widens or erodes a moat depending on switching costs and whether AI improves or replaces the product.

	High switching cost	**Low switching cost**
AI improves product	moat widener (positive)	feature race (neutral)
AI replaces product	partial risk (mixed)	moat eroder (negative)

7.2 Talk is not conversion: an AI-pivot event study

A natural worry is that we are simply rewarding firms that talked about AI. We measure AI salience and stance from earnings-call and filing language, following Hassan et al. [4] and Loughran and McDonald [5], define "pivoters" as the top tercile of 2023–2024 AI salience (86

pivoters vs 170 controls), and compare their subsequent fundamental change in a difference-in-differences (DiD) design (Table 15). The result is sobering for the narrative: pivoters did not convert AI talk into superior cash economics. Their median SBC-adjusted FCF-margin and ROIC changes are statistically indistinguishable from controls, and their gross-margin change is significantly *worse* (DiD -2.7 pp, $t = -2.87$).

Margin Mechanics and Power. The negative gross margin DiD strongly suggests a compute tax: firms are absorbing hyperscaler inference costs into COGS before they can pass those costs to customers at scale. Additionally, the need to retain specialized AI talent risks structural SBC inflation, validating our choice to penalize SBC in the FCF calculation. On the non-significant FCF-margin DiD, we note a power limitation: for the non-missing SBC-adjusted FCF-margin outcome sample (79 treated and 162 controls), the 80%-power minimum detectable effect at $\alpha = 0.05$ is 12.5 pp. The observed effect is 2.5 pp, well below this threshold, so the null result is inconclusive rather than affirmative evidence of zero conversion; it rules out only a large immediate cash-flow windfall for pivoters.

Table 15. AI-pivot event study. Treatment = top-tercile 2023–2024 AI-call salience. Outcomes are changes from the 2023 baseline to the latest 2026 reading; the DiD mean uses deltas winsorized at the 2.5/97.5 percentiles. * denotes $p < 0.05$ (Welch t). Full design in Appendix B.

Outcome (Δ 2023→2026)	Pivoter med. (%)	Control med. (%)	DiD (mean) (%)	t	p
Gross margin	-0.5	+1.1	-2.7	-2.87	0.005*
SBC-adj FCF margin	+7.0	+6.3	+2.5	+0.55	0.583
ROIC	+9.6	+8.0	+22.1	+0.70	0.483

7.3 R&D productivity as a moat proxy

R&D is one factor here, not the whole paper. Building on the intangible-capital literature [8–11], we capitalise R&D with a perpetual-inventory stock and measure productivity as gross profit created per cumulative R&D dollar,

$$\text{R\&D Productivity}_i = \frac{\Delta\text{Gross Profit}_i}{\text{Cumulative R\&D Spend}_i}. \tag{5}$$

We use it as a durability and quality signal in the spirit of gross-profitability investing [12]; where Sehgal [13] prices R&D *intensity* across the market, here we test whether that spending *converts* into profit, with three disciplines stated up front: it is useful for durability; it is *not* by itself a cheapness signal; and part of any growth correlation is mechanical because gross-profit growth sits in the numerator. Table 16 makes the discipline concrete: higher R&D-productivity quartiles have higher revenue growth, but *not* higher margins of safety: the highest-productivity quartile is the most expensive, not the cheapest. R&D productivity identifies quality; it does not identify bargains (Figure 7).

Table 16. Outcomes by R&D-productivity quartile (FCF-positive firms). Quality rises with productivity; margin of safety does not; high productivity is priced in.

R&D productivity quartile	n	Med. prod.	ROIC (%)	SBC-FCF (%)	Rev. CAGR (%)	MoS (%)
Q1 low	67	-0.12	-21.5	-17.2	-5.4	-49.8
Q2	67	0.18	7.5	5.0	7.9	-26.8
Q3	67	0.4	7.9	5.3	16.4	-39.3
Q4 high	67	0.94	1.6	2.4	24.4	-53.0

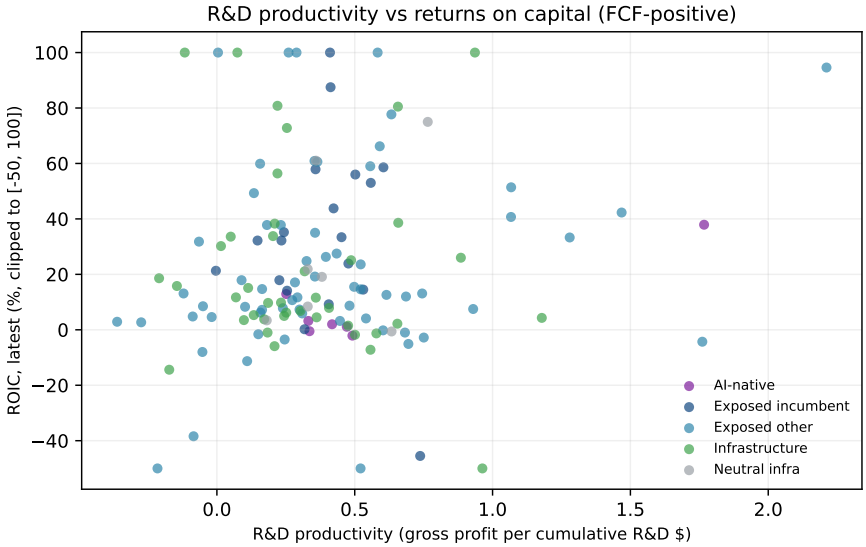

Figure 7. R&D productivity versus returns on capital (FCF-positive firms; ROIC clipped for readability). Productivity is a durability/quality signal, weakly related to current returns and not a standalone cheapness signal.

Case study

Duolingo (DUOL): why the framework does not buy every cheap, high-growth name. DUOL screens well on growth, Rule-of-40 and even the over-corrected screen. But its AI exposure sign is arguably *negative*: a general model can substitute for a consumer language-learning product, rather than deepen an enterprise workflow that is expensive to rip out. Signed exposure is what separates a consumer app facing substitution from an embedded enterprise system facing a feature race. DUOL therefore sits in the statistical-value tier, not the core-thesis tier, despite a strong screen.

8 Portfolio Construction

The previous sections narrowed the claim: not broad SaaS, not AI talk, not every cheap stock, not every high-R&D firm. What is left is the part a portfolio manager can act on. The question is no longer "is SaaS cheap" but "which names survive the full evidence chain strongly enough to become a watchlist."

8.1 Gates and score

A core long candidate must pass all of: (i) FCF-positive status; (ii) exposed-incumbent / workflow software classification rather than broad software cheapness; (iii) improving SBC-adjusted FCF margin; (iv) positive or neutral text stance; (v) positive base-case margin of safety; (vi) a moat/quality gate (R&D productivity ≥ 0.25 or Rule-of-40 $\geq 20\%$); (vii) no material survival, dilution or leverage red flag; and (viii) investable size. The raw over-corrected and path-to-profit screens remain diagnostic tables; they do not automatically enter the core portfolio. An avoid/short candidate has negative signed exposure, weak or deteriorating fundamentals, a rich valuation, low switching costs, or pre-FCF burn with low runway and high dilution. Conceptually the investable object is a product, not a sum:

$$\text{Investable Mispricing}_i = \text{Mispricing}_i \times \text{ExposureSign}_i \times \text{Moat}_i \times \text{Survival}_i \times \text{ValuationConfidence}_i. \tag{6}$$

A composite score makes the trade-offs explicit (Table 17); the weights are a stated prior, not an estimate, and the ranking is robust to reasonable reweighting.

Table 17. Composite portfolio score: components and stated weights.

Component	Weight	Reason
Margin of safety	35%	value-investor anchor
Fundamental trajectory	20%	confirms reality
Moat / resilience	25%	determines exposure sign
Survival / balance sheet	10%	protects downside
Valuation uncertainty	10%	avoids false precision

8.2 First-principles overlay from filings

To reduce model assumptions, we read the actual filings for the portfolio candidates rather than inferring. Table 18 reports source coverage for the 32 candidate firms. The honest result is a mix: 32/32 annual filings fetched and 32/32 with explicit AI-risk language, but only 6/32 disclose net revenue retention, 5/32 disclose customer concentration, 1/32 disclose undrawn revolver availability, and 0/32 disclose a strict contract term. We do not infer the missing fields. The teachable limitation is that AI-risk language is now near-universal, so its mere existence does not differentiate; the useful signal is the exact sentence and whether disclosed retention or concentration confirms the moat thesis (Table 19).

Table 18. First-principles overlay: source coverage across the 32 portfolio candidates. Fields are populated only from disclosed language; missing values are not inferred.

First-principles field	Source rule	Coverage
Annual filing fetched	SEC 10-K / 20-F / 40-F	32/32
Explicit AI-risk language	filing text, raw sentence stored	32/32
Net revenue retention	disclosed retention only	6/32
Customer concentration	disclosed concentration only	5/32
Contract term	strict term language	0/32
Revolver availability	disclosed undrawn facility	1/32
Long-term growth (LTG)	vendor consensus	0/32 (no source)

Table 19. Disclosed datapoints extracted verbatim from filings (the firms that actually disclose). Operators: > lower-bound, ∼ approximate.

Disclosed net revenue retention	
Ticker	Retention
FIVN	105%
GDDY	91%
GTM	90%
PAYC	91%
PCTY	>92%
WDAY	∼97%
Disclosed customer concentration (top customers)	
Ticker	Concentration
FRSH	5%
PD	2%
RAMP	30%
RNG	10%
SPSC	1%

8.3 Ranked candidates

Table 20 ranks only the core-thesis candidates that clear the stricter gates: FCF-positive, exposed-incumbent, improving, non-negative text stance, investable, and positive base-case margin of safety. Tier 1 is cheap even on the conservative valuation band; Tier 2 is cheap on the base case only. Figure 8 shows the wider candidate fair-value bands against current price for context. The list is the reproducible output of the framework, not a set of tips.

Table 20. Core-thesis ranked candidates after stricter gates. "Up. (base)"/"Up. (cons.)" are base-case and conservative upside to fair value; NRR shows the disclosed operator; "Conc." is disclosed top-customer concentration; "Text flag" is the call/filing stance overlay (+ supportive, 0 neutral). Dashes denote no disclosure.

Ticker	Cohort	Up. (base) (%)	Up. (cons.) (%)	Rev. CAGR (%)	Rule-40 (%)	R&D prod.	NRR	Conc.	Text flag
Tier 1: core thesis candidates, cheap on conservative valuation (n=2)									
FRSH	Exposed incumbent	65	23	18	28	0.53	–	5%	+
DOCU	Exposed incumbent	11	11	8	23	0.23	–	–	0
Tier 2: core thesis watchlist, cheap on base case only									
PCTY	Exposed incumbent	46	-12	14	30	0.56	>92%	–	0
WDAY	Exposed incumbent	8	-6	15	29	0.25	∼97%	–	+

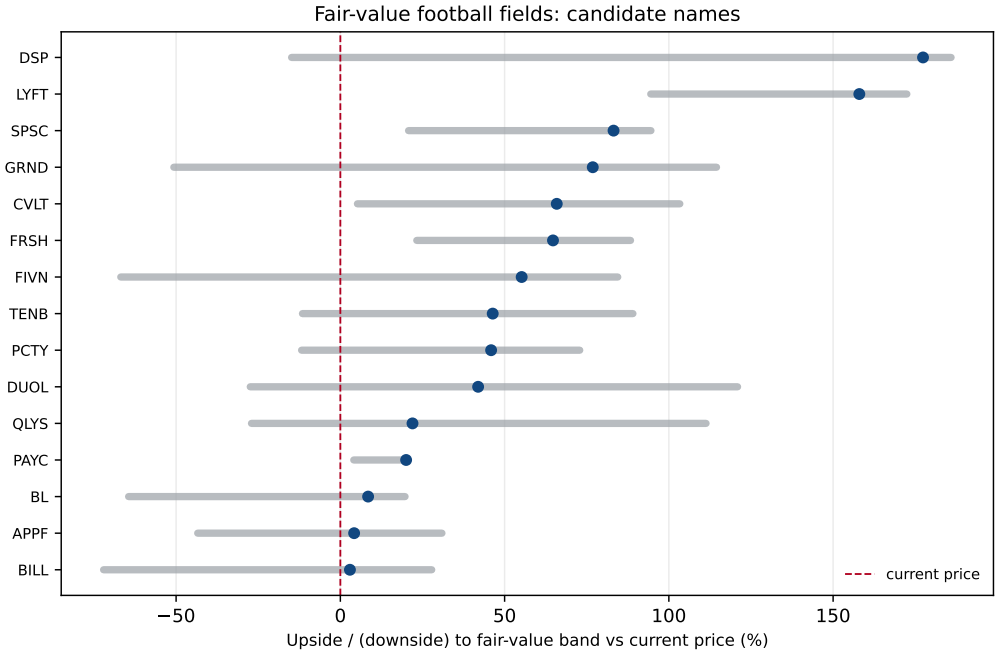

Figure 8. Fair-value football fields for candidate names. Grey bars span the bear–bull fair-value band relative to current price; the dot is the base-case fair value. Names left of the dashed line trade above their base-case fair value.

8.4 Valuation-gap profile and naive benchmarks

To evaluate whether the gates add practical underwriting value, we benchmark the core watchlist against naive screens (Table 22) and compute expected upside profiles (Table 21). The naive screens can show higher raw median MoS, but they do so by admitting names outside the signed-exposure thesis or with weaker business-model fit. The core gates deliberately trade some raw upside for thesis purity: FCF-positive status, exposed-incumbent workflow economics, improving cash conversion, non-negative text stance, moat quality, low dilution and investable size. Implementation should apply the ℓ_1 turnover penalty framework of Novy-Marx and Velikov [22]; the current two-endpoint snapshot does not itself estimate realised turnover.

Table 21 should therefore be read as a valuation-gap profile, not as a return forecast. A name can close the gap, remain cheap, or prove that the fair-value estimate was too high. The research contribution is the auditable underwriting sequence and the kill conditions in Table 26, not a claim that the watchlist has already produced alpha.

Table 21. Valuation-gap profile for the core watchlist. Upside is to estimated fair value, not a forecast realised return.

Watchlist Group	Firms	Median Base Upside	Median Cons. Upside	Win Rate (Base > 0)
Tier 1 (Core)	2	37.6	16.9	100.0
Tier 2 (Watchlist)	2	27.0	-8.9	100.0
Combined Core Watchlist	4	28.2	2.2	100.0

Table 22. Naive benchmark horse-race. Naive screens can show higher raw MoS; the gated watchlist prioritises thesis fit, signed exposure and investability over raw cheapness.

Metric	Gated Core Watchlist	Naive: FCF-Pos, MoS > 0	Naive: Rule-of-40 $\geq 20\%$
Firms	4	32	22
Median MoS	28.2	41.6	39.0
% Undervalued (Base)	100.0	100.0	100.0
Median R&D Prod.	0.39	0.3	0.41
Median Dilution	-3.1	-2.5	-2.3

8.5 Industry concentration

The gated core list is small, so concentration risk is worth stating plainly. The four Tier 1/Tier 2 names (Table 20) are not clustered in a single sub-vertical: DOCU is horizontal agreement/workflow software, FRSH is IT and customer-experience workflow, and PCTY and WDAY are human-capital and back-office systems. The wider investable over-corrected screen (Table 11) diversifies further into infrastructure and security names such as SPSC, QLYS, CHKP and GDDY. This spread reduces the risk that the valuation gap is a single-vertical headwind masquerading as a firm-level opportunity, though the four-name core is too small to rely on for diversification by itself.

8.6 How a portfolio reader should use the output

The portfolio output has three uses. First, it identifies a *long-research list*: profitable workflow incumbents where the model says price is discounting more impairment than the accounts show. Second, it identifies an *avoid list*: cash-burning, dilutive or substitution-exposed firms where cheapness is not enough. Third, it defines an *update protocol*: when the next filings arrive, the reader should refresh the same gates: SBC-adjusted FCF margin, ROIC, realised growth, dilution, runway, disclosed retention and reverse-DCF implied growth, rather than re-arguing the narrative. That is the practical value of the framework. It turns a noisy AI story into a repeatable underwriting checklist.

Quarterly Monitoring Checklist. Each earnings season, the thesis requires refreshing:

1. **SBC-adjusted FCF margin delta:** Is cash generation still improving?
2. **ROIC vs WACC spread:** Is the firm earning above its cost of capital?
3. **Realised vs Implied Growth:** Did revenue growth cross below reverse-DCF implied growth?
4. **Dilution rate:** Is share issuance accelerating to retain talent?
5. **Disclosed NRR:** Did net revenue retention break below stated targets?
6. **AI-risk language evolution:** Did the filing text pivot from feature integration to competitive threat?

Case study

Freshworks (FRSH): what a framework-consistent opportunity looks like. FRSH is the cleanest example of the core thesis. It is FCF-positive with improving margins, clears the conservative valuation band, and carries a supportive (+) call/filing stance in our text overlay, consistent with AI framed as a productivity layer rather than a replacement. Its disclosed top-customer concentration is low (5%; Table 19), insulating it from single-client churn. Unlike pre-FCF flyers or expensive AI-natives, it represents an embedded workflow

incumbent priced as if its moat is eroding while the statements show durability.

Name-level reads. A handful of other names clarify the framework's boundaries. **DOCU** clears the conservative band but with a smaller margin of safety. **PCTY** and **WDAY** remain watchlist names: they satisfy the core economic thesis but do not clear the conservative valuation band. Conversely, several names that screen statistically cheap on the over-corrected table (Table 11) are still excluded from the core-thesis list because they do not match the workflow-embedded enterprise-SaaS exposure being underwritten: **DUOL** is consumer-facing and substitution-exposed, while **SPSC**, **QLYS** and **CHKP** are infrastructure or security names rather than seat- or workflow-priced incumbents.

8.7 An out-of-sample factor-spanning test: does the signal pay?

The cross-sectional valuation gap is an ex-ante statement. The decisive question for a portfolio manager is whether a systematic version of the signal earns risk-adjusted returns. We test it directly. Using the point-in-time firm-quarter signal (a resilient-cheap composite, degrossed by the disruption posterior, $\alpha_i = (1 - 2\hat{p}_i^{\text{dis}})(y_i + D_i)$, where $\hat{p}_i^{\text{dis}}$ is the smoothed Markov regime-switching posterior of Hamilton [26] riding on a Kalman-filtered [27] latent resilience state, and y_i and D_i are the durability and priced-discount components), we form a sector-neutral, monthly-rebalanced quintile long/short portfolio, long resilient-and-cheap names, short disrupted/expensive ones, mapping each firm's signal to the return grid only *after* its filing date plus a 60-day lag. Monthly returns are adjusted closes for the 386 of 473 universe names with public price histories (81.6% coverage; source and survivorship caveat in Appendix C). We then regress the long/short spread on the Fama–French five factors plus momentum with Newey–West [28] standard errors (Table 23).

The result is a clean null, and it is the honest one. The raw long/short spread is economically small and statistically indistinguishable from zero (3.8% annualised, Newey–West $t = 0.24$, 38 months; long-only tilt versus equal-weight software 4.5%, $t = 0.45$). After factor adjustment the intercept is -14.0% annualised ($t = -0.97$) over 31 months, with no statistically significant factor loadings and an R^2 of only 0.035; so the sleeve is neither a disguised factor bet nor a source of alpha over this window. This corroborates, with a proper factor model, what the cross-section already implied: the expectations gap is real and signed, but it had *not* been harvestable as factor-adjusted alpha by the data freeze. We grant the signal full-sample parameter estimation, so this is a best-case feasibility check rather than a pseudo-out-of-sample forecast; even so, it does not pay. A paper that claimed alpha here would be overfitting a 31-month window.

Table 23. Out-of-sample factor-spanning test. The point-in-time resilient-cheap long/short spread regressed on FF5 + momentum (Newey–West SE). The test is whether the intercept $\alpha \neq 0$; it is not. * denotes $|t| \geq 1.96$.

Factor	Loading	t
α (monthly)	-1.17%	(−0.97)
α (annualised)	-14.0%	
MKT-RF	0.332	(0.63)
SMB	-0.084	(−0.23)
HML	-0.362	(−0.66)
RMW	0.249	(0.31)
CMA	-0.140	(−0.24)
MOM	-0.200	(−0.38)
R^2	0.035	
Months	31	

Finding

The strongest paper-consistent angle is a dispersion trade, not a sector call: own the over-discounted, cash-generative, moat-widening incumbents where price implies impairment; avoid the cash-burners, the moat-losers, the broad statistical screens that do not match the thesis, and the richly priced AI-native narratives. The stricter core list is small by construction.

9 A Generalized Narrative-Shock Framework

SaaS-AI is the worked example, not the only use case. The contribution is a general, factorized way to analyse short-horizon narrative shocks before the full long-term outcome is visible. We call it the *Signed-Exposure Intrinsic-Value Repricing* framework. It has six stages: (1) define the narrative shock; (2) build the full investable universe with a transparent funnel; (3) measure business reality across the six factor families; (4) estimate intrinsic value independently of price, routed by business type; (5) attach a *signed* exposure; and (6) convert to portfolio action via Equation (6).

The structure exports cleanly. For each shock the analyst asks the same three questions: an exposure question, a moat question, and a portfolio question (Table 24). The mechanics of Equations (2)–(6) are unchanged; only the exposure sign and the survival constraints are re-specified per shock.

Table 24. Exporting the framework to other narrative shocks. The valuation and gate mechanics are identical; the exposure sign and survival constraints are re-specified per shock.

Shock	Exposure question	Moat question	Portfolio question
AI in SaaS	Does AI substitute or enhance workflow software?	Are switching costs real?	Which firms were overdiscounted?
Quantum	Does it break security or enable new infrastructure?	Who sells protection vs. who is exposed?	Long enablers, avoid exposed legacy
Robotics	Does automation replace labor or create operating leverage?	Who owns data, hardware, distribution, service lock-in?	Long scalable adopters and enablers
Regulation	Does it impose cost or entrench scale leaders?	Who absorbs compliance better?	Long advantaged incumbents

Short-horizon caveat. In 2023–2026 data we observe the early fundamental response and the valuation gaps, but not the full five-to-seven-year realised-alpha cycle. The framework can therefore demonstrate ex-ante mispricing and a valuation-logic portfolio advantage; it must not claim fully proven long-run realised alpha without an out-of-sample backtest.

10 Robustness, Falsification and Limitations

10.1 Robustness

The headline conclusions are designed to survive reasonable variation in modelling choices. We re-run the valuation under (i) alternative discount rates: a flat 10% WACC, the per-firm CAPM WACC, and a +200 bps stress; (ii) alternative fair-value weights: DCF-only, multiple-only, and median triangulation; (iii) alternative cohort definitions: strict incumbents only versus the broad exposed cohort; (iv) alternative SBC treatment: as-reported versus SBC-burdened FCF and dilution-adjusted per-share value; (v) liquidity screens at \$50M, \$500M and \$1B; and (vi) time-window sensitivity across baselines. The two-worlds split and the localized exposed-incumbent gap are robust to these choices; what moves is the precise width of the gap, not its location. Appendix E states each alternative and its expected direction of effect.

Table 25. Observed robustness matrix from the frozen snapshot. Gaps are median valuation gaps under alternative lenses or screens; the table is descriptive and complements the qualitative specification discussion in Appendix E.

Specification	Firms	Median gap (%)	% positive gap
Baseline triangulated MoS (all FCF-positive)	167	-17.7	38.9
Baseline triangulated MoS (exposed incumbents)	19	2.9	57.9
Market cap $\geq$ \$500M	134	-22.1	34.3
Market cap $\geq$ \$1B	121	-23.0	33.9
Market cap $\geq$ \$5B	65	-40.3	21.5
DCF-only EV lens	167	-21.3	37.7
Normalized-DCF EV lens	167	-4.5	47.9
Peer-multiple EV lens	167	-5.5	46.1
Strict incumbent-only screen	19	2.9	57.9

10.2 Identification: the AI narrative versus the rate shock

The first-order threat to interpretation is that 2023–2026 SaaS compression reflects high interest rates punishing long-duration cash flows rather than anything AI-specific. The rate regime undoubtedly lowered the sector's valuation ceiling, but a discount-rate shock is a *common factor*: to first order it derates all long-duration cash flows in the same direction. It therefore cannot generate the two cross-sectional features that carry this paper. First, the reverse-DCF expectations gap is *signed and opposite across cohorts within the same rate environment*: the market implies 43.6% growth for AI-native firms (against realized 25.5%) but only 8.6% for exposed incumbents (against realized 13.8%). A uniform duration shock cannot make implied growth too *high* for one cohort and too *low* for another at the same time; an AI-narrative reallocation of expectations can. Second, the AI-pivot event study compares firms that face the *same* rates and differ only in AI salience, and finds the high-salience group's gross margin fell relative to controls (DiD -2.7 pp, $t = -2.87$): a cost/substitution signature, not a discount-rate one. Rates set the level; the AI narrative set the cross-sectional dispersion, and it is the dispersion, not the level, that this paper identifies and trades against. We further stress the discount rate directly (flat 10% WACC and a +200 bps shock) in the robustness matrix above, and the cohort ordering is preserved.

10.3 Falsification

Because H1–H5 are stated as refutable claims (Section 2), the central thesis is explicitly killable. It weakens if FCF-positive incumbents later show margin compression, if realised growth falls below reverse-DCF implied growth, if R&D productivity fails to convert into ROIC or FCF, if AI-native firms sustain superior fundamentals while incumbents lag, or if the long list stays cheap because fair value was overestimated. We state these conditions in advance, consistent with the multiple-testing discipline of Harvey et al. [24] and the false-discovery-rate control of Benjamini and Hochberg [25].

Table 26. Quarterly falsification and monitoring map. These are not trading signals; they are the conditions under which the paper's ex-ante valuation-gap thesis should be reduced, paused or killed.

Monitor	Thesis-supporting evidence	Thesis-weakening evidence
SBC-adjusted FCF margin	Flat to improving cash conversion after SBC.	Margin compression or cash conversion reversal.
ROIC versus WACC	ROIC remains above or moves toward cost of capital.	ROIC falls below WACC or improvement reverses.
Realised versus implied growth	Realised growth remains above reverse-DCF implied growth.	Realised growth falls below implied growth; price no longer under-expects the business.
Dilution and SBC	Dilution stable or falling while FCF improves.	Share issuance accelerates or SBC absorbs the operating leverage.
Retention and customer risk	NRR / concentration disclosures support workflow embeddedness.	NRR falls, concentration risk rises, or churn indicators worsen.
AI-risk language	Filings describe AI as product enhancement / efficiency layer.	Filings shift toward substitution, pricing pressure or competitive displacement.

10.4 Limitations

This study is built to be transparent and replicable, but several limitations remain and should be weighed when interpreting the results:

1. **Horizon.** This is a short-horizon study (median span 2.25 years) of a still-unfolding correction. It identifies where price appears to have moved further than value, but the re-rating is not yet realised; the portfolio output is an *ex-ante* valuation gap, not demonstrated realized alpha.
2. **Factor test: a null, with three caveats.** We do run the monthly factor-spanning test (Section 8.7) and find no significant factor-adjusted alpha. Three honest caveats apply. (i) *Price source and survivorship*: the return panel uses public adjusted closes (Yahoo) for the 81.6% of names with histories, so firms acquired or taken private during the window are absent and the test carries a survivorship tilt (Appendix C). (ii) *Window*: the factor overlap is only 31 months, which is short for a powered alpha test. (iii) *In-sample signal*: the signal's model parameters are estimated over the full sample, so this is a best-case feasibility check, not a true pseudo-out-of-sample forecast. None of these manufacture alpha; the point estimate is negative and insignificant, but they bound how strongly the null can be read.
3. **Data tier.** Fundamentals and prices come from a single accessible tier (Tier-1 (Sharadar + SEC EDGAR + Ken French)), not the CRSP/Compustat academic standard. Coverage, delisting treatment and historical corporate actions are less complete; we mitigate via point-in-time fields and a frozen, hashed snapshot, but do not claim CRSP-grade survivorship handling (Appendix C).
4. **Sparse disclosure.** Several first-principles fields are disclosed by only a minority of firms: net revenue retention (6/32), customer concentration (5/32) and contract terms (0/32). We use them where available and *do not infer* them elsewhere, which limits the overlay's cross-sectional power.
5. **No consensus growth input.** We have no local consensus long-term-growth feed, so terminal growth is a disciplined assumption rather than an observed input; this is the single largest source of fair-value uncertainty and is the reason valuation is triangulated and banded.
6. **Unobservable operating signals.** Real churn, win/loss, pricing actions and channel checks are not in public filings and remain a follow-up underwriting requirement before any position.
7. **Model sensitivity at the tails.** Extreme small-cap margins of safety are model-sensitive; the size floors ($500M / $1bn) exist precisely to keep the investable conclusions away from those tails.
8. **Survivorship and Private Equity Exits.** By requiring a 2026 endpoint, the study systematically drops firms that were acquired and delisted during the window. Between 2023 and 2026, private equity sponsors acquired software firms that fit the core thesis: compressed multiples but sticky, FCF-positive workflows. Missing those takeover outcomes means the historical opportunity set may *understate* the potential returns of the resilient-incumbent cohort.
9. **Not advice.** Nothing here is investment advice; cohorts and names illustrate the framework, not recommendations.

11 Replicability

An independent reader can regenerate every number, table and figure in this paper from the source data. The design follows a single rule: **no headline number is hand-typed.** All

reported figures resolve from `data/metrics_v2.tex`, which is generated deterministically by the `paper_assets` build package from three frozen artifacts: the firm-level valuation run, the first-principles overlay, and the AI-text panel, identified by snapshot hash `8a7ec05c7873`.

Deterministic pipeline. The build is idempotent: re-running it on the same snapshot reproduces byte-identical macros, tables and figures. Updating the snapshot and rebuilding updates the entire manuscript at once, so prose and data cannot silently diverge.

Steps. A reader with the snapshot reproduces the manuscript as follows:

1. Build the assets:
   `PYTHONPATH=scripts python3 scripts/saas_ai/analysis/build_paper_assets.py`
2. Compile the manuscript:
   `tectonic paper_latex/saas_ai_repricing/paper_v2.tex` (or any standard LaTeX distribution).
3. Audit the headline claims:
   `python3 scripts/saas_ai/analysis/validate_final_findings.py` (hard-fails on any drift).
4. Run the unit tests for the computational logic:
   `python3 -m pytest scripts/saas_ai/analysis/test_paper_assets.py`

What the repository contains. The public research repository at https://gitlab.com/vastdreams/saas-ai-repricing-research contains the published PDF and LaTeX source, the asset builder (`paper_assets`: `loaders`, `macros`, `params`, `tables`, `figures`, `eventstudy`); validation scripts; unit tests; generated publication tables and figures; data schemas; and rebuild instructions. It deliberately does *not* redistribute raw vendor data or the frozen vendor-derived CSV snapshot. Machine-readable provenance is in `SAAS_DATA_PROVENANCE.md`, data access and licensing in `SAAS_DATA_AVAILABILITY.md`, and citation metadata in `SAAS_CITATION.cff`. Rebuilding the snapshot from scratch requires the reader's own Sharadar/Nasdaq Data Link subscription and configured research database (Appendix C).

12 Conclusion

The repricing was largely consistent with fundamentals, with one systematic exception. Compressing the multiples of software firms that were still burning cash, diluting shareholders, and depending on future capital markets was consistent with their deteriorating economics, and AI-native firms were not, on the whole, left cheap. The exception is the cohort of profitable, improving, workflow-embedded incumbents whose business economics held up better than their prices; that cohort is the investable focus of the paper.

The result is not a sector call and not a proof of realised alpha. It is a first-principles underwriting map. First measure whether the business got worse. Then value it independently of price. Then ask what growth the market is implying. Then decide whether AI widens or erodes the moat. Only then build a watchlist. In this sample, that process rejects broad SaaS, rejects most pre-FCF names, rejects AI talk as a standalone signal, rejects high R&D productivity as a cheapness signal, and leaves a small incumbent watchlist.

The two lead findings are an expectations gap and a conversion test; the remaining results provide the backdrop and the price-space confirmation. Mapped to the hypotheses of Section 2:

1. **Under-expected incumbents: the central result (H3).** Reverse DCF shows the repricing mis-set expectations with opposite signs: AI-native firms are priced for more growth than they realised (implied 43.6% vs realised 25.5%), while exposed workflow incumbents are priced for less (implied 8.6% vs realised 13.8%). The market over-paid the narrative and under-paid the cash flow.
2. **Talk is not conversion (H4).** High-salience "pivoters" did not convert AI language into superior economics; their gross-margin change was significantly worse than controls (DiD -2.7 pp, $t = -2.87$). AI salience is a cost signal before it is a moat signal.
3. **Two worlds (H1).** The 167 FCF-positive firms improved (median SBC-adjusted FCF margin 5.5%→11.6%, ROIC 5.7%→14.0%) while the 168 pre-FCF firms deteriorated (FCF margin -27.8%→-33.6%, median runway 0.97 years). "AI hurt SaaS" is the wrong unit of analysis.
4. **Localized over-reaction, price-space confirmation (H2).** Broad SaaS is not cheap (median MoS -17.7% FCF-positive, -84.4% pre-FCF); the only cohort with a positive, if modest, median margin of safety is the same exposed-incumbent group (19 firms, median MoS 2.9%), which corroborates the expectations gap rather than standing alone.
5. **Quality is priced (H5).** R&D productivity identifies durable firms but not cheap ones; the highest-productivity quartile is the most expensive.

Summary. The results do not support a general claim that AI benefits SaaS or that SaaS is cheap. They indicate that AI exposure differentiated software businesses by moat quality: the repricing was consistent with fundamentals for weaker firms but over-discounted the stronger, workflow-embedded incumbents. The investable opportunity is confined to that second group.

Acknowledgments

This research was independently conceived and funded by FSE Research & Investments Pty Ltd, of which the author is the founder and sole researcher, with support from Eye Dream Pty Ltd trading as Finsoeasy (ABN 95 650 714 060). The author thanks the open-source scientific Python and LaTeX communities, and the maintainers of the Kenneth French data library, for the tools that made this analysis possible. All errors and omissions are the author's own.

Data Availability Statement

The public landing page for this research is https://research.finsoeasy.com. The open repository at https://gitlab.com/vastdreams/saas-ai-repricing-research contains the published PDF, LaTeX source, analysis code, generated publication tables/figures, schemas and rebuild instructions documented in `SAAS_DATA_AVAILABILITY.md`. It does *not* redistribute raw vendor data or the frozen vendor-derived CSV snapshot (`8a7ec05c7873`) because Sharadar/Nasdaq Data Link data must be obtained under each reader's own licence. Readers who wish to recompute every reported value may supply their own licensed fundamentals/prices, public SEC EDGAR filings, Kenneth French factors and public adjusted-close return panel, then run the documented build and validation pipeline in `SAAS_DATA_AVAILABILITY.md`.

How to Cite

APA.

Sehgal, A. (2026). *Did Fundamentals Justify the AI Repricing of Enterprise Software? Evidence*

from 2023–2026 (Working Paper). Eye Dream Pty Ltd (trading as Finsoeasy) / FSE Research & Investments Pty Ltd.

BibTeX.

```
@techreport{sehgal2026airepricing,
  author      = {Sehgal, Abhishek},
  title       = {{Did Fundamentals Justify the AI Repricing of
                 Enterprise Software? Evidence from 2023--2026}},
  year        = {2026},
  institution = {Eye Dream Pty Ltd (trading as Finsoeasy) /
                 FSE Research \& Investments Pty Ltd},
  type        = {Working Paper},
  url         = {https://research.finsoeasy.com},
  note        = {ORCID: 0009-0000-9424-4695}
}
```

Machine-readable citation metadata is provided in `SAAS_CITATION.cff`.

Disclosure Statement

The author is the founder of FSE Research & Investments Pty Ltd, which funded this research. This paper is for educational and informational purposes, does not constitute investment advice, and is not a recommendation to buy or sell any security. Cohorts and individual names are used to illustrate the framework, not as recommendations. The author has no other conflicts of interest to disclose.

Use of AI tools. Generative AI tools were used to assist with editing, prose clarity, formatting and code-level refactoring. All research design, data construction, modelling choices, numerical results and conclusions are the author's own. Every reported figure resolves deterministically from the validated data pipeline rather than from any language model, and the author has reviewed and verified the full manuscript and all results. The author takes full responsibility for the content.

References

[1] Bruce C. N. Greenwald, Judd Kahn, Paul D. Sonkin, and Michael van Biema. *Value Investing: From Graham to Buffett and Beyond*. John Wiley & Sons, New York, NY, 2001.

[2] Aswath Damodaran. *Investment Valuation: Tools and Techniques for Determining the Value of Any Asset*. John Wiley & Sons, Hoboken, NJ, 3rd edition, 2012.

[3] Alfred Rappaport and Michael J. Mauboussin. *Expectations Investing: Reading Stock Prices for Better Returns*. Harvard Business School Press, Boston, MA, 2001.

[4] Tarek A. Hassan, Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. Firm-level political risk: Measurement and effects. *The Quarterly Journal of Economics*, 134(4):2135–2202, 2019. doi: 10.1093/qje/qjz021.

[5] Tim Loughran and Bill McDonald. When is a liability not a liability? textual analysis, dictionaries, and 10-ks. *The Journal of Finance*, 66(1):35–65, 2011. doi: 10.1111/j.1540-6261.2010.01625.x.

[6] Tania Babina, Anastassia Fedyk, Alex He, and James Hodson. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745, 2024. doi: 10.1016/j.jfineco.2023.103745.

[7] Daron Acemoglu and Pascual Restrepo. Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244, 2020. doi: 10.1086/705716.

[8] Baruch Lev and Theodore Sougiannis. The capitalization, amortization, and value-relevance of r&d. *Journal of Accounting and Economics*, 21(1):107–138, 1996. doi: 10.1016/0165-4101(95)00410-6.

[9] Louis K. C. Chan, Josef Lakonishok, and Theodore Sougiannis. The stock market valuation of research and development expenditures. *The Journal of Finance*, 56(6):2431–2456, 2001. doi: 10.1111/0022-1082.00411.

[10] Ryan H. Peters and Lucian A. Taylor. Intangible capital and the investment-q relation. *Journal of Financial Economics*, 123(2):251–272, 2017. doi: 10.1016/j.jfineco.2016.03.011.

[11] Andrea L. Eisfeldt and Dimitris Papanikolaou. Organization capital and the cross-section of expected returns. *The Journal of Finance*, 68(4):1365–1406, 2013. doi: 10.1111/jofi.12034.

[12] Robert Novy-Marx. The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1):1–28, 2013. doi: 10.1016/j.jfineco.2013.01.003.

[13] Abhishek Sehgal. R&D Alpha: Investment Intensity and Long-Term Stock Returns. Working paper, FSE Research & Investments Pty Ltd, 2026. URL https://ssrn.com/abstract=6002295. Research landing page: https://research.finsoeasy.com.

[14] Andrea L. Eisfeldt, Gregor Schubert, and Miao Ben Zhang. Generative ai and firm values. NBER Working Paper 31222, National Bureau of Economic Research, 2023.

[15] Feng Li. The information content of forward-looking statements in corporate filings: A naïve bayesian machine learning approach. *Journal of Accounting Research*, 48(5):1049–1102, 2010. doi: 10.1111/j.1475-679X.2010.00382.x.

[16] Gerard Hoberg and Gordon Phillips. Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5):1423–1465, 2016. doi: 10.1086/688176.

[17] Eugene F. Fama and Kenneth R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993. doi: 10.1016/0304-405X(93)90023-5.

[18] Eugene F. Fama and Kenneth R. French. A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22, 2015. doi: 10.1016/j.jfineco.2014.10.010.

[19] Mark M. Carhart. On persistence in mutual fund performance. *The Journal of Finance*, 52(1):57–82, 1997. doi: 10.1111/j.1540-6261.1997.tb03808.x.

[20] Eugene F. Fama and James D. MacBeth. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3):607–636, 1973. doi: 10.1086/260061.

[21] Clifford S. Asness, Andrea Frazzini, and Lasse Heje Pedersen. Quality minus junk. *Review of Accounting Studies*, 24:34–112, 2019. doi: 10.1007/s11142-018-9470-2.

[22] Robert Novy-Marx and Mihail Velikov. A taxonomy of anomalies and their trading costs. *The Review of Financial Studies*, 29(1):104–147, 2016. doi: 10.1093/rfs/hhv063.

[23] Robert F. Stambaugh, Jianfeng Yu, and Yu Yuan. The short of it: Investor sentiment and anomalies. *Journal of Financial Economics*, 104(2):288–302, 2012. doi: 10.1016/j.jfineco.2011.12.001.

[24] Campbell R. Harvey, Yan Liu, and Heqing Zhu. ... and the cross-section of expected returns. *The Review of Financial Studies*, 29(1):5–68, 2016. doi: 10.1093/rfs/hhv059.

[25] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57(1):289–300, 1995. doi: 10.1111/j.2517-6161.1995.tb02031.x.

[26] James D. Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384, 1989. doi: 10.2307/1912559.

[27] Rudolph E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960. doi: 10.1115/1.3662552.

[28] Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987. doi: 10.2307/1913610.

[29] Peter J. Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964. doi: 10.1214/aoms/1177703732.

A Reproducibility and Audit

Every headline number resolves from a frozen, validated artifact, not from prose. The macros in `data/metrics_v2.tex` are generated by `scripts/saas_ai/analysis/build_paper_assets.py` (the `paper_assets` package) directly from `data/fundamental_value_run.csv`, `data/first_principles_overlay.csv` and `data/text_panel.csv`. The headline figures are independently asserted by `scripts/saas_ai/analysis/validate_final_findings.py`, which hard-fails if any claim drifts from the data, and the computational logic (event study, macro/params builders) carries unit tests in `scripts/saas_ai/analysis/test_paper_assets.py`.

Audit checklist. To reproduce: (1) confirm 335 SaaS rows after the payments carve-out and the 167/168 split; (2) recompute the median trajectories in Tables 5–7; (3) recompute the cohort margins of safety in Table 9, including the exposed-incumbent cohort (19 firms, median MoS 2.9%); (4) recompute the quadrant counts in Table 10, the reverse-DCF table (Table 8) and the event study (Table 15); (5) confirm overlay coverage in Table 18; and (6) run the validation script and confirm a passing status. All figures are produced by the same build, so they cannot diverge from the tables.

Table 27. Full SaaS firm-level manifest from the frozen snapshot.

Ticker	Cohort	Group	Mkt cap ($bn)	Rev. CAGR (%)	SBC-FCF (%)	ROIC (%)	MoS (%)
AI	AI-native	pre-fcf	1.35	2	-128	-59	-94
CFLT	AI-native	pre-fcf	10.91	23	-31	-10	-83
CRWD	AI-native	fcf-positive	183.06	26	6	1	-90
DDOG	AI-native	fcf-positive	67.18	27	5	3	-83
ESTC	AI-native	fcf-positive	6.46	17	1	2	-12
GTLB	AI-native	fcf-positive	5.37	28	5	-2	-26
MDB	AI-native	fcf-positive	26.99	21	1	-0	-81

Ticker	Cohort	Group	Mkt cap ($bn)	Rev. CAGR (%)	SBC-FCF (%)	ROIC (%)	MoS (%)
NET	AI-native	pre-fcf	69.33	30	-8	-1	-94
PATH	AI-native	fcf-positive	6.05	12	6	13	-6
PLTR	AI-native	fcf-positive	325.82	46	37	38	-95
S	AI-native	pre-fcf	6.18	26	-25	-47	-54
SNOW	AI-native	pre-fcf	88.57	30	-9	-31	-84
SOUN	AI-native	pre-fcf	3.66	86	-106	-284	-95
ADBE	Exposed incumbent	fcf-positive	95.85	11	34	134	-8
ADSK	Exposed incumbent	fcf-positive	48.81	15	27	88	-40
APPF	Exposed incumbent	fcf-positive	5.39	23	16	56	4
ASAN	Exposed incumbent	pre-fcf	1.54	10	-12	-36	60
BILL	Exposed incumbent	fcf-positive	4.17	14	6	0	3
BOX	Exposed incumbent	fcf-positive	3.42	7	8	18	19
BRZE	Exposed incumbent	pre-fcf	2.63	26	-11	-38	-38
CRM	Exposed incumbent	fcf-positive	144.28	10	26	44	-6
DBX	Exposed incumbent	fcf-positive	6.74	0	27	21	2
DOCU	Exposed incumbent	fcf-positive	9.02	8	15	32	11
FRSH	Exposed incumbent	fcf-positive	2.54	18	10	14	65
HUBS	Exposed incumbent	fcf-positive	12.48	20	2	9	-21
INTU	Exposed incumbent	fcf-positive	105.02	16	27	59	-17
MNDY	Exposed incumbent	fcf-positive	3.69	29	10	-46	98
NICE	Exposed incumbent	fcf-positive	5.85	11	12	58	52
NOW	Exposed incumbent	fcf-positive	87.41	22	19	24	-10
PAYC	Exposed incumbent	fcf-positive	6.59	10	16	33	20
PCTY	Exposed incumbent	fcf-positive	5.94	14	16	53	46
PEGA	Exposed incumbent	fcf-positive	6.59	8	20	35	-21
TEAM	Exposed incumbent	pre-fcf	22.55	23	-6	-283	16
WDAY	Exposed incumbent	fcf-positive	31.65	15	14	14	8
ZM	Exposed incumbent	fcf-positive	30.98	4	25	32	-46
ACCS	Exposed other	pre-fcf	0.03	-16	-2	-84	37
ADEA	Exposed other	fcf-positive	3.22	8	27	26	-65
AEYE	Exposed other	pre-fcf	0.09	13	-3	-33	5
AGYS	Exposed other	fcf-positive	2.24	16	15	61	-40
AIXI	Exposed other	pre-fcf	0.01	-54	-77	770	-95
ALIT	Exposed other	fcf-positive	0.47	-17	11	-105	114
ALKT	Exposed other	pre-fcf	1.69	29	-9	-8	-40
ALRM	Exposed other	fcf-positive	2.32	8	13	15	22
AMOD	Exposed other	pre-fcf	0.01	nan	nan	265	-95
AMPL	Exposed other	pre-fcf	0.78	12	-20	-62	-17
AMST	Exposed other	pre-fcf	0.00	-10	-791	2818	-95
ANY	Exposed other	pre-fcf	0.01	-29	-224	-82	-95
API	Exposed other	pre-fcf	0.37	1	-15	2	-77
APP	Exposed other	fcf-positive	157.50	32	68	95	-83
APPS	Exposed other	pre-fcf	0.58	-1	-1	6	-69
ASUR	Exposed other	fcf-positive	0.26	10	1	-3	-6
AUUD	Exposed other	pre-fcf	0.01	nan	nan	3525	-95
AWRE	Exposed other	pre-fcf	0.02	-3	-38	-33	102
BIYA	Exposed other	pre-fcf	0.00	19	-82	-87	-89
BL	Exposed other	fcf-positive	1.83	9	6	6	8
BLKB	Exposed other	fcf-positive	1.79	1	14	60	15
BLND	Exposed other	pre-fcf	0.39	-9	-31	-6	-26
BMBL	Exposed other	fcf-positive	0.52	-5	25	-99	206
BMR	Exposed other	pre-fcf	0.03	3	-175	-179	-86
BNZI	Exposed other	pre-fcf	0.01	52	-171	229	-69
BSY	Exposed other	fcf-positive	10.21	11	27	35	-54
BTDR	Exposed other	pre-fcf	3.44	36	-278	-23	-95
BULL	Exposed other	pre-fcf	3.51	nan	nan	-2	-88
BZAI	Exposed other	pre-fcf	0.25	nan	-241	-991	-95
CALX	Exposed other	fcf-positive	2.72	1	2	6	-60
CCC	Exposed other	fcf-positive	3.08	11	10	11	-19
CDNS	Exposed other	fcf-positive	94.04	14	17	38	-76
CMRC	Exposed other	fcf-positive	0.29	5	3	-2	127
COMP	Exposed other	pre-fcf	6.52	27	-3	-19	-83
CREX	Exposed other	pre-fcf	0.04	17	-19	-27	49
CRNC	Exposed other	fcf-positive	0.48	-6	16	8	-57
CURR	Exposed other	fcf-positive	0.27	nan	9	25	-95
CVLT	Exposed other	fcf-positive	4.29	18	13	12	66
CWAN	Exposed other	fcf-positive	7.25	43	2	-0	-85
CXAI	Exposed other	pre-fcf	0.01	nan	-329	-309	-95
CXM	Exposed other	fcf-positive	1.26	8	5	18	110
CYN	Exposed other	pre-fcf	0.02	-53	-10007	-58	-95
DAVE	Exposed other	fcf-positive	3.35	46	48	53	-28
DBD	Exposed other	fcf-positive	2.66	nan	6	14	39
DCBO	Exposed other	pre-fcf	0.54	nan	nan	nan	nan
DFNS	Exposed other	pre-fcf	0.02	-47	-835	-111	-95
DHX	Exposed other	fcf-positive	0.10	-8	12	0	70
DJCO	Exposed other	fcf-positive	0.65	13	15	4	-41
DMRC	Exposed other	pre-fcf	0.23	-4	-69	-211	-85
DOMO	Exposed other	pre-fcf	0.15	-0	-18	-164	74
DSGX	Exposed other	fcf-positive	6.69	13	34	-603	-38
DSP	Exposed other	fcf-positive	0.69	24	4	-4	177
DT	Exposed other	fcf-positive	11.62	19	11	66	-20
DUOL	Exposed other	fcf-positive	4.85	38	24	59	42
DUOT	Exposed other	pre-fcf	0.26	71	-322	-27	-95
DV	Exposed other	fcf-positive	1.71	14	4	19	64
EGAN	Exposed other	fcf-positive	0.18	-2	12	261	82
EGHT	Exposed other	fcf-positive	0.31	0	3	5	74
ENRD	Exposed other	pre-fcf	0.28	nan	nan	4	-95
EXFY	Exposed other	pre-fcf	0.11	-3	-10	-24	28

Ticker	Cohort	Group	Mkt cap ($bn)	Rev. CAGR (%)	SBC-FCF (%)	ROIC (%)	MoS (%)
FICO	Exposed other	fcf-positive	23.43	18	31	33	-65
FRGT	Exposed other	pre-fcf	0.00	-12	-44	-80	-73
FTFT	Exposed other	pre-fcf	0.01	-64	-105	-3	-95
FUSE	Exposed other	pre-fcf	0.05	nan	-109	12	-95
GBTG	Exposed other	pre-fcf	4.90	12	-2	8	-15
GRAB	Exposed other	pre-fcf	15.05	20	-8	11	-53
GREE	Exposed other	pre-fcf	0.02	-7	-40	15	-95
GRND	Exposed other	fcf-positive	2.68	31	17	42	77
GTM	Exposed other	fcf-positive	1.78	1	21	5	68
GWRE	Exposed other	fcf-positive	11.33	21	10	9	-65
HIT	Exposed other	pre-fcf	0.08	29	-6	-28	-49
HKIT	Exposed other	pre-fcf	0.00	20	-23	1	nan
HTCR	Exposed other	pre-fcf	0.00	-27	-27	115	-50
IBTA	Exposed other	fcf-positive	0.86	3	7	-8	-4
IDAI	Exposed other	pre-fcf	0.01	-55	-1063	-362	-95
IDN	Exposed other	fcf-positive	0.15	10	14	347	-60
IFBD	Exposed other	pre-fcf	0.01	457	-6	-2268	-72
IMMR	Exposed other	pre-fcf	0.22	616	-0	2	500
INLX	Exposed other	pre-fcf	0.03	-2	-7	-89	-25
INTA	Exposed other	fcf-positive	1.86	17	3	28	-10
IPM	Exposed other	pre-fcf	0.02	43	-6	-398	479
JTAI	Exposed other	pre-fcf	0.01	-20	-146	38	nan
KARO	Exposed other	fcf-positive	1.47	15	36	41	-25
KC	Exposed other	fcf-positive	3.93	18	46	-4	-95
KLTR	Exposed other	pre-fcf	0.20	1	-0	-11	52
LAW	Exposed other	pre-fcf	0.24	7	-27	-37	30
LGCL	Exposed other	pre-fcf	0.01	-16	-5	4	367
LIF	Exposed other	fcf-positive	3.51	28	5	4	-70
LPSN	Exposed other	pre-fcf	0.03	-21	-17	-5	-42
LSPD	Exposed other	pre-fcf	1.17	17	-5	-88	87
LYFT	Exposed other	fcf-positive	5.45	19	13	-1	158
MANH	Exposed other	fcf-positive	8.42	8	25	426	-66
MAPS	Exposed other	fcf-positive	nan	-5	0	3	nan
MITK	Exposed other	fcf-positive	0.69	7	15	38	-22
MKTW	Exposed other	fcf-positive	0.26	-14	9	-5	268
MNDO	Exposed other	fcf-positive	0.02	-4	19	32	105
MRT	Exposed other	pre-fcf	0.17	nan	nan	nan	nan
MSTR	Exposed other	pre-fcf	65.47	-1	-29	-20	-95
MTC	Exposed other	pre-fcf	0.54	-4	-457	-94	-95
MTLS	Exposed other	fcf-positive	0.32	7	3	8	109
MYSE	Exposed other	pre-fcf	0.01	-9	-935651	-84	-95
MYSZ	Exposed other	pre-fcf	0.00	19	-64	-243	-49
NATL	Exposed other	fcf-positive	3.26	2	2	13	-61
NCNO	Exposed other	fcf-positive	1.66	12	6	8	-13
NRDY	Exposed other	pre-fcf	0.17	-3	-25	-348	23
NVNI	Exposed other	fcf-positive	0.01	8	4	17	-26
NXPL	Exposed other	pre-fcf	0.02	13	-15	-123	-30
NXTT	Exposed other	pre-fcf	0.11	97	-707	-32	nan
NXXT	Exposed other	pre-fcf	0.04	80	-69	-709	-87
OLB	Exposed other	pre-fcf	0.01	-45	-67	235	-95
ONTF	Exposed other	pre-fcf	0.34	-8	-18	-24	-16
ORKT	Exposed other	pre-fcf	0.01	-3	-87	-76	4
OTEX	Exposed other	fcf-positive	5.75	-4	14	13	51
PAR	Exposed other	pre-fcf	0.62	6	-14	-15	-90
PCOR	Exposed other	pre-fcf	7.99	18	-2	-56	-46
PD	Exposed other	fcf-positive	0.57	6	5	7	199
PDFS	Exposed other	pre-fcf	1.90	16	-19	7	-74
PERF	Exposed other	fcf-positive	0.17	13	17	25	nan
PHUN	Exposed other	pre-fcf	0.04	-27	-513	319	nan
PLUS	Exposed other	pre-fcf	2.32	6	-5	39	20
PRCH	Exposed other	fcf-positive	1.10	7	9	8	-89
PRGS	Exposed other	fcf-positive	1.08	17	19	16	462
PSQH	Exposed other	pre-fcf	0.04	100	-117	-90	-80
PTC	Exposed other	fcf-positive	17.05	15	23	78	-29
PUBM	Exposed other	fcf-positive	0.48	2	4	-11	-11
QH	Exposed other	pre-fcf	0.00	-17	-2	-44	-95
QTWO	Exposed other	fcf-positive	3.29	13	11	61	-59
QXO	Exposed other	fcf-positive	12.86	848	0	-7	-93
RCT	Exposed other	pre-fcf	0.04	56	-93	-498	-65
RDVT	Exposed other	fcf-positive	0.62	22	13	80	-44
RDZN	Exposed other	pre-fcf	0.11	nan	-43	143	-62
RMNI	Exposed other	fcf-positive	0.32	-1	8	103	72
RNG	Exposed other	fcf-positive	3.74	7	12	12	41
ROP	Exposed other	fcf-positive	36.15	13	28	28	-17
RSKD	Exposed other	pre-fcf	0.69	8	-3	-14	2
RSSS	Exposed other	fcf-positive	0.08	8	10	-28	61
RTB	Exposed other	pre-fcf	0.01	-55	-99	184	-90
RUM	Exposed other	pre-fcf	3.56	11	-96	-645	-95
RYDE	Exposed other	pre-fcf	0.15	20	-277	-145	-95
SAGT	Exposed other	pre-fcf	0.04	63	-18	15	-90
SAP	Exposed other	fcf-positive	201.77	8	18	49	-37
SEMR	Exposed other	pre-fcf	1.78	20	-2	-10	-24
SHOP	Exposed other	fcf-positive	139.67	28	13	15	-78
SMRT	Exposed other	pre-fcf	0.23	-18	-20	-49	-5
SMSI	Exposed other	pre-fcf	0.02	-32	-80	-899	-88
SMWB	Exposed other	pre-fcf	0.25	13	-4	233	127
SPPL	Exposed other	pre-fcf	0.01	12	-40	-32966	-93
SPT	Exposed other	pre-fcf	0.44	16	-6	-303	175

Ticker	Cohort	Group	Mkt cap ($bn)	Rev. CAGR (%)	SBC-FCF (%)	ROIC (%)	MoS (%)
SRAD	Exposed other	fcf-positive	3.68	20	26	-442	-69
SSNC	Exposed other	fcf-positive	16.69	7	19	17	-52
SSTI	Exposed other	pre-fcf	0.09	3	-6	-228	82
SURG	Exposed other	pre-fcf	0.01	-30	-37	-5268	-94
SVCO	Exposed other	pre-fcf	0.39	10	-84	-141	-91
SWVL	Exposed other	pre-fcf	0.02	3	-12	62	-79
TGL	Exposed other	pre-fcf	0.01	-69	-281	-266	-95
TONX	Exposed other	pre-fcf	0.20	1098	-262	-77	-95
TRAK	Exposed other	fcf-positive	0.17	8	30	574	-4
TROO	Exposed other	fcf-positive	0.50	119	4	-66	-95
TTAN	Exposed other	pre-fcf	7.38	25	-11	-64	-79
TTD	Exposed other	fcf-positive	11.04	21	12	24	-5
TWAV	Exposed other	pre-fcf	0.01	-17	-128	-13	-93
TYL	Exposed other	fcf-positive	15.01	9	22	310	-59
U	Exposed other	fcf-positive	11.67	-6	5	-38	-94
UBER	Exposed other	fcf-positive	161.16	18	15	13	-36
UPBD	Exposed other	fcf-positive	1.12	8	5	6	449
UPLD	Exposed other	fcf-positive	0.02	-16	7	3	500
VERX	Exposed other	fcf-positive	2.37	14	2	-5	-23
VS	Exposed other	pre-fcf	0.01	143	-68	-147	-72
VTEX	Exposed other	fcf-positive	0.69	9	9	7	-15
WALD	Exposed other	pre-fcf	0.20	12	-10	-111	136
WBTN	Exposed other	pre-fcf	1.80	3	-2	-148	-59
WCT	Exposed other	pre-fcf	0.01	-26	-549	-344	-95
WFCF	Exposed other	fcf-positive	0.05	-0	4	51	-23
WK	Exposed other	fcf-positive	3.11	19	5	3	-7
WRD	Exposed other	pre-fcf	2.62	30	nan	-93	-95
XPER	Exposed other	pre-fcf	0.37	-6	-12	-6	57
XTIA	Exposed other	pre-fcf	0.07	189	-98	487	-94
YALA	Exposed other	pre-fcf	0.97	2	nan	39	15
YMM	Exposed other	pre-fcf	9.17	20	nan	16	-95
YOU	Exposed other	fcf-positive	8.14	21	42	51	-43
YXT	Exposed other	pre-fcf	0.02	-10	nan	-163	19
ZETA	Exposed other	pre-fcf	4.64	35	-1	-11	-43
ACIW	Infrastructure	fcf-positive	4.74	10	12	21	-31
AFRM	Infrastructure	fcf-positive	22.56	38	12	2	-87
AIFC	Infrastructure	pre-fcf	0.13	nan	-99	-68	-95
AISP	Infrastructure	pre-fcf	0.09	13	-43	-4935	-91
AKAM	Infrastructure	fcf-positive	21.47	5	7	5	-76
ALAR	Infrastructure	pre-fcf	0.07	27	nan	4	-92
ALLT	Infrastructure	fcf-positive	0.41	6	18	6	-78
AMBR	Infrastructure	pre-fcf	0.15	nan	nan	nan	nan
APPN	Infrastructure	fcf-positive	1.74	16	3	7	-23
ARBE	Infrastructure	pre-fcf	0.15	-1	-3355	-76	-95
ARQQ	Infrastructure	pre-fcf	0.27	24	-3870	1664	-95
ATCH	Infrastructure	pre-fcf	0.04	12	-35	-56	nan
ATEN	Infrastructure	fcf-positive	1.96	8	11	10	-69
AUID	Infrastructure	pre-fcf	0.02	199	-769	12492	-95
AVPT	Infrastructure	fcf-positive	2.19	24	15	-1045	-48
BAND	Infrastructure	fcf-positive	1.18	13	2	-1	-78
BB	Infrastructure	fcf-positive	2.25	-17	4	16	-72
BKKT	Infrastructure	pre-fcf	0.44	80	-5	-412	66
BLIN	Infrastructure	pre-fcf	0.01	-0	-10	74	11
BLZE	Infrastructure	pre-fcf	0.28	19	-11	-15	-76
BNAI	Infrastructure	pre-fcf	0.15	185	-1841	145	-95
CCSI	Infrastructure	fcf-positive	0.63	-1	26	22	-20
CETX	Infrastructure	pre-fcf	0.01	4	-11	-43	-11
CGNT	Infrastructure	fcf-positive	0.68	13	1	5	7
CHKP	Infrastructure	fcf-positive	11.88	6	35	34	18
CINT	Infrastructure	fcf-positive	0.52	13	11	nan	63
CISO	Infrastructure	pre-fcf	0.01	-30	-35	108	-92
CLBT	Infrastructure	fcf-positive	3.23	21	22	25	-73
CORZ	Infrastructure	pre-fcf	7.83	-14	-156	-62	-95
CPAY	Infrastructure	fcf-positive	22.48	11	25	56	11
CSAI	Infrastructure	pre-fcf	0.01	189	-189	-2904	-95
CSGS	Infrastructure	fcf-positive	2.29	3	7	15	-18
CYAB	Infrastructure	pre-fcf	0.01	nan	-250	-623	-85
CYBR	Infrastructure	fcf-positive	20.59	35	2	-7	-89
DAVA	Infrastructure	fcf-positive	0.17	-2	5	-90	330
DFDV	Infrastructure	pre-fcf	0.13	136	-213	-40	-95
DOCN	Infrastructure	pre-fcf	15.94	15	-5	9	-95
DOX	Infrastructure	fcf-positive	6.59	-3	12	30	-4
DTSS	Infrastructure	pre-fcf	0.01	55	-3	-194	17
EEFT	Infrastructure	fcf-positive	2.73	8	6	28	181
EVCM	Infrastructure	fcf-positive	2.09	-6	7	12	-70
EXOD	Infrastructure	pre-fcf	0.23	34	-38	-18	81
FFIV	Infrastructure	fcf-positive	19.18	6	22	73	-40
FIVN	Infrastructure	fcf-positive	1.32	12	3	4	55
GCT	Infrastructure	fcf-positive	1.58	35	11	18	45
GDDY	Infrastructure	fcf-positive	11.49	8	26	38	41
GEN	Infrastructure	fcf-positive	14.97	13	26	26	32
GGRP	Infrastructure	pre-fcf	0.01	-21	-67	-5128	-63
GOAI	Infrastructure	pre-fcf	0.13	68	-56	-20	-81
GRRR	Infrastructure	pre-fcf	0.43	27	-34	-44	-93
IIIV	Infrastructure	fcf-positive	0.59	-22	7	25	-64
IMXI	Infrastructure	pre-fcf	0.48	-5	-12	16	19
INTZ	Infrastructure	pre-fcf	0.02	5	-159	-129	-93
IOT	Infrastructure	pre-fcf	19.59	31	-5	5	-74

Ticker	Cohort	Group	Mkt cap ($bn)	Rev. CAGR (%)	SBC-FCF (%)	ROIC (%)	MoS (%)
JG	Infrastructure	pre-fcf	0.03	13	nan	-8	-66
KPLT	Infrastructure	pre-fcf	0.03	14	-3	62	296
KSPI	Infrastructure	fcf-positive	16.83	43	16	27	182
KVYO	Infrastructure	fcf-positive	6.98	32	3	-1	-49
LIDR	Infrastructure	pre-fcf	0.10	-53	-12364	-114	-95
LSAK	Infrastructure	fcf-positive	0.42	12	1	-12	-39
MLGO	Infrastructure	fcf-positive	0.05	-15	4	19	nan
MSAI	Infrastructure	pre-fcf	0.01	5	-146	-130	nan
MSFT	Infrastructure	fcf-positive	3153.07	16	19	39	-60
NN	Infrastructure	pre-fcf	3.03	2	-1639	-32	-95
NTCT	Infrastructure	fcf-positive	2.78	1	26	253	18
NTNX	Infrastructure	fcf-positive	14.08	15	15	12	-23
ODD	Infrastructure	pre-fcf	0.55	18	-10	7	226
ORCL	Infrastructure	pre-fcf	469.14	10	-46	11	-92
ORIO	Infrastructure	pre-fcf	0.02	2	-10	3	150
OS	Infrastructure	pre-fcf	5.80	27	-3	70	-77
OSPN	Infrastructure	fcf-positive	0.43	2	15	56	24
PAYO	Infrastructure	fcf-positive	1.69	12	6	2	-95
PAYS	Infrastructure	fcf-positive	0.33	34	71	-13	65
PDYN	Infrastructure	pre-fcf	0.29	6	-501	-49	-95
PGY	Infrastructure	fcf-positive	1.28	25	14	4	-39
PRTH	Infrastructure	fcf-positive	0.49	12	8	12	31
PSFE	Infrastructure	fcf-positive	0.48	4	6	2	279
RAMP	Infrastructure	fcf-positive	2.27	12	10	81	25
RDWR	Infrastructure	fcf-positive	1.21	8	4	4	-27
REKR	Infrastructure	pre-fcf	0.12	17	-42	-67	-92
RPD	Infrastructure	fcf-positive	0.45	5	5	4	209
RXT	Infrastructure	fcf-positive	1.37	-4	1	-1	-95
RZLV	Infrastructure	pre-fcf	0.85	nan	-143	740	-95
SAIL	Infrastructure	pre-fcf	8.34	23	-3	-303	-74
SANG	Infrastructure	fcf-positive	0.14	-7	8	-14	72
SKYA	Infrastructure	pre-fcf	0.08	nan	-1169	-172	-95
SNPS	Infrastructure	fcf-positive	100.70	17	19	10	-90
SOS	Infrastructure	pre-fcf	0.02	29	-164	-38	-95
SPSC	Infrastructure	fcf-positive	2.06	17	14	114	83
STEM	Infrastructure	pre-fcf	0.08	-39	-15	40	-95
TAOP	Infrastructure	pre-fcf	0.01	-11	-17	-34	-89
TCX	Infrastructure	fcf-positive	0.18	7	1	-2	-95
TDC	Infrastructure	fcf-positive	2.85	-4	33	107	37
TLS	Infrastructure	pre-fcf	0.32	10	-1	-96	-31
TUYA	Infrastructure	pre-fcf	1.51	nan	nan	nan	nan
UBXG	Infrastructure	pre-fcf	0.05	-40	-33	-91	-95
USIO	Infrastructure	pre-fcf	0.04	3	-3	-10	90
VERI	Infrastructure	pre-fcf	0.19	-14	-64	-168	-94
VHC	Infrastructure	pre-fcf	0.06	305	-11322	-117	-95
VRNS	Infrastructure	pre-fcf	3.13	13	-2	-12	-49
VRSN	Infrastructure	fcf-positive	25.20	5	58	80	-78
WIX	Infrastructure	fcf-positive	3.12	13	13	-6	120
XBP	Infrastructure	pre-fcf	0.02	nan	nan	nan	nan
XYZ	Infrastructure	fcf-positive	41.75	5	8	8	-55
YEXT	Infrastructure	fcf-positive	0.42	4	2	34	213
ZENA	Infrastructure	pre-fcf	0.09	191	-375	-66	-95
ZPTA	Infrastructure	pre-fcf	nan	-100	nan	nan	nan
ZSQR	Infrastructure	pre-fcf	0.48	nan	-709	-139	-95
FROG	Neutral infra	pre-fcf	8.54	24	-1	-12	-86
FSLY	Neutral infra	pre-fcf	4.94	12	-13	-11	-88
FTNT	Neutral infra	fcf-positive	83.57	14	30	75	-62
OKTA	Neutral infra	fcf-positive	21.43	13	12	22	-58
PANW	Neutral infra	fcf-positive	228.55	16	20	19	-93
QLYS	Neutral infra	fcf-positive	3.24	10	31	61	22
RBRK	Neutral infra	pre-fcf	15.11	44	-4	-13	-89
TENB	Neutral infra	fcf-positive	2.32	12	7	8	46
TWLO	Neutral infra	fcf-positive	27.83	11	6	3	-73
ZS	Neutral infra	fcf-positive	29.85	26	3	-1	-71

B Statistical and Valuation Methodology Details

This appendix gives the explicit mechanics behind each model so the analysis can be reimplemented independently. Notation: r_f risk-free rate, ERP equity risk premium, τ tax rate, β_i CAPM beta, GM_i gross margin, g growth, W WACC, T explicit forecast horizon.

B.1 Cost of capital (CAPM / WACC)

The cost of equity is $r_i^e = r_f + \beta_i \,\mathrm{ERP}$, with β_i from 5 years of monthly returns regressed on the market factor. The after-tax cost of debt is $r_i^d(1-\tau)$, and the firm WACC weights the two by market equity and net debt (restating Equation (3)):

$$W_i = w_i^e\,(r_f + \beta_i\,\mathrm{ERP}) + w_i^d\,r_i^d(1-\tau), \qquad w_i^e + w_i^d = 1.$$

We use $r_f = 4.30\%$, ERP $= 5.0\%$ and $\tau = 21\%$. Median estimated W is 9.5%. Implementation details that affect replication are explicit: betas are clipped to 0.30–2.50, firms with fewer than 24 monthly observations use the cross-sectional median beta, the cost of debt is bounded between r_f and 15%, and WACC is clipped to 7–18%. These bounds prevent small-history or distressed observations from dominating fair value.

B.2 Owner-earnings DCF (Route A)

For FCF-positive firms, enterprise value is a two-stage discounted cash flow on SBC-adjusted free cash flow F_0 over a ten-year explicit horizon ($T = 10$) with terminal rate $g_\infty = 3\% < W$. Stage-1 growth does not step down discretely; it glides linearly from g_1 to g_∞ across the horizon (an H-model-style fade), which is deliberately conservative for high-growth software:

$$\mathrm{EV}_i = \sum_{t=1}^{T} \frac{F_0 \prod_{s=1}^{t}(1+g_s)}{(1+W_i)^t} + \frac{F_0 \prod_{s=1}^{T}(1+g_s)\,(1+g_\infty)}{(W_i - g_\infty)\,(1+W_i)^T}, \qquad g_t = g_1 + (g_\infty - g_1)\,\frac{t-1}{T-1}. \quad (7)$$

Equity value is EV_i plus net cash; per-share fair value divides by projected diluted shares (current shares inflated by capped annual dilution over the horizon). We also compute a *normalized* variant that first converts gross profit to a steady-state owner-earnings margin before applying Equation (7), which stabilises firms whose current FCF is depressed by reinvestment.

B.3 Path-to-profitability DCF (Route B)

For pre-FCF firms, the FCF margin glides linearly from its current (negative) level m_0 to a target $m_i^\star = k \cdot \mathrm{GM}_i$ over $G = 7$ years, where $k = 0.19$ is the data-driven steady-state conversion (median target margin 10.6%):

$$m_{i,t} = m_0 + \frac{\min(t, G)}{G}\,(m_i^\star - m_0), \qquad F_{i,t} = \mathrm{Rev}_{i,t}\, m_{i,t}, \quad (8)$$

so the margin reaches $m_i^\star$ at year G and holds thereafter; revenue compounds at a rate that fades from realised CAGR toward g_∞ as in Route A, with the whole path discounted at the firm's CAPM WACC plus a two-point execution-risk premium (capped at 20%; a 12% flat fallback where beta is unavailable). A survival/runway gate and a dilution drag are applied, and a contamination guard removes firms whose implied current FCF yield is implausibly high (which would otherwise route them incorrectly to Route A). Revenue growth is clipped to the -5–30% range in the valuation engine, the target margin is clipped to 5–40%, annual dilution is capped at 25% for the per-share bridge, and displayed MoS is clipped to -95–500%. The contamination guard drops the actual-FCF lens when EV/FCF is below 5x or the latest SBC-adjusted FCF margin exceeds 50%, because those values are usually payment-float, one-off-cash-flow or data-contamination artifacts rather than durable SaaS owner earnings.

B.4 Fair value and margin of safety

Per firm we triangulate the point estimate and a bear–base–bull band as the median of the DCF, the normalized DCF and a peer EV/Sales multiple conditioned on the Rule-of-40, and define the margin of safety as in Equation (4): $\mathrm{MoS}_i = \mathrm{FairValue}_i/\mathrm{Price}_i - 1$.

B.5 Reverse DCF (implied growth)

For the lead FCF-positive reverse-DCF tables, the market-implied growth is the single stage-1 rate that sets enterprise value equal to the current market EV using latest SBC-adjusted FCF as

F_0. Following the standard reverse-DCF convention we hold this rate *constant* over the explicit horizon, rather than gliding it as in the forward fair value of Equation (7), so each firm's implied growth is one interpretable number, and bisection-solve

$$\hat{g}_i^{\text{impl}} : \ \text{EV}_i(\hat{g}_i^{\text{impl}}) = \text{EV}_i^{\text{market}}, \qquad \text{EV}_i(g) = \sum_{t=1}^{T} \frac{F_0(1+g)^t}{(1+W_i)^t} + \frac{F_0(1+g)^T(1+g_\infty)}{(W_i - g_\infty)(1+W_i)^T}, \tag{9}$$

holding W_i, g_∞ and T fixed over the search bracket $g \in [-50\%, 150\%]$. We compare $\hat{g}_i^{\text{impl}}$ to realised revenue CAGR; "realised − implied" > 0 means the market pays for less growth than the firm delivered.

For pre-FCF firms the reported normalised reverse-DCF diagnostic is separate: F_0 is normalised owner earnings, defined as latest revenue times the routed target FCF margin, and the solve uses the same linearly fading cash-flow path as Equation (7). This is why the pre-FCF diagnostic is reported as normalised implied growth rather than mixed into the lead actual-FCF reverse-DCF table.

B.6 R&D capitalization and productivity

R&D capital follows a perpetual-inventory recursion with depreciation $\delta = 20\%$:

$$K_{i,t} = (1-\delta)\, K_{i,t-1} + R_{i,t}, \tag{10}$$

following Peters and Taylor [10]. R&D productivity is gross profit created per cumulative R&D dollar (Equation (5)). The Rule-of-40 is revenue growth plus SBC-adjusted FCF margin.

B.7 Event study: difference-in-differences

For each outcome y (gross margin, SBC-adjusted FCF margin, ROIC) we form the firm-level change $\Delta y_i = y_{i,\text{latest}} - y_{i,\text{base}}$. Treatment is the top tercile (top 33%) of mean 2023–2024 AI-call salience. The difference-in-differences estimate is

$$\widehat{\text{DiD}}_y = \overline{\Delta y}\big|_{\text{treated}} - \overline{\Delta y}\big|_{\text{control}}, \tag{11}$$

computed on deltas winsorized at the 2.5/97.5 percentiles so the mean is not outlier-driven [29]. Significance uses the Welch (unequal-variance) t:

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s_1^2/n_1 + s_2^2/n_2}}. \tag{12}$$

B.8 Quartile sorts

For the R&D-productivity analysis we sort FCF-positive firms with finite productivity into quartiles and report group medians of returns on capital, cash margin, growth and margin of safety. Quartile sorts are descriptive: they show whether quality and cheapness move together, not a tradeable factor.

C Data Sources and Quality

C.1 Primary fundamentals and prices: Sharadar

Firm fundamentals are from the Sharadar Core US Fundamentals feed (point-in-time income statement, balance sheet and cash-flow items); prices and split/dividend-adjusted returns are from Sharadar Equity Prices. Coverage is comprehensive for listed US software firms over the study window. As an accessible commercial tier this is not the CRSP/Compustat academic standard: delisting treatment and deep historical corporate-action adjustments are less complete. We mitigate by using point-in-time fields, requiring both endpoints, converting foreign-currency line items to USD before forming ratios, and pinning everything to a hashed snapshot.

C.2 Factors and risk-free rate: Kenneth French

The market factor and the risk-free rate (used for CAPM betas and WACC) are from the Kenneth French data library, the standard academic source, aligned to the firms' monthly return windows. The same library supplies the five factors plus momentum used in the spanning test of Section 8.7; the monthly series is frozen to `data/ff_factors_monthly.csv` and its coverage ends in late 2025, which sets the 31-month overlap of the factor regression.

C.3 Return panel for the spanning test: public adjusted closes

The valuation and cross-sectional results use Sharadar prices. The monthly *return panel* used only for the factor-spanning test (Section 8.7) is built from public split- and dividend-adjusted monthly closes (Yahoo Finance chart API) and frozen to `data/returns_monthly.csv`, because the Sharadar monthly price panel was not part of the released snapshot. We disclose this mixed-source choice plainly. It carries one material limitation: public histories resolve for 386 of the 473 signal-panel names (81.6% coverage), and firms acquired or taken private during the window typically have no continuing series, so the return test inherits a survivorship tilt that the cross-sectional valuation results do not. This is the main reason the factor test should be read as a feasibility check rather than a definitive anomaly test.

C.4 Qualitative fields: SEC EDGAR

First-principles overlay fields are extracted from as-filed annual reports on SEC EDGAR (`10-K`, `20-F`, `40-F`); the extraction rules are in Appendix D. We populate fields only from disclosed language and never infer missing values.

C.5 AI-text panel

The AI-salience and stance panel is built from earnings-call and filing language following Hassan et al. [4] and Loughran and McDonald [5].

C.6 Snapshot integrity

All results are pinned to a frozen snapshot so every number is traceable, versionable and transparent:

- **Snapshot id:** `8a7ec05c7873` **Built:** 17 June 2026
- **Data tier:** Tier-1 (Sharadar + SEC EDGAR + Ken French)
- **Study window:** Sep 2023–Mar 2026 (median span 2.25 years)

D EDGAR Extractor Methodology

The first-principles overlay reads as-filed annual reports rather than inferring fields, so each extractor is a documented, auditable rule. Coverage across the 32 portfolio candidates is reported in Table 18; the rules are:

Filing retrieval
Resolve each ticker's CIK and fetch the latest annual report (`10-K/20-F/40-F`) via the EDGAR full-text and submissions endpoints; store the source URL. Coverage 32/32.

AI-risk language
Flag explicit AI/automation risk-factor language and store the raw sentence verbatim. Because such language is now near-universal (32/32), its mere presence is not a differentiator; the useful signal is the specific sentence and stance.

Net revenue retention (NRR)
Capture disclosed dollar-/net-retention figures with the disclosed operator ($=$, $>$ lower-bound, $\sim$ approximate). We do not estimate NRR when undisclosed. Coverage 6/32.

Customer concentration
Capture disclosed top-customer or top-10-customer revenue concentration. Coverage 5/32.

Contract term
Capture strict contract-term language (e.g. minimum committed term) only when explicitly stated. Coverage 0/32.

Revolver availability
Capture disclosed undrawn revolving-credit availability as a liquidity backstop. Coverage 1/32.

The teachable limitation is that filings disclose far less than an underwriter needs: retention and concentration, the fields that would most directly confirm or refute the moat thesis, are disclosed by a minority of firms. The overlay therefore strengthens conviction where firms disclose and stays silent where they do not.

E Alternative Specifications

This appendix states the alternative modelling choices behind the robustness discussion (Section 10) and the expected direction of effect for each.

Discount rate
Per-firm CAPM WACC (baseline) versus a flat 10% and a +200 bps stress. A higher rate lowers fair value and shrinks every margin of safety roughly in parallel, so it moves the *level* of cheapness but not the cross-sectional *ranking* that defines the localized gap.

Fair-value weighting
Median triangulation (baseline) versus DCF-only and multiple-only. DCF-only widens dispersion (more sensitive to growth assumptions); multiple-only compresses it toward sector norms. The exposed-incumbent cohort remains the cheapest under all three.

Cohort definition
Strict incumbents only versus the broad exposed cohort. Tightening the definition raises the incumbent median MoS (fewer marginal names); loosening it dilutes the signal but does not reverse its sign.

SBC treatment
SBC-adjusted FCF (baseline) versus as-reported FCF with separate dilution tracking. As-

reported FCF flatters cash-burning firms most; the baseline is deliberately conservative and is where many pre-FCF valuations break.

Liquidity screens

Floors at \$50M, \$500M and \$1B. Lower floors admit model-sensitive small-cap tails with extreme margins of safety; the investable conclusions use the higher floors precisely to exclude them.

Time window

Sensitivity across baseline choices within the 2023–2026 span. Earlier baselines deepen the measured price drawdown; the two-worlds split and the incumbent gap persist.

F Glossary of Terms

ARR	Annual recurring revenue; the annualised value of subscription contracts.
Cohort	An economic grouping by relationship to the AI narrative (AI-native, exposed incumbent, exposed other, infrastructure/neutral), carrying a signed exposure.
DCF	Discounted cash flow; the present value of projected future cash flows.
Diff-in-diff (DiD)	Difference between the average outcome change of treated and control firms.
EV	Enterprise value; market equity plus net debt.
FCF (SBC-adjusted)	Free cash flow with stock-based compensation treated as a cash-equivalent cost; the basis for valuation in this paper.
Margin of safety (MoS)	Fair value divided by price, minus one; positive means undervalued.
Moat	A durable competitive advantage (switching costs, data, workflow embeddedness).
NRR	Net revenue retention; revenue from existing customers period-over-period, net of churn.
Owner earnings	Cash an owner could withdraw without impairing the business; here SBC-adjusted FCF.
Pre-FCF	A firm with non-positive SBC-adjusted free cash flow today (Equation (1)).
Point-in-time	Using data as it was reported and public at each historical date.
Reverse DCF	Solving a DCF for the growth rate implied by the current price (Eq. (9)).
ROIC	Return on invested capital; after-tax operating profit over invested capital.
Rule-of-40	Revenue growth plus FCF margin; a SaaS health heuristic.
SBC	Stock-based compensation; a non-cash but economically real expense and source of dilution.
Signed exposure	AI exposure with a direction (moat widener +, eroder −, neutral 0).
Snapshot	The frozen set of source artifacts (hash `8a7ec05c7873`) all numbers resolve from.
WACC	Weighted average cost of capital (Equation (3)).
Welch t	An unequal-variance t-test (Equation (12)).
Winsorization	Clipping extreme values to chosen percentiles to limit outlier influence.