
Two Clocks, One Stock: Measuring When Software Valuations Diverge from Operational Reality

By the end of this paper, you will have a systematic way to answer one question about any software stock: is the current pricing consistent with what the operational record actually shows — or has narrative moved ahead of the data?

UNIVERSE	PERIOD	PUBLISHED
IGV constituents (point-in-time)	2020 – Q1 2026	April 2026

SECTION

Executive summary

01 The observation — and the question

02 The V and Q scores: what the map measures

03 The V×Q map: four regimes

04 From map to evidence: testing whether the map measures something real

05 Two readings of the map: Signal A and Signal B

06 Validation: what the historical record shows

07 Using the map: implications for operators and investors

EXECUTIVE SUMMARY

The Question	Software valuations move faster and further than the businesses beneath them. When a company's stock re-rates 30% in a quarter while its revenue, margins, and cash flow change by single digits, is there any systematic pattern to what follows? This paper investigates that empirical question across 100+ IGV software names over 2020–Q1 2026. It does not predict regime change or business model disruption. It reports what the historical record shows when specific combinations of market pricing and operational metrics have co-occurred.
The Approach	<p>We define two scores, both grounded in statistical observations. The Value score (V) measures how cheap a company's trailing PE sits in its own three-year history — a rising V score means multiple compression, the market's current belief about the company is more pessimistic than its own recent history. The Quality composite (Q) measures where three operational metrics (sales efficiency, revenue growth, FCF margin) sit in their own three-year quarterly history.</p> <p>Together, V and Q place every IGV constituent on a two-dimensional map with four regime states, updated weekly.</p>
The Result	<p>To test whether the map measures something real, we implement a systematic backtest varying V thresholds, Q composite definitions, and holding periods. Two distinct V×Q configurations emerge where the pattern is clearest — we call each a Signal.</p> <p>Signal A fires when all three operational metrics are simultaneously at historical highs (Q>87) and the stock is not at a historically elevated valuation: 91% historical win rate vs IGV benchmark, +47% annualised per-trade median excess return.</p> <p>Signal B fires when the rate of change of revenue growth and operating margin is accelerating above historical norms (delta Q>75) with the same valuation guard rail: 67% win rate, +19% annualised.</p> <p>Both are statistically significant; Signal A at the 1% level, Signal B at the 5% level.</p>
The Limitation	<p>The V×Q system is a diagnostic instrument. It tells you where each company currently sits relative to its own history on two dimensions. It cannot tell you whether the current narrative about AI and software business models will ultimately prove correct, whether a company's operational trajectory will inflect in the next quarter, or whether any specific map position guarantees a future outcome. The historical pattern is evidence that the map readings have corresponded to observable outcomes in the past — it is not a forecast.</p>

The observation — and the question

Software valuations fluctuate more than the underlying businesses do. In any given quarter, a company's stock might re-rate by 30% while its revenue growth, margins, and cash flow change by single digits. Everyone who covers software knows this. It is visible in the AI re-rating of 2023, in the rate-shock correction of 2022, in the COVID crash and recovery of 2020, and in the 2026 sell-off driven by fears that AI will render traditional SaaS models obsolete. The amplitude of narrative-driven price moves is a defining feature of the sector, not a bug within it.

What is less obvious is whether those divergences between pricing and operations follow any pattern. This paper investigates that question empirically across 100+ IGV software names over 2020–Q1 2026. It does not resolve whether any given narrative will ultimately prove correct. What it does is narrower: build a systematic measurement tool for the divergence between market pricing and operational reality, apply it across the full software universe, and report what the historical record shows.

The question is not whether the current narrative is right. The question is whether the gap between a stock's pricing history and its operational history carries any systematic information.

Analytical approach

The intellectual lineage of this approach is factor investing — the insight that systematic differences in stock returns can be explained by measurable company characteristics rather than qualitative judgement alone. We apply this logic at the sector level, with an important caveat: Fama-French operates across thousands of stocks and decades of data; this framework covers 100+ software names over six years. The statistical power is correspondingly limited. The value is specificity — these factors are calibrated to software business dynamics in a way that broad cross-asset models are not. Unlike cross-sectional screens, analyst targets, or earnings transcripts, the V×Q map is two-dimensional, self-referential, and updated weekly, measuring each company against its own history rather than against peers or a qualitative view.

The approach is entirely statistical and backward-looking. We define two quantifiable factors, score every company against them weekly, and examine what the historical record shows when each factor reaches specific thresholds. No qualitative views are formed on individual companies. Every entry in the backtest is triggered by a quantitative threshold crossing, and every exit is mechanical.

The V and Q scores: what the map measures

The Value score (V): a valuation guard rail

The V score measures how cheap a company’s current trailing PE sits within its own three-year history. It is a self-referential percentile: what fraction of weekly observations over the past 36 months had a higher PE than today’s reading? A high V score means the stock is cheap relative to its own recent history. A V score of 80 means the stock is below 80% of its own weekly history over the past three years. A low V score means the stock is at or above its own recent pricing norm.

V does not compare Salesforce to ServiceNow — it compares Salesforce today to Salesforce over the last three years. A company that has always traded at high multiples will show a high V score only when the market has moved away from those expectations specifically. Crucially, the backtest data shows win rates are positive across the entire V range from V=50 (at-median valuation) upward. V functions as a valuation guard rail that excludes names where the market is already pricing in exceptional outcomes — not a deep-value filter requiring an extreme discount. V high does not mean “cheap.” It means “not at a historically elevated multiple.”

The Quality composite (Q): two readings of operational strength

The Q composite measures where a company’s current operational metrics sit within their own three-year quarterly history — a percentile score on a 0–100 scale computed from quarterly earnings filings using a 12-quarter lookback. The composite is the arithmetic average of three individual metric percentiles. The two signals our backtest identified use different Q composites: Signal A uses absolute levels (where are the metrics today vs their own history?); Signal B uses rates of change (how fast are they improving?).

Metric	Signal A — absolute level (SE + RY + FCF)	Signal B — rate of change (SE + δ RY + δ OM)
Sales efficiency	Revenue growth per dollar of S&M spend. Above-historical efficiency means the growth engine is working, not just spending.	Same absolute metric as Signal A.
Revenue / margin metric	Revenue growth YoY (RY): quarterly revenue vs same quarter prior year. The most direct test of demand health.	Delta revenue growth (δ RY): this quarter’s YoY growth rate minus last quarter’s. Measures whether revenue acceleration is above the company’s own historical rate of change.
Quality / cash metric	FCF margin: TTM free cash flow as % of revenue. Confirms pricing power and cost discipline are intact.	Delta operating margin (δ OM): this quarter’s operating margin minus last quarter’s. Captures earnings quality of the acceleration — revenue growing while margins improve.

Signal A metrics scored as percentile vs trailing 12 quarterly observations. Signal B δ metrics measure quarter-on-quarter rate of change, scored against the company’s own 3yr history of such changes. All composites are arithmetic averages. Detailed definitions in methodology note.

Both Q composites are entirely backward-looking. High Q means the recent past was strong — it does not mean the future is safe. A Q>87 composite means the arithmetic average of the three

percentiles exceeds 87; it requires broad-based strength across all three dimensions, not a perfect score on any single one. In structural disruption, Q can remain historically high through several quarters before inflection appears in filed metrics. The map cannot distinguish “narrative ahead of reality” from “narrative ahead of reality but correct with a lag.”

The table below shows how the Q composite works in practice, using a sample of IGV names as of the most recent quarter. Each metric is expressed as a percentile against the company's own trailing 12 quarterly observations — so 95.8 means that metric is stronger today than in 95.8% of its own recent quarters, while 4.5 means it is near a three-year low. The composite is their arithmetic average. Notice the difference between Salesforce (95.8 across all three — broad-based historical strength) and Microsoft (43.9 composite, dragged down by an FCF margin near a three-year low despite reasonable scores elsewhere). Same sector, same week, very different operational readings.

Company	FCF Margin (Pctl)	Sales Efficiency (Pctl)	Rev Growth (Pctl)	Q Composite
Adobe	95.8	62.5	87.5	81.9
Intuit	95.8	70.8	70.8	79.1
Microsoft	4.5	68.2	59.1	43.9
Salesforce	95.8	95.8	95.8	95.8
Atlassian	4.5	86.4	77.3	56.1
Workday	95.8	62.5	29.2	62.5

The V×Q map: four regimes

V and Q together define a two-dimensional space. Every software company, at any point in time, occupies a position in that space. The map has four regimes, each with a different interpretation and a different historical pattern.

Exhibit 1: The V×Q regime map

	Q LOW — operations below historical norms	Q HIGH — operations at historical highs
V LOW Multiple at or above own norm	<p>NARRATIVE EXPANSION</p> <p>Multiples elevated relative to own history; business hasn't caught up with priced-in expectations. Risk of disappointment is highest here.</p> <p><i>Typical: 2021 AI bull, 2023–2024 AI re-rating</i></p>	<p>GROWTH REGIME</p> <p>Market and business aligned. Strong operations, fairly to richly priced. No asymmetric entry opportunity. Monitor for V movement as leading indicator of regime shift.</p> <p><i>Typical: well-functioning growth market</i></p>
V HIGH Multiple compressed vs own norm	<p>DETERIORATION</p> <p>Market pessimism accompanied by genuine operational weakness in trailing metrics. Cheap for a reason. Signal A correctly suppressed — Q has fallen below threshold.</p> <p><i>Typical: genuine business slowdown or model disruption</i></p>	<p>★ DISLOCATION</p> <p>Multiple compressed while trailing operations remain strong. The operational record has not confirmed the narrative. Signal A fires when $Q > 87$; Signal B may also fire if $\delta Q > 75$. Historically resolved toward positive relative returns.</p>

From map to evidence

The $V \times Q$ map produces a visual framework for placing software companies in a two-dimensional space. But a framework that generates plausible-sounding quadrant labels is not the same as a framework that measures something real. The question is whether the map readings actually correspond to any systematic pattern in what happens to stock prices in the weeks that follow.

To answer that rigorously, we apply the map in a systematic backtest across the full IGV universe. The outcome variable is defined in advance and measured mechanically: 13-week relative return versus the IGV ETF benchmark from the date a map reading is triggered. We vary V and Q thresholds across a grid of combinations to find which specific map positions have the strongest and most consistent historical relationship with the outcome. This is not vibe-scoring. The map either corresponds to a measurable pattern in subsequent returns, or it does not. The backtest determines which. A signal, in this context, is simply a label for a specific set of map coordinates — a combination of V threshold and Q threshold — where the backtest evidence is clearest. It is not a recommendation to trade.

Two readings of the map: Signal A and Signal B

The map produces two distinct positions where the historical evidence is clearest, each corresponding to a different interpretation of operational strength.

SIGNAL A — Quality confirmation

V guard rail ($V \geq 65$) + $Q > 87$

Absolute level composite: $SE + RY + FCF$

All three operational metrics are simultaneously averaging in the top decile of their own 3-year quarterly history, while the stock is not at a historically elevated multiple. The market's current pricing is inconsistent with the operational record.

Fires: selectively, in dislocation regimes and idiosyncratic corrections. Average 2–3 entries per active year.

SIGNAL B — Momentum acceleration

V guard rail ($V \geq 55$) + delta $Q > 75$

Momentum composite: $SE + \delta RY + \delta OM$

Revenue growth and operating margin are both accelerating at above-historical rates, with above-average sales efficiency, and the stock is not at a historically elevated multiple. Earnings momentum is running ahead of current pricing.


Fires: consistently, across all market regimes including expansion years. Average 4–5 entries per active year.

The two signals are complementary, not correlated. Signal A requires absolute operational highs — rare, concentrating in dislocation environments. Signal B requires acceleration above historical norms — occurring more broadly, including when absolute levels are still recovering from a trough. Signal B fires in every calendar year; Signal A fires only in four of seven years, but with higher per-entry conviction.

Validation: what the historical record shows


Both signal readings are tested through a systematic backtest applied point-in-time across the full IGV universe over 2020–Q1 2026. The purpose is to examine whether the map's Dislocation readings correspond to any systematic pattern in subsequent relative returns — not to prove they will do so in future, but to report what the historical record shows.

Signal A results



95% Wilson CI: [62%–98%]. Excluding 2020: 8 entries, 87.5% WR, p = 0.035 (exact binomial). Directional signal persists without 2020; return magnitude is 2020-weighted. All p-values use exact binomial; normal approximation overstates significance at small n.

Signal B results



95% Wilson CI: [49%–81%]. Significant at 5% level. Value is consistent annual coverage, not high per-entry WR.

Year-by-year breakdown

Exhibit 2: Annual detail — Signal A and Signal B, 13-week hold vs. IGV ETF

Year	Market regime	A: entries	A: win rate	B: entries	B: win rate	Coverage
2020	COVID dislocation	3	100%	6	83%	Both active
2021	Bull market (earnings acceleration)	3	100%	3	50%	Both — A fires in bull market
2022	Rate shock correction	0	—	4	50%	B only
2023	AI narrative expansion	0	—	3	100%	B only
2024	AI re-rating peak	0	—	4	50%	B only
2025	Post-AI dislocation	3	67%	9	63%	Both active
Q1 2026*	Software sell-off (active)	2	100%	1	100%	Both active

*Q1 2026 partially seasoned. Signal A 2025: 2 of 3 entries outperformed (67%). Signal B fires every year — the three years Signal A is silent (2022–2024) are covered by Signal B.

The 2021 result for Signal A — 100% WR in a broadly constructive market — confirms the signal is not a crash-recovery strategy. Company-specific dislocations in constructive markets produce the same pattern as macro dislocations.

CURRENT MAP STATUS: Q1 2026

The 2026 sell-off has moved a meaningful portion of the IGV universe into elevated V territory. Based on trailing data at time of writing: approximately 20–25 IGV constituents have V scores above 60; of those, roughly 8–12 have Q composites above 75; and 3–5 have crossed the Signal A production threshold simultaneously. The map is in active Dislocation territory for a subset of the universe.

Using the map

The V×Q map has two layers of value. The first is the two signals: when Signal A or Signal B thresholds are met, the historical record shows a pattern of positive relative returns. The second — and more continuously useful — is the weekly diagnostic: the full sector map, updated from every quarterly filing and every week's prices, tells you which regime each name currently occupies and where the sector as a whole is positioned.

For tech operators: your stock has a map position every week

Most public company management teams track stock price and operational metrics separately — in different meetings, on different timescales, without a systematic framework for connecting them. The V×Q map overlays those two views simultaneously, for your own company and every peer.

If your company is in the Dislocation quadrant (V high, Q high), the trailing data does not confirm the market's pessimism. That is actionable for investor relations, buyback decisions, and how to frame guidance. If your company is in the Deterioration quadrant (V high, Q falling), the cheap stock is accompanied by genuine operational weakness — a confirmation, not a mispricing, that calls for a different set of operational questions.

The map's full value emerges when read across the whole universe. If 15 of your 20 closest peers are in elevated V territory, you are in a sector-wide narrative event, not a company-specific one — a distinction that matters for how you communicate to investors. If your Q composite has been declining for three quarters while peers' scores are stable, your operational trajectory is diverging from the sector in a way not yet visible in stock price alone.

For investors: a weekly regime diagnostic for software

The map replaces unstructured qualitative judgement with a calibrated weekly read. In early 2026, with IGV down 25%, an investor reading the full V×Q table can immediately answer: of the 80 IGV names, how many are in each quadrant? If the majority are in Deterioration, the sell-off is tracking genuine operational weakness. If the majority are in Dislocation, the sell-off has moved ahead of the trailing data. If evenly split, stock selection matters more than sector exposure. That analysis is available from the table in five minutes; it currently requires hours of individual earnings review without it.

The weekly V×Q table

The map will be published weekly on our website as a full scored table of every IGV constituent — V score, Q composite, individual metric scores, Signal A and B status, quadrant classification, and 4-week trend direction.

METHODOLOGY AND LIMITATIONS

Construction, constraints, and caveats

Construction summary

Component	Definition
Universe	IGV ETF constituents, point-in-time (end-of-month holdings). ~40–50 names in early years; 80–100+ later.
V score	Mid-rank percentile of trailing PE in 36-month rolling window: (obs below + 0.5×equal) / n × 100. Min. 45 obs per year.
Signal A Q	Arithmetic average of SE, RY, and FCF margin percentiles against trailing 12 unique quarterly obs.
Signal B Q	Arithmetic average of SE, δ RY, and δ OM percentiles against trailing 12 quarterly obs.
Sales efficiency	TTM revenue growth rate ÷ (S&M expense / TTM revenue).
RY / δ RY	RY: quarterly revenue ÷ same quarter prior year – 1. δ RY: this quarter's RY minus prior quarter's RY.
FCF margin / δ OM	FCF: TTM free cash flow ÷ TTM revenue. δ OM: this quarter's operating margin minus prior quarter's.
Signal A entry	$V \geq 65$ and $Q(SE+RY+FCF) > 87$ simultaneously on point-in-time data.
Signal B entry	$V \geq 55$ and $Q(SE+\delta RY+\delta OM) > 75$ simultaneously.
Exit / benchmark	13 weeks (91 calendar days) from entry. Benchmark: IGV ETF total return over identical window.
Gap rule	Min. 26 weeks between re-entries in same symbol.

Key limitations

V is a valuation guard rail, not a deep-value filter. Historical win rates are positive from V=50 (at-median) upward for both signals; the V threshold was a calibration choice within a flat zone, not an economically motivated cheapness requirement. Q is entirely backward-looking: high Q means the recent past was strong, not that the future is safe. In structural disruption, Q can remain at historical highs through several quarters before inflection appears in filed metrics. The three Q metrics were chosen for universal availability across the full IGV universe from standard filings — more precise metrics (NRR, CAC payback, Rule of 40) are not consistently available at scale.

Legal notice

This paper is produced for informational and educational purposes only. It does not constitute investment advice, a solicitation to buy or sell any security, or a recommendation of any investment strategy. The analysis is based on publicly available data and backtested models and is not a guarantee of future results. Recipients should conduct their own due diligence and consult their own legal, financial, and tax advisers before making any investment decisions. This paper has not been reviewed or approved by any regulatory authority. The author holds no investment adviser licence in any jurisdiction. Past performance is not indicative of future results.

About

ARPU is a weekly publication covering software and semiconductor businesses for operators, investors, and researchers. Beginning in 2026, our newsletter will publish a weekly V×Q scan of the IGV universe. For enquiries, email us at: arpu@hedder.com