The SaaS Illusion

When Useful Finance Becomes Financial Engineering

Hiram Samel Aditya Chaudhry

June 2026 v3

MIT SLOAN SCHOOL OF MANAGEMENT

Corresponding author: Hiram Samel, hsamel@mit.edu

Abstract

In the first weeks of 2026, roughly two trillion dollars in market capitalization disappeared from the U.S. software sector. The consensus explanation is that artificial intelligence disrupted the subscription model. This paper asks the prior question: were these valuations ever justified? Drawing on a panel of 115 publicly traded U.S. software companies tracked from 2016 through 2025, we argue that the SaaS sector constructed a financial architecture that allowed it to present results systematically disconnected from the underlying economics of its businesses. The central mechanism is the exclusion of stock-based compensation, the sector's largest routinely excluded cost, from the adjusted figures the market used to value these firms. Over the decade the sector expensed $153 billion in stock-based compensation while excluding it from reported performance, carried more than $320 billion in cumulative paid-in capital on its balance sheets by 2025, and accumulated roughly $75 billion in net losses. Across the panel the median GAAP operating margin was negative while the median adjusted margin, with that cost removed, was positive. We trace this permission structure to a four-decade genealogy in which useful financial conventions migrated beyond their original purpose, and we show how zero interest rate policy suppressed the market's incentive to question it. The beneficiaries sat inside the system that measured performance; the costs fell on the passive holders outside it. AI did not destroy two trillion dollars of value. It revealed that much of it was never there.

JEL Classification: G12, G14, G24, G32, M41

Keywords: stock-based compensation; non-GAAP measures; software-as-a-service; valuation multiples; permission structure; zero interest rate policy; passive investing

Acknowledgements: We are grateful to Don Lessard, Sanjay Sarma and Ben Armstrong for comments on earlier drafts. In addition, we thank Claude (Anthropic) and ChatGPT (OpenAI) for research assistance with literature organization, copy editing, structural review and document formatting; all analytical judgments, theoretical arguments, empirical specifications, and interpretations are the authors' own.

In the first six weeks of 2026, roughly two trillion dollars in market capitalization disappeared from the American software sector. Salesforce, the company that pioneered subscription software from a San Francisco apartment in 1999, lost more than thirty percent of its value in twelve months. The iShares Expanded Tech-Software Sector ETF (IGV) fell roughly 30 percent from its 52-week high. Median EV/revenue multiples across the panel (Table A5) compressed from double-digit levels in 2020–2021 (17.5x at the 2020 peak and 13.1x in 2021) to 4.4x by 2025, levels not seen since the mid-2010s. The financial press called it the SaaSpocalypse.

The consensus explanation is that artificial intelligence did this. Autonomous agents are beginning to compress product categories built around human seats. When an AI system can update a CRM, manage a project board, and triage a support queue without a human ever logging into the application, the per-seat subscription model breaks. This explanation is substantially correct. But it is also incomplete, because it assumes the two trillion dollars was actually there, the premise that what AI disrupted was a sector of genuine economic substance whose valuations reflected a reasonable assessment of what these businesses were worth.

The question nobody seems to be asking is the prior one: were these valuations ever justified?

The answer, on the evidence, is no. And the reasons why are not unique to software. They belong to a pattern that has recurred throughout American finance over the past four decades: useful financial tools and conventions, designed for specific analytical purposes, migrate beyond their original context, harden into an industry's accepted reporting practice, and eventually construct a permission structure that allows entire sectors to present financial results disconnected from the underlying economics of their businesses. We have seen this pattern in leveraged buyouts, where a legitimate debt-evaluation metric became a license to load companies with unsustainable leverage. We have seen it in private equity, where the same conventions justified fee structures that flattered the general partner at the expense of the limited partner. We have seen it in asset management, where performance reporting norms obscured the gap between gross and net returns. Each time, the original tool was useful. Each time, the migration was gradual. Each time, the beneficiaries were those closest to the architecture.

The SaaS sector did not invent this pattern. But it may be its most complete expression. On the architectural mechanics through which such norms harden into binding pricing constraints, see the Samel and Chaudhry (2026) companion paper. Over two decades, the software-as-a-service industry constructed a financial architecture of extraordinary

sophistication, one that combined the exclusion of one of its largest costs from reported results, the invention of an entirely new measurement vocabulary, a cash-flow cycle that disguised the consequences of its own compensation practices, and a macroeconomic environment that eliminated the market's incentive to ask hard questions. The result: a sector of 115 public companies that has absorbed more than $320 billion in equity capital and has, in aggregate, accumulated approximately $75 billion in net losses. Over the past decade alone, these companies expensed $153 billion in stock-based compensation, and then excluded it from the results they presented to the market. AI simply made this harder to ignore.

The evidence in this paper draws on a panel of 115 publicly traded US software companies, identified through the Thomson Reuters Business Classification system and tracked from 2016 through 2025 using financial data from LSEG Refinitiv. The panel is unbalanced by construction: companies enter at the point of their initial public offering. Only thirty were public by 2016; the remaining eighty-five listed during the decade that followed, most of them during the years when the permission structure described here was most permissive and capital most abundant. That the panel grows as zero interest rate policy deepens is not incidental to the analysis. It is part of what the analysis documents.

I. The Genealogy of Permission

How did it become acceptable for an entire sector to exclude one of its largest recurring costs from its reported results? The answer lies in a four-decade genealogy in which each generation of financial practice made the next seem reasonable.

The first step was EBITDA. In the leveraged buyout wave of the 1980s, firms like KKR needed a metric that stripped out financing decisions so they could evaluate whether a target company could service the debt they planned to impose (Jensen 1986). EBITDA did this. It was a tool, designed for a specific purpose, and it served that purpose well. But it also established a principle: real costs could be excluded from reported performance if you had a coherent story about why they were not the operating business's concern. Interest was the bank's problem. Depreciation was the accountant's abstraction. A small door, opened for a defensible reason.

By the early 1990s, EBITDA had migrated from leveraged finance to general corporate reporting. Companies that had never been through a buyout headlined it in their earnings releases. Analysts modeled around it. The door stayed open.

The second step was the dot-com extension. The buyout practitioners had said: ignore our financing costs. The internet entrepreneurs of the late 1990s said something bolder: ignore that we are not profitable at all. Growth was what mattered. Was the reasoning entirely wrong? It was not. Amazon would vindicate it, for exactly one company. But the logic was applied to hundreds of businesses whose unit economics could not work at any scale, and it took the NASDAQ's loss of eighty percent of its value to demonstrate the difference (Ofek and Richardson 2003).

Did the system learn? Briefly. Congress passed Sarbanes-Oxley. The SEC introduced Regulation G in 2003, requiring GAAP figures alongside non-GAAP metrics. This was called reform. In practice, it stipulated that the fiction could continue as long as the facts appeared somewhere nearby. Companies headlined the adjusted figures. They buried the reconciliation. Within five or six years, a new generation of companies emerged from San Francisco and Silicon Valley with a financial architecture more permissive than anything that had come before.

II. The Mechanism

The SaaS product model was a genuine advance. Deliver software over the internet. Charge a monthly subscription. Update continuously. Better for customers. More predictable for the business. Salesforce proved the concept; hundreds followed. The product model was real. What grew up around it was something else.

The financial architecture of SaaS synthesized the EBITDA principle, the dot-com growth thesis, and a new mechanism. The mechanism was the treatment of stock-based compensation, and it is worth stating precisely, because precision is what the conventional treatment is designed to prevent.

A SaaS company hires an engineer for $400,000 a year. Half is cash. Half is stock. Under GAAP, both are expenses. Both hit the income statement. The company reports a loss. Then it releases its 'adjusted' results. The stock-based compensation has been removed, 'added back,' as though it were an item mistakenly included. The rationale: SBC is 'non-cash.' It does not involve a wire transfer.

Is this reasoning sound? It is not. The engineer's labor is real. The product cannot be built without it. When the company issues stock to pay for that labor, it is not avoiding a cost. It is transferring one: from the income statement, where it would reduce reported earnings, to the equity structure, where it appears as dilution. That the cost is settled in shares rather than cash is the defense usually offered for the exclusion, but it is also the point. Depreciation is non-cash too, and no one treats its removal as a measure of

profitability. The absence of a wire transfer does not make the expense less real; it changes only who bears it and when. The shareholder absorbs it through dilution rather than the company through cash, and the cost is deferred into the equity structure rather than eliminated from it. If a restaurant excluded the cost of food from its margins, we would know what to call it. When a software company excludes the cost of its workforce, we call it ‘adjusted EBITDA.’

The scale of the exclusion has grown steadily. In 2016, the median SaaS company spent 9.3 percent of its revenue on stock-based compensation. By 2022, the figure had more than doubled to 19.4 percent. By 2025, median SBC intensity had eased to 14.9 percent of revenue, with the seventy-fifth percentile at 20.3 percent. Across the full 115-company panel, SaaS firms spent over $28 billion on stock-based compensation in 2024 alone. Cumulatively, from 2016 through 2025, the sector expensed $153 billion in SBC, and excluded it all from the adjusted results the market used to value these companies.

The effect is to transform reported results. Across all company-years in the panel, the median GAAP operating margin is –8.6 percent. The median adjusted EBITDA margin, with SBC excluded, is +6.5 percent. The gap between these two figures, roughly fifteen percentage points, is the distance between what these businesses actually earn and what they report to the market. In 2025, half the sector has finally crossed to GAAP profitability, but only after the rate shock and AI threat revoked the permission structure that had sustained losses for a decade. The SBC exclusion remains the difference between reported ‘profit’ and actual economics. The convention was never universal. Microsoft, Netflix, and Intel had long included equity awards as costs in their non-GAAP results, and the SEC questioned the add-back as early as the mid-2000s on the grounds that paying employees in equity is plainly a performance incentive rather than an item to be set aside. What the panel documents is not the practice in its origins but the practice in retreat, and the retreat is itself evidence that the exclusion was always a choice the conditions permitted rather than an economic fact the numbers required.

A reasonable objection: GAAP itself understates the value of intangible-intensive firms by expensing internally generated intangibles immediately. Recent work by Ewens, Peters, and Wang has shown that capitalizing portions of R&D and SG&A materially improves the valuation relevance of accounting data for firms whose assets are primarily intangible. This is a real limitation of GAAP, and it applies to software companies. But it does not rescue the SaaS reporting architecture. The intangibles argument explains why GAAP operating margins might be too low for a company investing heavily in product development. It does not explain why the sector built an entire measurement vocabulary to avoid reporting GAAP margins at all. It does not

explain why the market rewarded unprofitable companies at higher multiples than profitable ones. And it does not explain why fifty-six percent of SaaS firms never sustained two consecutive years of profitability over a decade. The gap between “GAAP is an imperfect measure for intangible firms” and “therefore excluding one of your largest recurring costs from reported results is justified” is the gap in which the permission structure operates.

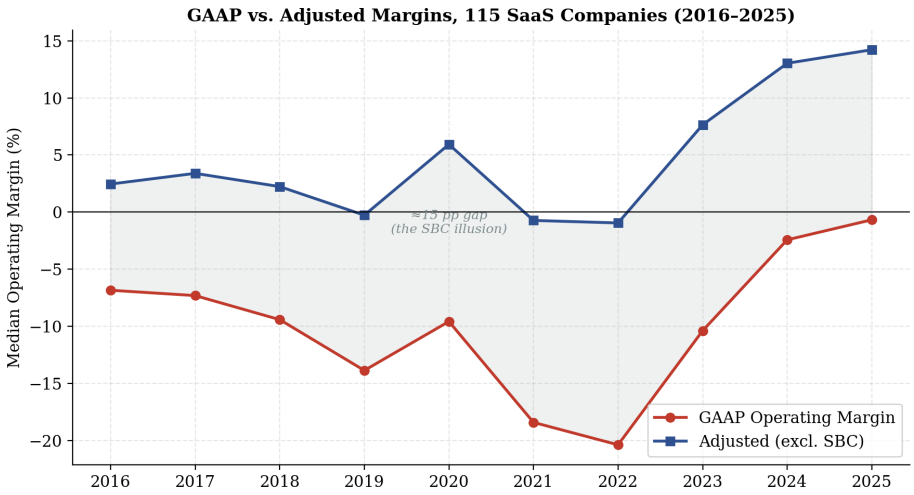

Figure 1. GAAP vs. Adjusted Margins, 115 SaaS Companies (2016–2025). Source: Author’s analysis.

III. The Alphabet of Illusion

The SBC exclusion was the foundation. But an illusion sustained across an entire sector for two decades required more than a single accounting convention. It required a measurement vocabulary, one sophisticated enough to give the appearance of analytical rigor while operating, at every level, on the adjusted figures that excluded one of its largest costs.

ARR , annual recurring revenue, became the sector’s preferred measure of scale. It annualizes a company’s current subscription run rate. But ARR is not a GAAP metric. It is not audited. Its calculation varies from company to company. And because it is always presented as a forward-looking figure, it tends to overstate the revenue a company will actually recognize.

NRR , net revenue retention, measures how much revenue a company retains from existing customers, including expansions and net of churn. As a diagnostic, NRR is useful. But it says nothing about profitability. A company can have spectacular retention while losing money on every customer it retains.

The Rule of 40 is the most revealing. It holds that a 'healthy' SaaS company's revenue growth rate plus its profit margin should exceed forty percent. This makes growth and profitability *interchangeable*, a company can score well while hemorrhaging cash, provided it is hemorrhaging fast enough. And the margin used is the adjusted margin, with stock-based compensation excluded.

How much difference does this make? In 2024, twenty-two percent of the 115 SaaS companies in our panel passed the Rule of 40 using adjusted EBITDA margins. Using GAAP operating margins, the figure fell to six percent. CrowdStrike typically scores in the 40s to 50s on the adjusted measure and in the 15s to 30s on GAAP, a gap of roughly 20 to 30 points, entirely SBC. Cloudflare and Samsara show gaps of similar magnitude, above 20 points each, across recent quarters. The convention raised the Rule of 40 pass rate from six percent under GAAP operating margins to twenty-two percent under adjusted EBITDA margins, but the headline is always the minority who 'pass.'

There were others: *LTV/CAC, the 'magic number,' 'months to recover CAC'*, each calibrated against adjusted figures excluding the same cost. Taken together, they formed a closed system of self-referencing measurements, internally consistent and externally disconnected from the question that should be prior to all others: does this business make money?

IV. Where the Cash Went

If a company reports that it is 'free cash flow positive' while excluding the cost of its workforce from its performance metrics, where does the cash actually go? Salesforce, the sector's largest, oldest, and most closely watched company, provides the most complete answer, and its story illustrates every feature of the architecture.

In early 2023, under pressure from activist investors, Salesforce expanded its share buyback program to $20 billion, the largest in the company's history. The financial press praised it as a pivot toward shareholder value. Analysts upgraded the stock. Here is what actually happened.

Over the 2016–2025 period, Salesforce generated $25.0 billion in cumulative GAAP operating profit. This was a genuine achievement, one of the few SaaS companies to produce sustained operating earnings at scale. But during the same period, the company expensed $21.7 billion in stock-based compensation. It then spent $28.1 billion on share buybacks, the vast majority after the 2023 announcement. The buybacks exceeded the cumulative SBC by $6.4 billion. Salesforce spent more than all

of its stock compensation and then some, just managing dilution. The ‘pivot to shareholder value’ was a treadmill.

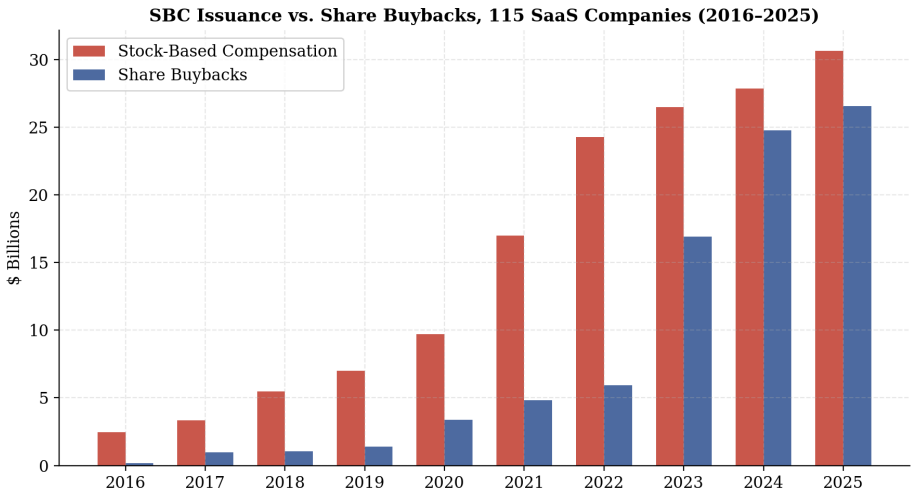

Figure 2. SBC Issuance vs. Share Buybacks, 115 SaaS Companies (2016–2025). Source: Author’s analysis.

But the buyback treadmill is only half of Salesforce’s capital story. The other half is acquisitions. In 2019, Salesforce acquired Tableau for $15.7 billion, entirely in stock. In 2021, it acquired Slack for $27.7 billion, in a roughly even combination of cash and stock. These transactions are visible in the balance sheet as two extraordinary jumps in paid-in capital: $18.2 billion in 2019, $15.3 billion in 2021. Together with smaller deals, stock-financed acquisitions account for roughly $33 billion of Salesforce’s $64.6 billion in lifetime paid-in capital.

Were these bad acquisitions? Not on the face of it. Tableau brought genuine data visualization capabilities. Slack brought workplace communication. Each made strategic sense in isolation. But neither improved the underlying economics of the business. Salesforce’s GAAP operating margins in the years following these acquisitions did not expand. What expanded was revenue, by adding the targets’ top line to Salesforce’s ARR while spreading dilution across a larger shareholder base, many of whom were Tableau and Slack shareholders who never chose to own Salesforce stock. The acquisitions did not build a more profitable company. They built a larger one, which is a different thing, and the conflation of the two is precisely what the permission structure enables.

Salesforce’s lifetime paid-in capital thus decomposes into three channels, each of which monetized the inflated equity: $21.7 billion in SBC (paying the workforce), roughly $33 billion in stock-financed acquisitions (buying revenue), and the remainder in IPO proceeds and other equity issuance (funding operations). Every channel was

individually defensible. Collectively, they produced a company that, after absorbing over $64 billion in equity, has $16.4 billion in retained earnings. A 25 percent return over a quarter century, from the most successful SaaS company in history.

In 2025, following the celebrated efficiency pivot, Salesforce reported $12.4 billion in free cash flow. But reported FCF benefits from treating SBC as a non-cash item: the company added back over $3 billion in stock-based compensation to arrive at that figure. If SBC is a real cost, and it is, because the shares are real and the dilution is real, then the underlying cash generation is closer to $9 billion. Of that, roughly $8 billion went to buybacks, mostly to offset the very dilution the SBC created. The residual: approximately $1 billion in genuinely free cash on a $245 billion enterprise value. The efficiency pivot is real. But much of the cash it produces goes straight back into servicing the equity structure. The discipline is a treadmill.

The pattern is not unique to Salesforce. Across the full panel, from 2018 through 2022, five years spanning the late ZIRP era and the rate shock, the entire SaaS sector's free cash flow minus SBC was approximately zero. Five years. Roughly $61 billion in reported free cash flow across the sector. Net cash generation after recognizing the actual cost of the workforce: slightly negative.

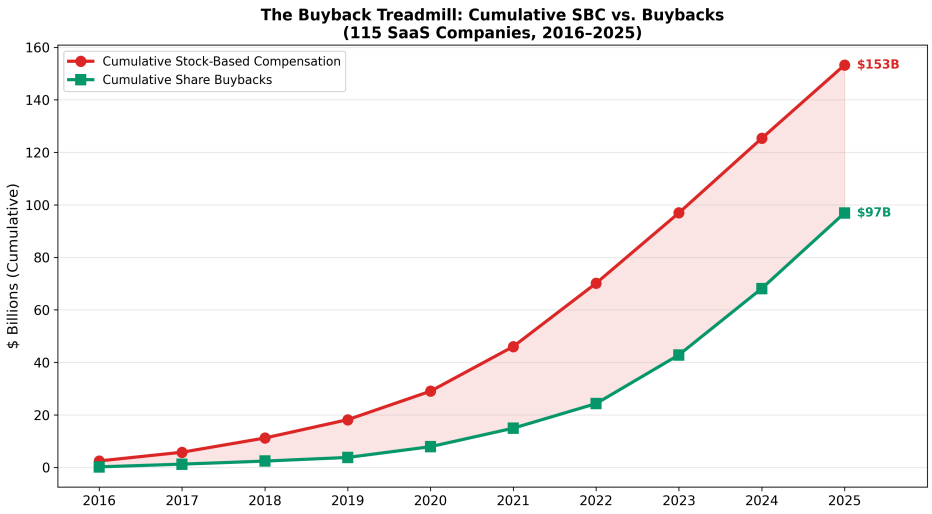

Figure 3. The Buyback Treadmill: Cumulative SBC vs. Buybacks, 2016–2025. Source: Author's analysis.

V. The Evidence

Salesforce, at least, eventually produced operating profits. Most SaaS companies never did. Of the 115 companies in our panel, 56 percent never achieved two consecutive years of positive GAAP operating income through 2025. At the three-year threshold, 70 percent never did. Measured by net income, the bottom line that accounts for all costs,

including taxes and interest, 82 percent never achieved three consecutive profitable years.

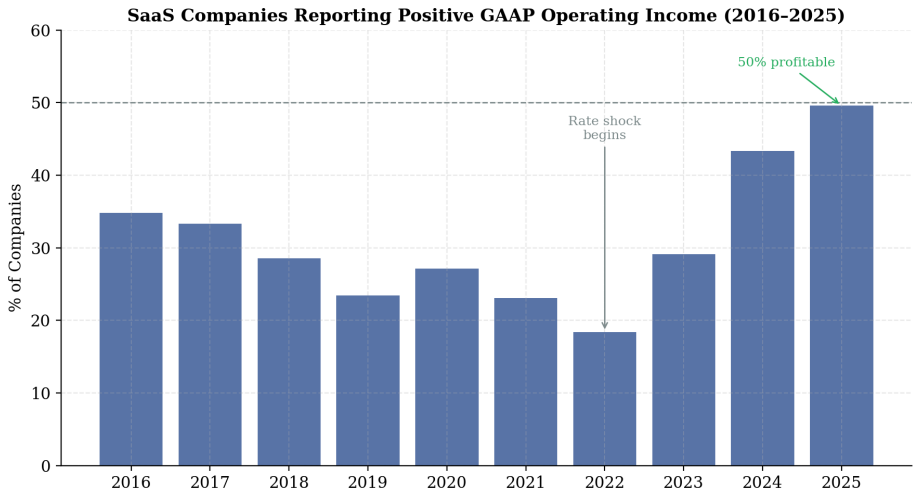

Figure 4. Percentage of SaaS Companies Reporting Positive GAAP Results (2016–2025). Source: Author's analysis.

Consider CrowdStrike: $4.6 billion in annual revenue, a $132 billion enterprise value, and a cumulative GAAP operating loss of $1 billion across its public life. Its cumulative SBC, $3.4 billion, more than triples the operating loss. Snowflake: $4.4 billion in revenue, $92 billion in enterprise value, a cumulative operating loss of $5.2 billion, and $5.6 billion in cumulative SBC. Cloudflare: $2.0 billion in revenue, $74.4 billion in enterprise value, never once profitable. Twilio: $4.9 billion in revenue, a cumulative operating loss of $3.1 billion, shareholders diluted 87 percent. These companies have been given time, capital, and every advantage the non-GAAP framework could provide. None of them has converted that advantage into sustained profitability.

The dilution these losses produce is persistent and compounding. The median SaaS company dilutes shareholders by 2 to 4 percent annually. Over a decade, this compounds relentlessly: Twilio diluted shareholders 87 percent, Okta 86 percent, MongoDB 61 percent, HubSpot 52 percent. The cost is borne disproportionately by passive fund beneficiaries, the holders of index funds, ETFs, and the retirement and pension accounts beneath them who own these stocks by virtue of sector weights rather than by conviction. The dilution they absorb was priced into valuations by conventions that excluded its cost. When the conventions contract, as they are now contracting, these same holders bear the revaluation. The beneficiaries of the convention during its long period of stability, the employees who received equity at the intensity the convention permitted, the founders whose concentration compounded through early issuance, the intermediaries whose fees depended on the volumes transacted, are identifiable, though no individual among them designed the arrangement that made

them beneficiaries. The asymmetry between who benefits from a permission structure and who pays when it contracts is not incidental to what a permission structure is. It is what a permission structure is.

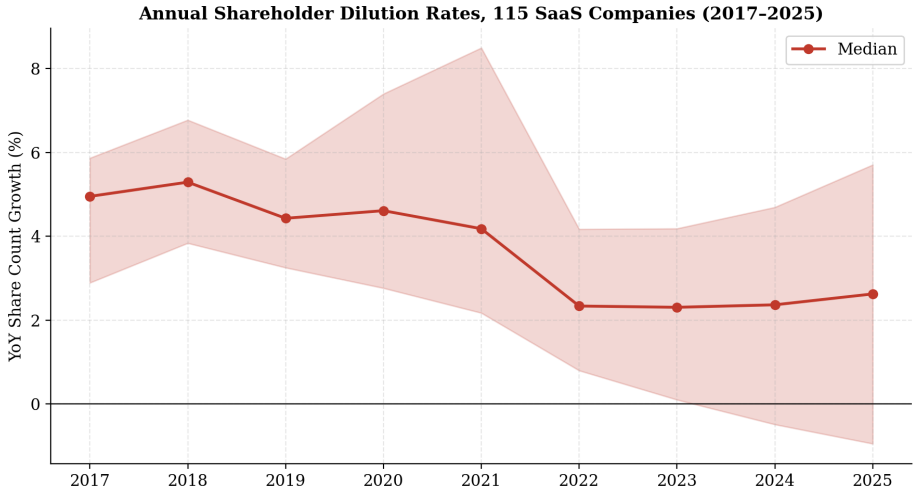

Figure 5. Annual Shareholder Dilution Rates, 115 SaaS Companies (2017–2025). Source: Author's analysis.

The retained earnings tell the aggregate story. In 2025, 84 percent of SaaS companies carried accumulated deficits on their balance sheets. The sector's total retained earnings: negative $74 billion, after having reached a peak deficit of nearly $80 billion at the end of 2023, before the efficiency pivot began to erode it. Against approximately $322 billion in paid-in capital. The sector has consumed $74 billion more than it has ever earned, and it has required $322 billion in equity to get there.

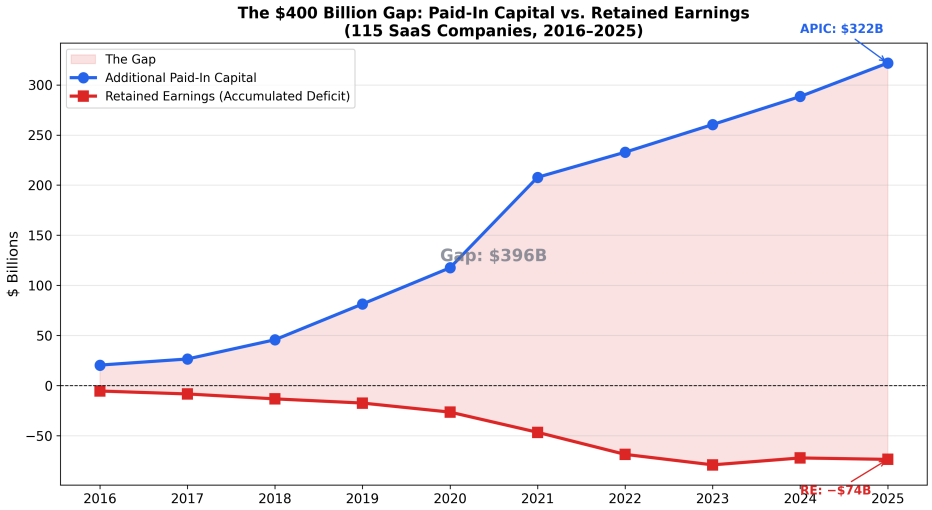

Figure 6. The $400 Billion Gap: Paid-in Capital vs. Retained Earnings, 115 SaaS Companies (2016–2025). Source: Author's analysis.

VI. The Interest Rate Machine

Conventions create the space for an illusion. They do not, by themselves, sustain it for fifteen years. What sustained it was the most extraordinary interest rate environment in the history of modern central banking.

From 2008 through 2022, U.S. monetary policy operated in an unusually low-rate regime, including long stretches at or near the zero lower bound. Why does this matter specifically for SaaS? Because of what low rates do to the present value of distant cash flows. At a ten percent discount rate, a dollar of profit expected a decade from now is worth thirty-nine cents today. At four percent, sixty-eight cents. As rates approach zero, the difference between a company that is profitable now and one that merely promises profitability in the future nearly disappears.

In 2020, at the peak of ZIRP, the median SaaS company traded at 17.5 times revenue. The seventy-fifth percentile: 34.0 times. By 2022, after the fastest rate-tightening cycle in forty years, the median had fallen to 5.0 times. Nothing about these companies' products or customers had changed. The cost of patience had changed. That was enough.

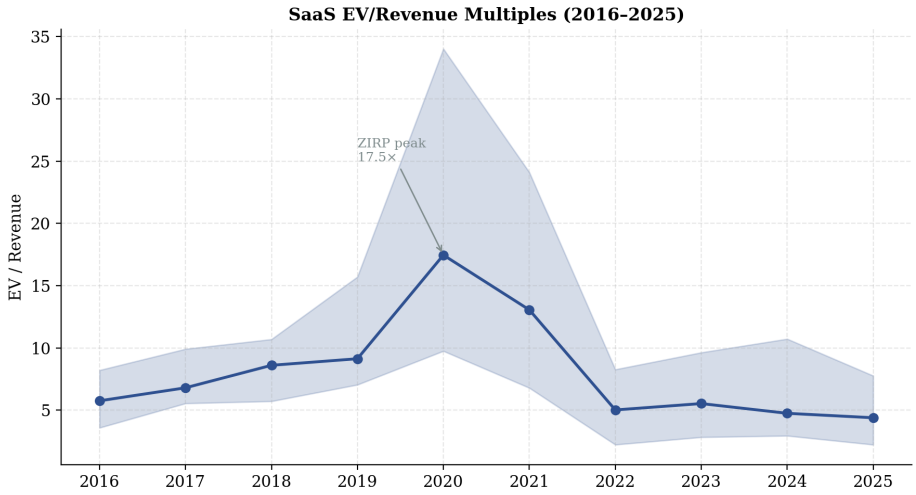

Figure 7. SaaS EV/Revenue Multiples (2016–2025). Interquartile range shaded. Source: Author's analysis.

The result was a reflexive loop. High valuations made SBC attractive. Attractive SBC recruited talent. Talent produced growth. Growth justified higher valuations. The loop did not require profitability at any point. Hanno Lustig and his co-authors have demonstrated that low interest rates systematically favor long-duration assets: claims on the distant future (Greenwald et al. 2021; Gormsen and Lazarus 2023). SaaS companies are long-duration assets by construction. When rates fell, the value of their promises inflated. When rates rose, the BVP Nasdaq Emerging Cloud Index fell more

than sixty percent from its November 2021 peak to its 2022 trough (Bessemer Venture Partners and Nasdaq, EMCLOUD index data). The round-trip in our panel is nearly symmetrical: median multiples of 5.8x in 2016, 17.5x at the 2020 peak, and 4.8x in 2024 (Table A5). A decade of valuation expansion, erased.

VI-b. Testing the Mechanism

The account so far presents two causes operating in parallel: a rate environment that inflated the value of long-duration claims on the future, and a permission structure that allowed an entire sector to exclude one of its largest costs from the metrics the market used to price those claims. A careful reader will ask whether this is two causes or one. If the rate environment alone did the work, then what the paper has documented is a duration phenomenon in need of a less institutional explanation. The question has an answer. It is not the answer one might expect.

Consider the period from the 2021-Q4 peak to the 2022-Q4 trough, the cleanest window in which the rate shock did its work. Median SaaS multiples fell from 13.0x revenue to 5.0x, a compression of roughly sixty percent. Regressing each firm's percentage compression in EV/revenue on SBC intensity at peak yields no significant relationship, on its own ($p = 0.30$) or after adding revenue growth and operating margin at peak ($p = 0.34$). Revenue growth absorbs the cross-section of who fell most; SBC intensity adds nothing once growth is in. For the event of repricing itself, the duration story is sufficient. Rates fell, duration amplified, rates rose, duration corrected. The permission structure did not, in the cross-section, add independent explanatory power to the fact of the correction.

This is worth stating plainly, because a careful reviewer running the same regression will reach the same conclusion. It is also not the paper's full claim. The correction was not the phenomenon. The correction was the moment at which the phenomenon ended. The phenomenon was the two decades during which it was possible to construct financial architectures that the correction would eventually unwind, and the relevant empirical test is not which firms fell most in the unwinding but which firms had constructed the most to unwind.

That question has a clear cross-sectional answer, and it is preserved even after the full set of controls. Firms with higher SBC intensity at the 2021 peak cut SBC intensity more aggressively through 2025, a coefficient of negative 0.65 that survives controls for growth, starting margin, size, leverage, cash position, stock return, and subsequent growth. A firm at thirty percent SBC intensity at peak cut SBC by roughly six and a half percentage points more of revenue than a firm at twenty percent. The effect is large, it is

stable across specifications, and it describes a real behavioral change rather than an accounting artifact.

What the effect describes more precisely is a compression of the permitted range of practice. At the 2021 peak, the spread between the tenth and ninetieth percentiles of SBC intensity across the sector was 33 percentage points of revenue. By 2025, that spread had compressed to 21 percentage points, and the compression was asymmetric: the ninetieth percentile fell from 38.7 percent to 26.1 percent, while the tenth percentile was essentially unchanged at around 5 percent. If this were pure mean reversion toward a stable sectoral equilibrium, both tails would move toward the center at similar rates. They did not. The top of the distribution compressed while the bottom remained where it was. That is the signature of a norm tightening rather than a symmetric duration unwinding.

The firms doing the cutting were not under financial duress. Firms in the top quintile of 2021 SBC intensity held median cash balances of roughly two years of revenue, three times the cash-to-revenue ratio of firms in the bottom quintile. Leverage was essentially zero across the distribution. These firms were not forced to cut by liquidity distress; they had ample runway to continue the practices the permission structure had allowed. They cut because the practices themselves were no longer defensible within a professional community whose tolerance had visibly contracted. What the measurement system had permitted, the measurement system could no longer permit. The rate environment made the question of profitability unavoidable. The permission structure determined how far from profitability firms had traveled while the question could be avoided.

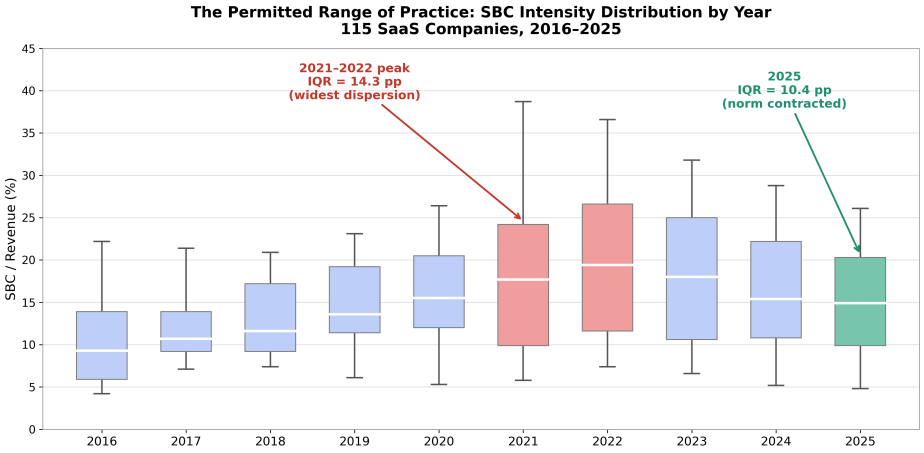

Figure 8. The Permitted Range of Practice: SBC Intensity Distribution by Year, 115 SaaS Companies (2016–2025). Box plots show 10th–90th percentile range with interquartile box and median line. The ninetieth percentile of SBC intensity fell from 38.7 percent of revenue in 2021 to 26.1 percent in 2025; the tenth percentile remained essentially unchanged. Source: Author's analysis.

The two mechanisms are complementary rather than competing. The rate environment provided the conditions under which the permission structure could operate unchecked for as long as it did. The permission structure, operating unchecked, produced the specific character of vulnerability the rate environment eventually exposed. Neither mechanism is sufficient on its own. What we observe in the cross-section is not the failure of any individual actor within the system. It is the behavior of a professional community that produced a measurement regime serving its own immediate interests, sustained it through a rate environment that rewarded its assumptions, and then renegotiated what was permitted when the assumptions no longer held. The community was not corrupt. The regime was not designed. The outcome was cumulative, and it was durable because no single participant had both the means and the incentive to disturb it: the same logic, at the professional-community level, that Shleifer and Vishny (1997) identified at the individual-arbitrageur level: the cost of challenging a persistent convention falls on the challenger, and the benefit, if any, is diffuse.

VII. What the Exceptions Reveal

At this point, a reasonable reader will object: some SaaS companies *are* profitable. Fortinet earned more than $7 billion in cumulative operating profit over the past decade while *shrinking* its share count. AppLovin operates at a 40 percent GAAP margin while buying back stock. The Trade Desk has been profitable for ten consecutive years. If the model works for them, perhaps the unprofitable companies simply need more time.

But the exceptions do not vindicate the system. They indict it. Because the financial architecture that sustained CrowdStrike's $132 billion valuation is the same architecture that sustained Fortinet's $62 billion valuation. Both reported adjusted metrics. Both were evaluated using the proprietary SaaS vocabulary. Both appeared in the same indices and the same analyst coverage. The difference is that Fortinet actually made money, and the conventions the market used could not distinguish between the two.

Consider the two companies side by side. Fortinet and CrowdStrike compete in adjacent cybersecurity markets, serve overlapping customer bases, and are tracked by the same sell-side analysts using the same vocabulary. Both appear in the iShares Expanded Tech-Software ETF. Over 2016–2025, Fortinet produced $7.4 billion in cumulative GAAP operating profit; CrowdStrike produced $1.0 billion in cumulative operating losses. Fortinet reduced its share count; CrowdStrike diluted shareholders by twelve percent. Yet through most of the period, CrowdStrike traded at roughly twice Fortinet's revenue multiple. A measurement regime that rewarded one over the other by that factor, in the face of such different economic performance, is not noisy. It is blind to

the thing that should have distinguished them. The problem is not that Fortinet is undervalued, and it is not that CrowdStrike is overvalued. The problem is that the metrics the market used to price both companies could not tell the difference between them.

This is the deeper problem. When measurement conventions cannot reliably differentiate between a company that earns real profits and one that has never sustained profitability, the market's capacity to allocate capital toward its most productive uses is structurally weakened. It is not that CrowdStrike is a bad company; its products may be excellent, its technology genuine. It is that the market lacked the mechanism to ask whether $132 billion was the right price for a company with $4.6 billion in revenue and $1 billion in cumulative losses. The Rule of 40, adjusted EBITDA, ARR growth: none of these tools could pose the question, because the question they were designed to avoid is the one that matters; does this business earn more than it spends?

If we believe in efficient markets, then we must ask why the market's correction mechanism was disabled for the better part of two decades. And the answer is the incentive structure. Venture firms whose returns depend on high-multiple exits do not demand profitability before IPO. Investment banks earning fees from SaaS offerings do not question the conventions. Sell-side analysts model around adjusted metrics because that is what every other analyst does. Compensation consultants benchmark pay against peers using the same figures. Auditors sign off on the GAAP numbers and maintain silence about the adjusted results the market actually uses. Passive index funds hold whatever enters the index, regardless of whether the company has ever earned a profit, a dynamic documented empirically by Jiang, Vayanos, and Zheng (2025) and extended to the valuation-misallocation argument in Samel and Chaudhry (2026). Everyone is acting within their professional norms. No one is corrupt. And the system produces an outcome none of them would defend if asked to describe it plainly: an outcome in which the beneficiaries sit inside the system that measures its performance, and the costs are borne by holders outside it who have no mechanism to observe what is being measured on their behalf.

VIII. Good Until It Isn't

Consider HubSpot. Founded in 2006. IPO in 2014. A company widely admired for its culture, its content marketing, and its product. For its first decade as a public company, HubSpot never achieved a GAAP operating profit. Its cumulative operating loss: $600

million. Its cumulative stock-based compensation: $2.2 billion, nearly four times the loss. Its 2025 enterprise value: $23.2 billion.

Is HubSpot a bad company? It is not. Its customers like the product. Its employees are well-treated. It has grown revenue from $186 million to $3.0 billion. Taken on its own terms, it is a successful business. Except that it has never made money. And it has never been required to.

This is what a benign environment produces. Not fraud. Not incompetence. A company that, had the measurement system made the question of profitability visible, would have faced it in year eight or ten. It did not, because the measurement system the market used did not frame profitability as a condition of continued access to capital. HubSpot is not the failure. The absence of the question is. The SaaS permission structure removed the question for twenty years, and the result is not a single bad actor but an entire sector in which profitability became, for most of its participants, optional.

The pattern has a precedent, and the precedent is instructive. In American manufacturing, the equivalent permission structure was the quarterly earnings metric. Shed assets, offshore production, strip out costs, and EPS rises, the stock price follows, the CEO is rewarded. Each decision was individually rational. Each was rewarded by the market. And collectively, over three decades, the practice hollowed out the productive capacity of the American industrial economy. By the time anyone noticed, the capabilities were gone. The market mechanism that should have said ‘you are consuming your own productive base’ could not say it, because the metric the market used treated asset-shedding as a positive.

The structure is the same. In manufacturing, the measurement system rewarded behavior that was individually enriching but collectively destructive. In SaaS, the measurement system rewarded growth-without-profitability that directed more than $320 billion in capital toward businesses that produced roughly $75 billion in cumulative losses, while the sectors that needed capital, energy, infrastructure, manufacturing, healthcare, were told they were not ‘innovative’ enough to deserve it. Samel and Chaudhry (2026) documents the mirror image of this misdirection on the manufacturing side: a capital-markets architecture that consistently underprices physically productive firms because the metrics it rewards were not designed to recognize them. In both cases, the conventions hardened into permission structures. In both cases, the people closest to the architecture benefited. In both cases, the costs were borne by people furthest from it: factory workers in one case, pension beneficiaries in the other. And in both cases, the correction came late, from outside the system, and destroyed value that was assumed to be permanent.

SaaS is good until it isn't. Offshoring is efficient until it isn't. The leveraged buyout creates value until it doesn't. Each is an instance of the same phenomenon: a locally correct observation, subscription software is better than license software, global labor arbitrage reduces costs, debt disciplines management, that hardens into an institutionalized permission structure that prevents the market from recognizing when the original logic no longer applies. The failure is not in the original insight. It is in the system's inability to know when to stop.

IX. What the SaaSpocalypse Revealed

The current selloff is being understood as a technology story. AI arrived, the seat-based model broke, and two trillion dollars in value was destroyed. This narrative is correct in its particulars. AI agents are real. Seat compression is real. Per-seat subscriptions depend on humans as the unit of consumption. When agents do the work, the seat disappears. With the seat goes the recurring revenue narrative. With the narrative goes the multiple. With the multiple goes the stock price on which the SBC was valued. And the reflexive loop comes apart.

But a disruption story assumes that what was disrupted had substance. A sector in which four out of five companies never sustained a profit by net income. In which $153 billion in stock-based compensation was expensed over a decade and excluded from every adjusted figure the market used. In which cash generated by the exclusion was recycled into nearly $100 billion in buybacks to offset the dilution the exclusion caused. In which $322 billion in equity capital produced roughly $75 billion in cumulative losses. This is not a sector that was destroyed by AI. It is a sector whose permission structure AI made harder to maintain.

The consequences extend well beyond shareholders. When an entire sector's reporting conventions obscure whether its companies make money, the distortion propagates. Capital flows toward the sector with the most permissive conventions rather than the most productive use. Talent follows the capital: a company that can offer $400,000 in total compensation, half in stock valued under conventions that exclude its cost, can outbid employers in virtually every other sector. And the conceptual cost may be the most enduring: when policymakers and investors equated innovation with software and growth with SaaS, the companies that build physical systems, employ people at scale, and maintain the infrastructure on which everything else depends receded to the margins of the conversation. The permission structure did not merely distort capital allocation. It distorted our understanding of what an economy is for.

The question this forces is not whether SaaS was good or bad. Subscription software is a better delivery model. Cloud computing works. Some companies built real businesses. The question is whether a market that cannot distinguish productive investment from capital consumption is functioning as a market at all. And if it is not, if the conventions that sustained two decades of misallocation could be individually defensible and collectively disconnected from reality, then we should ask where else this pattern is operating. We should ask it about the AI companies now trading at multiples that dwarf what SaaS ever achieved, using many of the same conventions. We should ask it about every sector in which the distance between reported results and economic reality has been quietly widening for years, sustained by an incentive structure that no individual participant has reason to challenge.

How can the market hold participants to account while fostering innovation? The answer is not to kill non-GAAP metrics: they have legitimate analytical uses. It is to restore the feedback loop. Productive patience requires a mechanism for verification. Amazon was productively patient: its losses were declining as a percentage of revenue, its unit economics were visibly improving, and there was a testable path to profitability. The SaaS conventions eliminated the verification mechanism. The market could not tell whether the path to profitability existed, because the metrics were designed to make the question irrelevant. Reform means making the GAAP figure the headline and the adjusted figure the footnote, rather than the reverse. It means restoring the market's ability to ask the question that every business, however innovative, must eventually answer.

There are signs the regime is shifting. On February 25, 2026, in its Q4 FY2026 earnings release, Nvidia, the most valuable company in the world, announced that beginning in Q1 FY2027 it would include stock-based compensation expense in its non-GAAP financial measures (NVIDIA 2026). The most important company in the technology sector just conceded the central premise: SBC is a real cost and excluding it distorts reported economics. If the convention's most powerful beneficiary has abandoned it, the permission structure is contracting.

The sector-wide data makes the same point the Nvidia concession makes at the level of a single firm. By 2025, half the SaaS sector is GAAP profitable for the first time in a decade. Median operating margins have recovered from approximately negative twenty percent in 2022 to near zero. Sector-level SBC intensity has declined from its peak, and as documented in Section VI-b, the decline is concentrated at the top of the distribution rather than the bottom: the signature of a norm tightening rather than a symmetric correction. The range of what is permitted has contracted; firms that had operated at the permitted edge have moved toward the new center. The speed of the adjustment now

shows that the firms were correctable throughout the preceding decade. What the measurement system enforced in three years through the mechanism of revoked permission, the measurement system had failed to enforce for a decade through the mechanism of market discipline, not because the firms were uncorrectable, but because the conventions the market used to evaluate them could not frame the question of correction. The answer arrived only when the conditions that had sustained the question's suppression no longer held.

The SaaS illusion is not the story of a single sector's financial misadventure. It is a case study in what happens when useful finance becomes financial engineering: when the tools designed to illuminate are repurposed to obscure, and the permission to obscure becomes the organizing principle of an industry. It is also one expression of a broader arrangement in which measurement regimes serving those who design them outlast the conditions that produced them, and in which the costs of their persistence fall disproportionately on the passive fund beneficiaries whose savings are allocated by index weight rather than by conviction. The pattern has run before. Its mirror, run in the opposite direction across American manufacturing, is documented in Samel and Chaudhry (2026). There will be other instances. There will, almost certainly, be other sectors in which the distance between reported economics and underlying reality is being quietly sustained by conventions that no individual participant has the standing to disturb. Identifying them is work for the next paper, and the one after that.

The platform was never built. It was always going to be built next year.

A note on sources

Aggregate SaaS-sector financial data (stock-based compensation, revenue, operating profit, buybacks, paid-in capital, retained earnings, share count, and enterprise value) are compiled from LSEG/Refinitiv fundamentals for the panel of 115 U.S.-listed software companies classified under the Thomson Reuters Business Classification (TRBC) SaaS code, covering fiscal years 2016 through Q3 2025 (firm fiscal-year conventions vary). Market capitalization and enterprise value figures reflect closing values as of the most recent reporting date available in the panel. The iShares Expanded Tech-Software Sector ETF (IGV) reference decline and median EV/revenue compression figures are drawn from the same LSEG panel; the "SaaSpocalypse" phrasing is media shorthand that appeared in financial press coverage of the 2026 software selloff. Nvidia's February 25, 2026 announcement is cited from the company's Q4 FY2026 earnings release (NVIDIA 2026). Company-level figures for Salesforce, CrowdStrike, Snowflake, HubSpot, Datadog, Cloudflare, Twilio, Okta, MongoDB, ServiceNow, Workday, Fortinet, AppLovin, and The Trade Desk in the body text and Table A4 are cumulative totals computed directly from the panel across each firm's period of public reporting within 2016–2025. The 2021–Q4 peak-to-2022–Q4 trough repricing regression and the 2021–Q4 to 2025–Q3 SBC-retreat regression in Section VI-b use the same panel; full specifications are reported in the Technical Appendix (Part C) that follows the references. Paid-in capital is reported as the panel-wide sum of additional paid-in capital (APIC) measured at each firm's most recent 2025 reporting date; it is a balance-sheet stock accumulated since each firm's founding, not a flow of capital raised within the 2016–2025 window. Of the roughly $322 billion total, approximately $42 billion was on the books at panel entry (incumbents' 2016 balances plus entrants' paid-in capital at IPO) and the remainder accumulated over the panel period, of which roughly $153 billion is capitalized stock-based compensation. Figures in this paper derive from the corrected fundamentals panel finalized in June 2026 and are aligned with the June 2026 manufacturing companion paper (Samel and Chaudhry 2026).

References

Bessemer Venture Partners, and Nasdaq. n.d. "BVP Nasdaq Emerging Cloud Index (EMCLOUD)." Index data series.

Ewens, Michael, Ryan H. Peters, and Sean Wang. 2025. "Measuring Intangible Capital with Market Prices." *Management Science* 71(1): 407–427.

Gormsen, Niels Joachim, and Eben Lazarus. 2023. “Duration-Driven Returns.” *Journal of Finance* 78(3): 1393–1447.

Greenwald, Daniel L., Matteo Leombroni, Hanno Lustig, and Stijn Van Nieuwerburgh. 2021. “Financial and Total Wealth Inequality with Declining Interest Rates.” NBER Working Paper 28613.

Jensen, Michael C. 1986. “Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers.” American Economic Review 76(2): 323–329.

Jiang, Hao, Dimitri Vayanos, and Lu Zheng. 2025. “Passive Investing and the Rise of Mega-Firms.” Review of Financial Studies 38(12): 3461–3496.

NVIDIA Corporation. 2026. “NVIDIA Announces Financial Results for Fourth Quarter and Fiscal 2026.” Press release, February 25, 2026. Santa Clara, CA.

Ofek, Eli, and Matthew Richardson. 2003. “DotCom Mania: The Rise and Fall of Internet Stock Prices.” Journal of Finance 58(3): 1113–1137.

Refinitiv/LSEG. 2025. Financial fundamentals panel for U.S.-listed software companies, 2016–2025. London Stock Exchange Group data service.

Salesforce, Inc. 2019. “Salesforce to Acquire Tableau.” Press release, June 10. San Francisco, CA.

Salesforce, Inc. 2020. “Salesforce Signs Definitive Agreement to Acquire Slack.” Press release, December 1. San Francisco, CA.

Salesforce, Inc. 2023. “Salesforce Announces Strong Fourth Quarter and Full Year Fiscal 2023 Results.” Press release, March 1. San Francisco, CA.

Samel, Hiram, and Aditya Chaudhry. 2026. “American Manufacturing and the Architecture of Capital.” MIT Initiative for New Manufacturing (INM) Working Paper.

Securities and Exchange Commission. 2003. “Conditions for Use of Non-GAAP Financial Measures.” Regulation G, Final Rule, 17 CFR Parts 228, 229, 244 and 249. January 22.

Shleifer, Andrei, and Robert W. Vishny. 1997. “The Limits of Arbitrage.” Journal of Finance 52(1): 35–55.

Thomson Reuters. 2024. Thomson Reuters Business Classification (TRBC) industry sector scheme. Refinitiv, London Stock Exchange Group.

Statistical Appendix

All tables based on author's analysis of 115 public SaaS companies, 2016–2025. Data from Refinitiv/LSEG.

Table A1. SBC as Percentage of Revenue: Distribution Over Time

Year	N	P10	P25	Median	P75	P90	Mean
2016	23	4.2%	5.9%	**9.3%**	13.9%	22.2%	11.1%
2017	26	7.1%	9.2%	**10.7%**	13.9%	21.4%	12.3%
2018	35	7.4%	9.2%	**11.6%**	17.2%	20.9%	14.5%
2019	43	6.1%	11.4%	**13.6%**	19.2%	23.1%	40.0%
2020	52	5.3%	12.0%	**15.5%**	20.5%	26.4%	38.8%
2021	88	5.8%	9.9%	**17.7%**	24.2%	38.7%	25.0%
2022	107	7.4%	11.6%	**19.4%**	26.6%	36.6%	36.8%
2023	107	6.6%	10.6%	**18.0%**	25.0%	31.8%	25.0%
2024	110	5.2%	10.8%	**15.4%**	22.2%	28.8%	21.5%
2025	112	4.8%	9.9%	**14.9%**	20.3%	26.1%	17.0%

Computed from TTM SBC / TTM Revenue at each firm's latest reporting quarter within the calendar year. Companies with SBC/Revenue ≤ 0 excluded from distribution.

The mean column is unwinsorized and is sensitive to a small number of low-revenue IPO-year observations with very high SBC/revenue ratios; the 2019 and 2020 means sit well above the corresponding ninetieth percentiles for this reason. The percentile columns are the reliable summary of the distribution.

Table A2. The Buyback Treadmill: SBC, Buybacks, and Free Cash Flow

Year	Total SBC	Buybacks	Reported FCF	FCF – SBC	BB/SBC
2016	$2.5B	$0.2B	$3.2B	$0.7B	8%
2017	$3.3B	$1.0B	$4.5B	$1.2B	31%
2018	$5.4B	$1.2B	$6.0B	$0.6B	21%
2019	$7.0B	$1.4B	$7.5B	$0.6B	20%
2020	$10.9B	$4.1B	$12.0B	$1.1B	38%
2021	$17.0B	$7.0B	$17.6B	$0.6B	41%
2022	$24.1B	$9.4B	$17.7B	–$6.4B	39%
2023	$26.8B	$18.6B	$33.7B	$6.9B	69%
2024	$28.3B	$25.1B	$46.3B	$18.0B	89%
2025	$27.8B	$28.8B	$48.3B	$20.5B	103%

Cumulative 2016–2025: SBC $153.1B; Buybacks $96.9B; Reported FCF $196.8B; FCF minus SBC $43.7B. In 2025, buybacks exceeded SBC at the sector level for the first time: the efficiency pivot made visible.

Table A3. Rule of 40 Pass Rates: Adjusted vs. GAAP

Year	N	Adj. Pass	GAAP Pass	Gap
2016	23	43%	22%	22 pp
2017	26	46%	12%	35 pp
2018	35	57%	23%	34 pp
2019	46	46%	20%	26 pp
2020	57	37%	19%	18 pp
2021	98	44%	27%	17 pp
2022	105	31%	11%	20 pp
2023	110	24%	5%	19 pp
2024	111	22%	6%	15 pp
2025	112	28%	10%	18 pp

Rule of 40 = YoY revenue growth + margin. Adjusted uses EBITDA + SBC add-back; GAAP uses operating profit margin. Gaps are computed from unrounded percentages.

Table A4. Selected SaaS Companies: Cumulative Financial Performance

Company	TTM 2025 Rev	2025 EV	Cum. Op	Cum. NI	Cum. SBC	Dilution
CrowdStrike	$4.6B	$132.3B	–$1.0B	–$1.0B	$3.4B	+12%
Snowflake	$4.4B	$92.1B	–$5.2B	–$4.7B	$5.6B	+21%
HubSpot	$3.0B	$23.2B	–$0.6B	–$0.7B	$2.2B	+52%
Datadog	$3.2B	$46.5B	–$0.1B	$+0.1B	$2.2B	+19%
Cloudflare	$2.0B	$74.4B	–$1.0B	–$1.1B	$1.3B	+19%
Twilio	$4.9B	$13.9B	–$3.1B	–$4.4B	$4.0B	+87%
Okta	$2.8B	$14.0B	–$2.5B	–$2.8B	$3.4B	+86%
MongoDB	$2.3B	$27.0B	–$1.7B	–$1.7B	$2.3B	+61%
ServiceNow	$12.7B	$187.1B	$+4.2B	$+3.7B	$10.1B	+29%
Workday	$9.2B	$60.2B	–$0.8B	–$0.6B	$10.0B	+35%
Salesforce	$40.3B	$245.0B	$+25.0B	$+17.9B	$21.7B	+40%
Fortinet	$6.6B	$62.3B	$+7.4B	$+5.6B	$1.9B	–11%

Cumulative figures over all available years (2016–2025). Dilution = split-adjusted cumulative share count growth from 2016 Q1. Revenue = TTM as of Q3 2025

Table A5. SaaS EV/Revenue Multiples Over Time

Year	N	P25	Median	P75	Mean
2016	23	3.6x	**5.8x**	8.2x	5.7x
2017	27	5.6x	**6.8x**	9.9x	7.4x
2018	35	5.7x	**8.6x**	10.7x	9.0x
2019	47	7.0x	**9.1x**	15.7x	13.4x
2020	57	9.8x	**17.5x**	34.0x	22.5x
2021	103	6.8x	**13.1x**	24.1x	19.2x
2022	109	2.2x	**5.0x**	8.3x	6.8x
2023	110	2.8x	**5.5x**	9.6x	7.3x

2024	113	3.0x	**4.8x**	10.7x	8.6x
2025	114	2.2x	**4.4x**	7.8x	7.0x

Enterprise value / trailing revenue. Companies with EV/Revenue > 100x excluded as outliers.

Table A6. Salesforce: Capital Structure Over Time (fiscal years ending January 31; FY2025 closed January 31, 2025)

Year	Revenue	GAAP Op	SBC	FCF	Buybacks	APIC	Ret. Earn.
FY2016	$6.7B	$0.1B	$0.6B	$1.0B	—	$5.7B	–$0.7B
FY2017	$8.4B	$0.2B	$0.8B	$1.7B	—	$8.0B	–$0.5B
FY2018	$10.5B	$0.5B	$1.0B	$2.2B	—	$9.8B	$0.6B
FY2019	$13.3B	$0.6B	$1.3B	$2.8B	—	$13.9B	$1.7B
FY2020	$17.1B	$0.5B	$1.8B	$3.7B	—	$32.1B	$1.9B
FY2021	$21.3B	$0.5B	$2.2B	$4.1B	—	$35.6B	$5.9B
FY2022	$26.5B	$0.6B	$2.8B	$5.3B	—	$50.9B	$7.4B
FY2023	$31.4B	$1.9B	$3.3B	$6.3B	$4.0B	$55.0B	$7.6B
FY2024	$34.9B	$6.0B	$2.8B	$9.5B	$7.6B	$59.8B	$11.7B
FY2025	$37.9B	$7.7B	$3.2B	$12.4B	$7.8B	$64.6B	$16.4B

APIC jumps: FY2020 (+$18.2B, Tableau acquisition), FY2022 (+$15.3B, Slack acquisition). Cumulative 2016–2025 (including FY2026 year-to-date data not shown in annual rows above): GAAP op $25.0B, SBC $21.7B, Buybacks $28.1B.

Technical Appendix

The SaaS Illusion: When Useful Finance Becomes Financial Engineering

Hiram Samel and Aditya Chaudhry
MIT Sloan School of Management
Working Paper — June 2026

This appendix documents the data construction, variable definitions, summary statistics, and full regression specifications underlying the empirical claims in the main text. It is designed to be self-contained: a reader with access to the replication archive can reproduce every reported figure. All monetary values are in USD unless otherwise noted. Figures derive from the corrected fundamentals panel finalized in June 2026 and are aligned with the companion manufacturing paper (Samel and Chaudhry 2026).

A. Data Construction

A1. Universe

The panel comprises 115 U.S. publicly traded software-as-a-service companies (114 with valid data for the balance-sheet series reported here), identified under the Thomson Reuters Business Classification (TRBC) 2012 “Software as a Service” code. The classification is deliberately conservative: diversified technology firms whose revenue is not predominantly subscription software (for example Microsoft, Oracle, and Alphabet) are excluded, so the panel captures pure-play SaaS rather than the broader software sector. Companies enter the panel at their initial public offering; the panel is therefore unbalanced by construction, with thirty firms public as of 2016 and the remaining eighty-five listing over the subsequent decade.

A2. Data source and field codes

All fundamentals are drawn from LSEG Refinitiv, pulled at fiscal-quarterly frequency. The table below lists the fields used and the specific Refinitiv field codes, including the corrected additional paid-in capital (APIC) field. A common pitfall is noted: the issued-and-paid common stock field returns par value (typically $1 per share), not paid-in capital; the correct APIC series uses the share-premium-including-option-reserve field.

Variable	Refinitiv field code	Statement
Retained earnings	TR.F.RetainedEarnTot	Balance sheet
Paid-in capital (APIC)	TR.F.COMSTOCKSHRPREMINCLOPTIONRSRV	Balance sheet
Net cash flow from operations	TR.F.NetCashFlowOp	Cash flow
Capital expenditures	TR.F.CAPEXTot	Cash flow

Variable	Refinitiv field code	Statement
Stock-based compensation	TR.F.SHRBASEDPAYMTCF	Cash flow (operating reconciliation)
Revenue	TR.F.TOTREVENUE	Income statement
Operating profit	TR.F.OPPROFBEFNONRECURINCEXPN	Income statement
Enterprise value	TR.F.EV	Valuation
Share buybacks	TR.F.STOCKCOMREPURCHRETIREDCF	Cash flow (financing)

Note: TR.F.ComStockIssuedPaid returns the par value of common stock, not paid-in capital. The correct APIC field is TR.F.COMSTOCKSHRPREMINCLOPTIONRSRV.

A3. Trailing-twelve-month construction

Flow variables (revenue, operating profit, stock-based compensation, free cash flow) are converted to trailing-twelve-month figures by summing the four most recent fiscal quarters. EV/revenue ratios computed against a single quarter would be inflated roughly fourfold and would carry seasonal distortion. Data is pulled from 2015 onward so that trailing figures are available beginning in 2016.

A4. Paid-in capital as a balance-sheet stock

Paid-in capital is reported throughout as the panel-wide sum of APIC measured at each firm's most recent reporting date. It is a balance-sheet stock accumulated since each firm's founding, not a flow of capital raised within the panel window. Of the roughly $322 billion total at the latest 2025 observation, approximately $42 billion was on the books at panel entry (incumbents' 2016 balances plus entrants' paid-in capital at IPO), and the remainder accumulated over the panel period, of which roughly $153 billion is capitalized stock-based compensation. This decomposition matters for interpretation: nearly half of the equity capital the sector has absorbed is workforce compensation paid in stock rather than cash raised from external investors.

A5. Winsorization

All derived ratio variables are winsorized at the first and ninety-ninth percentiles after replacing infinite values with missing. EV/revenue observations above 200× are dropped. Results are robust to alternative thresholds at 0.5/99.5 and 2.5/97.5. The cross-sectional regressions in Part C additionally require firms to have at least $10 million in trailing-twelve-month revenue at the 2021–Q4 peak, which excludes two non-economic micro-cap tickers (a sub-$1 million-revenue pink-sheet listing and a $0.13 million-revenue shell) whose stock-based-compensation and margin ratios would otherwise dominate the cross-section. The aggregate figures in Part B use the full panel and are unaffected.

B. Aggregate Findings

B1. Headline aggregates, 2016–2025

Metric	Value
SaaS firms in panel	115 (114 with balance-sheet data)
Cumulative reported free cash flow	$196.8B
Cumulative stock-based compensation	$153.1B
Free cash flow minus SBC	$43.7B
Cumulative buybacks	$96.9B
Total paid-in capital (APIC), 2025	$324.5B
Total retained earnings, 2025	–$73.7B
Firms carrying an accumulated deficit	84% (96 of 114)

Note. Paid-in capital is reported here as each firm's most recent APIC observation through the panel end (November 2025); the year-by-year table in B2 sums the latest observation within each calendar year, so its 2025 figure ($321.8B) differs slightly. The source note rounds the panel total to $322 billion, and the body reports it conservatively as more than $320 billion.

B2. The gap: paid-in capital versus retained earnings

The sector has absorbed paid-in capital that vastly exceeds cumulative earnings; the gap widens in every year of the panel. Figures are the panel-wide sum of the latest observation per firm within each year.

Year	APIC ($B)	Ret. earnings ($B)	Gap ($B)	% deficit
2016	20.4	–5.5	25.9	80%
2017	26.5	–8.4	34.9	80%
2018	45.7	–13.3	59.0	82%
2019	81.4	–17.5	98.8	82%
2020	117.5	–26.5	144.0	77%
2021	207.8	–46.7	254.5	87%
2022	232.9	–68.7	301.5	86%
2023	260.5	–79.1	339.6	88%
2024	288.4	–72.2	360.7	83%
2025	321.8	–73.7	395.5	84%

B3. Free cash flow versus free cash flow minus SBC

Reported free cash flow obscures the cost of equity compensation. Cumulative reported FCF over the panel is $196.8 billion; after subtracting stock-based compensation, cumulative cash generation is $43.7 billion. The sector's FCF minus SBC was approximately zero or negative through 2022 and turned materially positive only in

2023–2025, as cash generation grew into and past a roughly flat SBC bill rather than as a result of SBC restraint.

Year	FCF ($B)	SBC ($B)	FCF–SBC ($B)
2016	3.2	2.5	0.7
2017	4.5	3.3	1.2
2018	6.0	5.4	0.6
2019	7.5	7.0	0.6
2020	12.0	10.9	1.1
2021	17.6	17.0	0.6
2022	17.7	24.1	–6.4
2023	33.7	26.8	6.9
2024	46.3	28.3	18.0
2025	48.3	27.8	20.5

B4. Valuation multiples

Median EV/revenue across the panel compressed from double-digit levels at the 2020–2021 peak to mid-single digits by 2025: 17.5× in 2020, 13.1× in 2021, 4.8× in 2024, and 4.4× in 2025. Full annual percentiles are reported in Table A5 of the main text.

C. Regression Specifications (Section VI-b)

The two regressions reported in Section VI-b of the main text are run on the SaaS panel described above. The first tests whether stock-based-compensation intensity explains the cross-section of multiple compression during the 2021–Q4 to 2022–Q4 repricing; the second tests whether firms that built the most SBC intensity at the peak subsequently retreated from it most. Both ladders and the quintile characterization below were regenerated from master_panel.parquet for this version. The first regression finds that SBC intensity does not explain the cross-section of compression once revenue growth is controlled; the second finds that peak SBC intensity predicts subsequent retreat at a coefficient that is stable and significant across the full control ladder. Each table reports the reproduced coefficients, standard errors, sample sizes, and fit, with the estimation sample and standard-error treatment noted beneath it.

C1. Repricing regression (2021–Q4 peak to 2022–Q4 trough)

Over the cleanest window of the rate shock, median SaaS EV/revenue fell from roughly 13.0× to 5.0×, a compression of about sixty percent. This regression tests whether stock-based-compensation intensity explains the cross-section of that compression. The

dependent variable is the firm-level percentage change in EV/revenue from the 2021–Q4 peak to the 2022–Q4 trough (trough divided by peak, minus one), which places every firm's compression on the same footing rather than letting the highest-multiple names dominate an absolute change. The key regressor is SBC intensity (SBC/revenue) at the 2021–Q4 peak. SBC intensity does not explain the cross-section of compression: its coefficient is statistically insignificant on its own ($p = 0.30$) and remains insignificant once revenue growth and operating margin at peak are added ($p = 0.34$). Revenue growth carries the relationship (coefficient –0.41, $p < 0.001$) and starting margin enters positively ($p = 0.002$). For the event of repricing itself, the duration channel subsumes the permission-structure channel; the cross-section of who fell most is explained by growth, not by SBC intensity. This is reported in the main text precisely because it is the result a referee running the same regression would reach, and because the paper's claim concerns who built the largest architecture over the prior decade, not who fell most during the unwinding (see C2).

Specification ladder (coefficient on peak SBC intensity):

Model	SBC coef. (SE)	p-value	N; R^2
(1) SBC intensity only	–0.21 (0.20)	0.30	85; 0.02
(2) + revenue growth + operating margin (at peak)	+0.27 (0.29)	0.34	82; 0.34

Reproduction note. Regenerated from master_panel.parquet on the SaaS cross-section present at both the 2021–Q4 peak and the 2022–Q4 trough, restricted to firms with TTM revenue of at least $10 million at the peak (this removes two non-economic micro-cap tickers; see A5). Ratio and change variables are winsorized at 1/99. Standard errors are heteroskedasticity-robust (HC1); with one observation per firm, firm-level clustering reduces to the same. Model 1 N = 85; model 2 N = 82.

C2. SBC-retreat regression (2021–Q4 to 2025–Q3)

This is the paper's central cross-sectional test. The dependent variable is the change in SBC intensity from the 2021–Q4 peak through 2025–Q3. The key regressor is SBC intensity at the 2021–Q4 peak. The coefficient is stable at approximately –0.65 across a ladder of controls: revenue growth, starting operating margin, size, leverage, cash position, stock return, and subsequent growth. The interpretation: a firm with ten percentage points higher SBC intensity at peak cut SBC by roughly six and a half percentage points more of revenue through the correction, independent of every control. The coefficient survives all specifications at p below 0.001.

Specification ladder (coefficient on peak SBC intensity):

Model (controls added cumulatively)	SBC coef. (SE)	p-value	R^2 (N)
(1) SBC intensity only	–0.65 (0.10)	< 0.001	0.59 (84)
(2) + growth + starting margin	–0.68 (0.09)	< 0.001	0.57 (81)

Model (controls added cumulatively)	SBC coef. (SE)	p-value	R^2 (N)
(3) + size	–0.66 (0.10)	< 0.001	0.57 (81)
(4) + leverage + cash	–0.67 (0.09)	< 0.001	0.71 (71)
(5) + stock return	–0.64 (0.08)	< 0.001	0.72 (70)
(6) + subsequent growth	–0.67 (0.08)	< 0.001	0.73 (70)

Quintile characterization (firms sorted by SBC intensity at the 2021 peak):

Quintile	SBC/rev	Op margin	Growth	Leverage	Cash/rev	Stock return
Q1 (low)	4.8%	–2%	42%	0.04	0.90	–0.57
Q5 (high)	39.3%	–66%	58%	0.05	2.47	–0.58

The quintile table addresses the natural objection that high-SBC firms cut under financial distress. They did not: the highest-SBC quintile carried nearly three times the cash-to-revenue ratio of the lowest (2.47 versus 0.90), essentially zero leverage in both, and peak-to-trough stock declines no deeper than the rest (–0.58 versus –0.57). Firms with more than two years of revenue in cash chose to retreat from a practice they no longer had permission to maintain. One control coefficient is worth flagging for the full table: leverage carries a significant negative coefficient on SBC cuts (roughly –0.20 to –0.27, $p < 0.01$), which the quintile characterization contextualizes rather than contradicts.

Reproduction note. Regenerated from master_panel.parquet on the SaaS cross-section present at both the 2021–Q4 peak and 2025–Q3, restricted to firms with TTM revenue of at least $10 million at the peak (see A5). Ratio and change variables are winsorized at 1/99. Standard errors are heteroskedasticity-robust (HC1); with one observation per firm, firm-level clustering reduces to the same. The sample falls from N = 84 to N = 70 across the ladder as the leverage, cash, return, and subsequent-growth controls require complete data on each; N is reported per model. Q1 and Q5 are equal-count quintiles of peak SBC intensity (seventeen firms each).

C3. Variable definitions

Variable	Construction
SBC intensity	TTM stock-based compensation / TTM revenue
EV/revenue	Enterprise value / TTM revenue
Multiple compression	Percentage change in EV/revenue, peak to trough (trough/peak – 1)
Revenue growth (YoY)	TTM revenue_t / TTM revenue_{t–4} – 1
Operating margin	TTM operating profit / TTM revenue (GAAP, SBC included)
Leverage	Total debt / (total debt + market capitalization)
Cash position	Cash and short-term investments / TTM revenue

Definitions confirmed against 14_saas_vib_regressions.py and constructed in master_panel.parquet. Size is log TTM revenue at the peak (results are unchanged using log market capitalization). Stock return is a price return: in the C2 ladder it is the peak-to-2025–Q3 holding-period return, and in the quintile table the reported figure is the peak-to-trough drawdown (2021–Q4 to 2022–Q4). Subsequent growth is the change in TTM revenue from the peak to 2025–Q3. All are cited to master_panel.parquet rather than saas_paper_panel.parquet.

D. Replication

The replication archive contains the data-construction scripts, the SaaS panel file, and the analysis scripts that produce the figures and regressions reported here. Standard dependencies: pandas, numpy, statsmodels, pyarrow.

The replication archive contains 12_saas_paper_analysis.py, which builds saas_paper_panel.parquet and the aggregate figures in Part B, and 14_saas_vib_regressions.py, which constructs the Section VI-b variables and runs the C1 repricing regression, the C2 retreat ladder, and the quintile characterization against master_panel.parquet, reproducing the coefficients reported here.