WINDSOR DRAKE

WINDSOR DRAKE RESEARCH · 2026

The Rule of 40 Premium

Windsor Drake is a sell-side M&A advisory firm focused on software, financial technology, artificial intelligence, and other technology-enabled businesses in the lower middle market. The firm publishes research on valuation, transaction structure, and deal activity for founders, boards, and investors.

About this benchmark

This is a cross-sector valuation study. Rather than profile a single industry, it compares how one widely used measure of company quality, the Rule of 40, translates into enterprise value across four software-adjacent verticals. It is intended to be read alongside Windsor Drake's sector-specific valuation work, which examines each market in depth.

CONTENTS

Contents

THE FINDING

The finding, in one chart

The Rule of 40 is the most cited shorthand in software valuation: a company whose revenue growth rate plus its profit margin clears 40 is considered healthy, and one that falls short is not.[1] It is a useful discipline. It is also, as a guide to what a company is worth, badly incomplete, because the price the market pays for clearing the bar is not a constant. It depends almost entirely on what the company is called.

EXHIBIT 1

Same quality, different price: EV/revenue against Rule of 40, by vertical

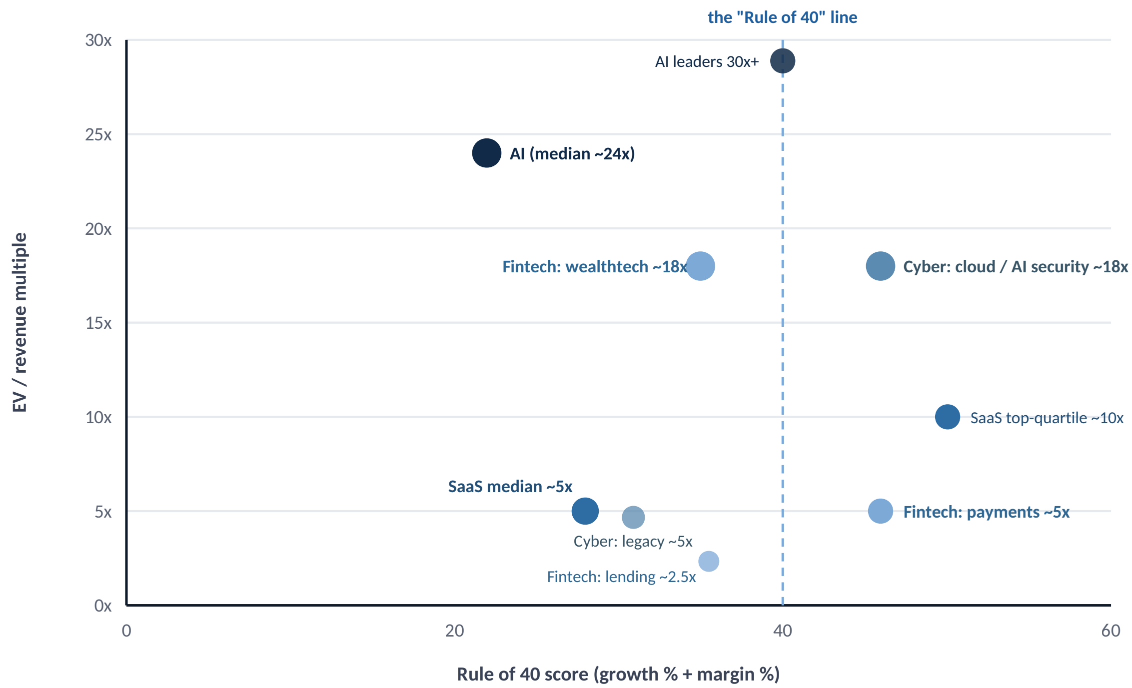

Source: Windsor Drake analysis of public market data and disclosed transactions, 2025–2026, drawing on Meritech, Aventis Advisors, and sector comps.[1,2,3,4] Positions are representative central tendencies; ranges within each category are wide.

Read the chart by looking down a vertical slice. At a Rule of 40 score of roughly 45, a payments company trades near five times revenue, a cloud-security company near eighteen, and the two are doing equally well by the measure that supposedly governs software value. Now look at the far left. An AI company growing fast but unprofitable, sitting well below the 40 line, still trades above two dozen times revenue, four to five times the multiple of the payments company that comfortably passes the test.[2,3,4] The rule says one company is healthy and the other is not. The market says the unhealthy one is worth far more.

~4x

An AI company that fails the Rule of 40 still tends to out-multiple a payments company that passes it, on revenue.

This is the premium the report is named for. There is a real reward for clearing the Rule of 40, but it sits on top of a much larger effect: a base multiple set by the company's category and the architecture of its revenue. The base dominates the bonus. A founder who optimizes the score while ignoring the category is tuning the smaller of the two dials. The pages that follow quantify both, vertical by vertical, and set out what a seller should do with the result.

SUMMARY

Key findings

1. **The Rule of 40 is a quality test, not a price.** Clearing it is associated with materially higher multiples within a vertical, but the size of that reward, and the base it sits on, differ so much across verticals that the score alone explains only a fraction of value.[1,2]

2. **Category sets the base multiple before quality adjusts it.** Median EV/revenue runs around 5x in broad SaaS, near 6x in public fintech, around 8x in cyber, and roughly 24x in AI. A company inherits its starting multiple from what it is called.[2,3,4]

3. **Revenue architecture matters as much as the vertical label.** Inside fintech, compounding models such as wealthtech trade near 18x while transaction-processing models such as payments and lending trade at 2.5x to 6x. The question is whether revenue deepens with the customer or merely scales with volume.[3]

4. **The within-vertical premium is real and roughly linear.** In software, each ten-point gain in Rule of 40 is associated with about a 1.1x uplift in EV/revenue, and companies above the line trade at roughly double those below it.[1,2]

5. **In AI, narrative and defensibility have displaced the rule.** AI companies trade at 10x to 50x revenue with a median near 24x, frequently while failing the Rule of 40 outright, because the market is pricing future defensibility rather than current growth-plus-margin.[4]

6. **Strategic M&A pays a second premium on top.** In cyber, public names trade near 8x while disclosed acquisitions average closer to 19x, with strategic deals running higher still. Where you sell, and to whom, can move the number more than the fundamentals do.[2]

7 **The practical implication is positioning, not just performance.** For many lower-middle-market companies, how the business is categorized and which buyer pool it is taken to will determine the multiple more than another few points of Rule of 40 ever could.

PART I

What the Rule of 40 measures, and what it hides

Before the rule can be useful, it has to be understood for what it is: a single number that compresses two of a company's most important traits into one, and in doing so throws away most of what determines value.

The Rule of 40 holds that a software company's revenue growth rate and its profit margin should sum to at least 40. A business growing 30 percent with a 10 percent margin passes. So does one growing 15 percent with a 25 percent margin, and one growing 50 percent while burning 10 points of margin. The appeal is that it captures the central trade-off in software, the exchange of profitability for growth, in a way a board can hold in its head.[1] The measure is most often computed on forward revenue growth and free cash flow margin, though definitions vary, and that variation is the first thing to keep in mind when comparing one source to another.[1]

As a management discipline the rule earns its place. As a valuation tool it carries three blind spots that this report is built around.

It treats a dollar of growth as identical everywhere

The rule weighs growth and margin equally and says nothing about the quality of the revenue underneath. A dollar of revenue that compounds, deepening as a customer stays and expands, is worth far more than a dollar that has to be re-won every period, yet both count the same in the score. As Part IV shows, this single distinction explains much of the gap between a wealthtech platform and a payments processor that post the same Rule of 40.[3]

It is blind to category

Two companies with identical scores can sit in markets that the public market values on entirely different scales. The rule cannot see that one is called a payments company and the other an AI company, and that this label, more than the score, sets the multiple. A framework that ignores the single largest driver of the multiple is not a pricing tool.[2,4]

It assumes the present predicts the price

The Rule of 40 is a snapshot of current performance. Increasingly, the market is not paying for the present. In AI especially, buyers underwrite a future state, defensibility, data advantage, a position in an emerging category, and will pay multiples that current growth-plus-margin cannot justify on any historical basis.[4] A backward-looking test cannot capture a forward-looking price.

> *The Rule of 40 answers a real question well. It just is not the question of what the company is worth.*

How a heuristic became a law

The rule's authority is worth understanding, because it explains why founders trust it past the point where it helps. The Rule of 40 spread because it solved a real problem at a specific moment. As the software market matured and pure growth multiples came under pressure, investors needed a single, comparable way to judge whether a company was trading growth for profit on acceptable terms. A two-variable sum that any operator could compute filled that need, and it propagated through board decks, investor updates, and banker pitches until it acquired the status of a law rather than a heuristic.

That history left two marks. First, the rule is calibrated to a particular kind of company, the recurring-revenue software business it was built to describe, and it travels poorly to businesses whose economics differ, which is most of fintech, much of cyber, and nearly all of AI. Second, the market that produced the rule has already moved past it. The years to 2021 priced growth almost regardless of margin; the years that followed re-rated toward profitability and made the rule briefly central; and the current period prices AI on a narrative the rule cannot read.[2,4] A measure that was central for one phase of the market is now one input among several, and the founders most exposed are those still treating it as the whole answer.

None of this makes the rule useless. Within a single category, among companies the market values on the same scale and with similar revenue quality, the Rule of 40 tracks valuation closely, and the next sections show exactly how closely. The error is in treating it as a law that holds across categories. It does not, and the cost of that error, paid in mispriced expectations and misdirected effort, is what the rest of this report quantifies.

PART II

The benchmark, vertical by vertical

Each vertical values software on its own scale. The four benchmarks below establish where each market clears today and how quality moves the number inside it. Read together, they make the cross-vertical premium in Part III impossible to miss.

Software, the base case

Public software is the reference point against which the others are read. By the end of 2025 the broad public software universe traded at roughly five times revenue, well off the 2021 peak, as the market shifted from rewarding growth at any cost to rewarding durable, profitable growth.[2] The Rule of 40 is most predictive here. The median public software company scores about 28, only around one in five clears 40, and the companies above the line trade at roughly double the multiple of those below it, with each ten-point gain in score associated with about a 1.1x uplift in EV/revenue.[1,2]

EXHIBIT 2

Software: where the market clears, and what quality adds

COHORT	TYPICAL EV/REVENUE	RULE OF 40 CONTEXT
Broad public software (median)	~5.1x	Median score ~28; ~20% clear 40
Above the Rule of 40 line	~2x the below-line median	Premium roughly linear: ~1.1x per 10 points
Below the line	Compressed	Harder fundraising and exit terms
Mature, profitable software	~26x EV/EBITDA	EBITDA becomes the operative metric

Source: Aventis Advisors and Meritech benchmarks, Q4 2025; Windsor Drake analysis.[1,2]

What moves a software company within this band is straightforward and well understood: net revenue retention, gross margin, and the durability of growth. Because the market values these companies on a shared scale, the Rule of 40 does much of its honest work here, and a seller can reasonably expect the score to track the multiple. The trap, addressed later, is assuming that this tight relationship inside software also holds when the company is compared to a fintech, a cyber, or an AI business that the market prices on a different scale entirely.

Fintech, where revenue architecture splits the field

Fintech does not trade as one market. Public fintech averages near six times revenue, but the dispersion underneath is the widest of any vertical here, and it is not random.[3] Models whose revenue compounds with the customer relationship command multiples several times those of models whose revenue scales with transaction volume or credit issued. Wealthtech and software-led enterprise fintech trade in the high teens; payments sit at four to six times; lending platforms cluster near two and a half. Two companies with the same Rule of 40 can sit at either end of that range depending only on how their revenue is built.[3]

EXHIBIT 3

Fintech: multiples track revenue architecture, not just the score

SUBSECTOR	TYPICAL EV/REVENUE	REVENUE CHARACTER
Wealthtech	~18–25x	Compounds; usage deepens with the relationship
SMB / enterprise fintech software	~15–17x	Compounds; recurring software economics
Payments	~4–6x	Processes; revenue scales with volume
Lending platforms	~2.5x	Processes; revenue tied to credit and risk
Public fintech (overall avg)	~5.9x	Blended; M&A has averaged ~4.4x in 2025

Source: Finro, Multiples.vc, and Windsor Drake fintech valuation research, 2025–2026.[3] Figures are central tendencies; individual companies vary widely.

The practical consequence for a fintech founder is that repositioning can be worth more than operating improvement. A company that processes payments but also sells a genuine software layer faces a choice in how it presents itself, and that choice can be worth several turns of revenue. The market is not being irrational when it pays more for compounding revenue; it is pricing the difference between a business that grows by deepening relationships and one that grows only by handling more volume at thinner spreads. A seller who can credibly demonstrate the former should never accept being priced as the latter, and the gap is large enough that the framing question deserves as much attention as the financials.

Cybersecurity, where the subsector is the story

Cyber's headline public multiple of roughly eight times revenue conceals a split as sharp as fintech's. Cloud and AI-native security command multiples in the high teens to low twenties; application and data security occupy a middle band; legacy hardware and services trade at four to six times.[2,5] Cyber also shows the clearest strategic-M&A premium of any vertical here: while public names trade near eight times, disclosed acquisitions have averaged closer to nineteen, and competitive strategic processes have cleared well above twenty.[5] The AI-security niche has been the most active, with strategic buyers paying up for model-security capabilities.[6]

EXHIBIT 4

Cybersecurity: a wide public range and a large M&A premium

SEGMENT	PUBLIC EV/REVENUE	NOTE
Cloud / AI-native security	~18–22x	The "haves"
Application security	~13–15x	High-growth
Data security / GRC	~6–8x	Middle band
Legacy hardware / services	~4–6x	The "have-nots"
Public median / M&A average	~7.8x / ~18.8x	Strategic deals price higher still

Source: Finro cybersecurity valuation data, Multiples.vc, and Windsor Drake cyber valuation research, Q4 2025–Q2 2026.[5,6]

Cyber also illustrates the buyer-pool effect better than any other vertical, which is why the M&A premium is so large. A security platform that would be merely a good public company can be a strategic necessity inside a larger vendor's suite, closing a gap a competitor might otherwise fill. That strategic value is not visible in the standalone financials and does not show up in the Rule of 40, yet it is what drives the difference between the public median near eight times and disclosed deals near nineteen.[5] For a cyber founder, the lesson is that the exit multiple is a function of how many strategic acquirers view the asset as must-own, and that this can be cultivated long before a process begins.

AI, where the rule stops applying

AI is the outlier that proves the report's point. AI companies trade at ten to fifty times revenue with a median near twenty-four, often while failing the Rule of 40 outright, because they are unprofitable by design and the market is underwriting something other than current performance.[4] Anthropic's 2025 financing implied roughly thirty-seven times revenue; Databricks has carried close to twenty-eight times ARR.[4] The valuation driver here has shifted from growth toward defensibility, proprietary data, and category position, which a growth-plus-margin score cannot measure.[4]

EXHIBIT 5

AI: multiples set by stage, defensibility, and narrative

COHORT	TYPICAL REVENUE MULTIPLE	PRIMARY DRIVER
AI overall (median)	~24x	Defensibility > growth rate
Seed-stage AI	~10–25x	Team, IP, category
Late-stage / public AI	~25–35x	Run-rate and moat
Reference points	Anthropic ~37x; Databricks ~28x ARR	Narrative-led

Source: Aventis Advisors AI valuation research, Value Add VC, and stage-level surveys, 2025–2026.[4] Ranges are wide and move quickly.

The AI benchmark should be read with the most caution of the four, for two reasons. It is the least settled, with multiples that move quickly and a sample dominated by private financings whose terms carry structure that headline numbers omit. And it is the most likely to re-rate, because so much of the base rests on expectations rather than realized economics. The point for this report is not whether the AI base is correct. It is that the AI base exists at all, sitting four to six times above the software base for companies that, on the Rule of 40, often look worse. That gap is the clearest possible demonstration that category, not score, sets the price.

PART III

The premium, quantified

With each vertical's scale established, the cross-cut becomes arithmetic. Take a company sitting exactly at the Rule of 40, the same quality by the standard measure, and ask what it is worth in each market. The answer is the premium this report set out to size.

The decomposition is simple and, once seen, hard to unsee. A company's multiple is the product of two things: a base multiple inherited from its category and revenue architecture, and a quality adjustment earned through its Rule of 40 score. The base is large and varies enormously across verticals. The adjustment is real but modest, on the order of doubling between the bottom and the top of the range within a vertical.[1,2] Because the base varies by roughly five times across these four markets while the quality adjustment varies by about two, the base does most of the work.

EXHIBIT 6

What a company at the Rule of 40 is worth, by vertical

VERTICAL	BASE MULTIPLE	AT RULE OF 40 (EST.)	WHAT SETS THE BASE
Payments (fintech)	~5x	~5–6x	Volume-linked, non-compounding revenue
Broad SaaS	~5x	~7–9x	Recurring software economics
Compounding fintech	~15x	~16–20x	Revenue deepens with the customer
Cloud / AI security	~16x	~18–22x	Mission-critical, high retention, strategic demand
AI	~24x	~25–35x	Defensibility and category narrative

Source: Windsor Drake estimates derived from the benchmarks in Part II. "At Rule of 40" applies the observed within-software quality premium to each vertical's base multiple; figures are illustrative analytical constructs, not quotes, and ranges are wide.

The table makes the central claim precise. A payments company that clears the Rule of 40 lands around five to six times revenue. An AI company that clears the same bar lands four to six times higher, and even an AI company well below the bar typically sits above the payments company that cleared it. The score the founder worked to achieve moved the multiple by a point or two. The category the company happens to occupy moved it by a factor of five.

The base multiple does roughly five times the work of the quality premium. Founders spend most of their attention on the smaller dial.

EXHIBIT 7

The four verticals at a glance

VERTICAL	BASE MULTIPLE	INTERNAL RANGE	RULE OF 40 FIT	DOMINANT VALUE DRIVER
Software / SaaS	~5x	~3–10x+	Strong	Recurring growth and margin
Fintech	~6x	~2.5–25x	Weak	Revenue architecture (compounds vs processes)
Cybersecurity	~8x	~4–22x	Moderate	Subsector and strategic demand
AI	~24x	~10–50x	Negligible	Defensibility and narrative

Source: Windsor Drake synthesis of the benchmarks in Part II. "Rule of 40 fit" describes how well the score tracks the multiple inside each vertical. Figures are central tendencies as of mid-2026.

Three head-to-head reads

The abstraction becomes concrete when two companies are placed side by side. Each pair below holds the Rule of 40 roughly constant, or even tilts it against the higher-valued company, and lets the category and revenue architecture do the talking.

EXHIBIT 8

The premium, three ways

COMPARISON	LOWER MULTIPLE	HIGHER MULTIPLE
Score held constant (both clear ~40)	Payments processor, ~5x revenue	Cloud-security platform, ~18x revenue
Score tilted against the winner	Payments processor at Rule of 40, ~5x	AI company below Rule of 40, ~24x
Same vertical, different architecture	Lending platform, ~2.5x revenue	Wealthtech platform, ~18–25x revenue

Source: Windsor Drake, from the benchmarks in Part II. Archetypes, not specific companies; illustrative of the structural gaps.

The first pair is the cleanest statement of the finding: two companies, equally healthy by the rule, three to four turns apart because one sells mission-critical, high-retention software and the other processes volume. The second pair is the provocative one, an AI company that fails the test trading at five times the multiple of a payments company that passes it, because the market is buying the AI company's future and the payments company's present. The third pair stays inside a single vertical and isolates revenue architecture alone, a lending business and a wealthtech business that may post similar scores and live in the same broad market, yet trade a full order of magnitude apart on revenue.[3,4,5]

None of these gaps is irrational. Each reflects a real difference in the durability, defensibility, or expandability of the revenue. The error the rule encourages is to look only at the score, see two healthy companies, and expect two similar outcomes. The market never promised that, and the seller who understands why will set expectations, and structure a process, very differently from the one who does not.

The same company, priced four ways

To make the stakes unmistakable, take one set of financials and read it through each vertical's lens. Consider a company with $20 million of revenue growing 30 percent with a 10 percent free cash flow margin. Its Rule of 40 score is exactly 40. It passes. Now ask what it is worth, applying the at-Rule-of-40 multiples from Exhibit 6.

EXHIBIT 9

One company at the Rule of 40, valued in four markets

IF THE MARKET PRICES IT AS…	MULTIPLE APPLIED	IMPLIED ENTERPRISE VALUE
A payments business	~5–6x	~$100–120M
A broad SaaS business	~7–9x	~$140–180M
A compounding fintech	~16–20x	~$320–400M
A cloud / AI security platform	~18–22x	~$360–440M
An AI company	~25–35x	~$500–700M

Source: Windsor Drake illustration applying Exhibit 6 multiples to identical financials ($20M revenue, 30% growth, 10% FCF margin; Rule of 40 = 40). For demonstration only; actual outcomes depend on the company's real category, revenue architecture, and buyer pool.

The same income statement supports a valuation of roughly $100 million or roughly $700 million, a seven-fold spread, with the Rule of 40 held perfectly constant the entire time. Every dollar of that spread is decided by the four drivers in Part V, not by the score. This is the practical meaning of the report's title: the premium for clearing the Rule of 40 is real, but it is dwarfed by the premium for being the right kind of company in the right market. A founder who treats valuation as a function of the score is solving for a variable that, in this example, does not move at all.

The illustration also clarifies what positioning can and cannot do. A company cannot simply declare itself an AI business and collect the AI multiple; the market prices the economics, not the label. But where a business genuinely possesses the characteristics of a higher-multiple category, compounding revenue, real defensibility, a credible position in a contested market, allowing it to be priced as the lower category is a self-inflicted discount. The work of a good process is to ensure the company is valued as what it actually is, which is often higher than the default framing a casual observer would assign.

The mirror image is just as instructive. Take a second company with the same $20 million of revenue, but growing 50 percent while burning 30 points of margin, for a Rule of 40 score of 20. By the rule, it fails outright. Priced as an AI company at roughly 24 times revenue, it is still worth on the order of $480 million, comfortably above the $100 to $120 million of the payments business that scored a perfect 40. The failing company is worth four to five times the passing one. No amount of operating discipline at the payments company closes that gap, because the gap was never about discipline. It was about category.

EXHIBIT 9B

When failing beats passing: two companies, same revenue

COMPANY	RULE OF 40	PRICED AS	IMPLIED EV
Disciplined payments business (30% growth, +10% margin)	40 (passes)	Payments	~$100–120M
Unprofitable AI business (50% growth, −30% margin)	20 (fails)	AI	~$480M

Source: Windsor Drake illustration; identical $20M revenue. Demonstrates that category can outweigh the score entirely. Not a quote for any company.

Two cautions belong with these exhibits. First, the "at Rule of 40" figures are analytical constructs built from the central tendencies in Part II, not transaction quotes; any individual company can sit far from them. Second, the premium is not static. The AI base in particular reflects a market pricing a future that may or may not arrive, and a correction there would compress the most extreme gaps shown here. The structure of the finding, that category dominates score, is durable. The exact magnitudes are a snapshot of mid-2026.

PART IV

Where the premium is moving

A benchmark is a photograph of a market in motion. The relationships in this report are stable in structure but shifting in magnitude, and a seller timing a process needs to know which way each base is trending.

Software: the profitability re-rating is now priced in

The large move in software has already happened. The market spent 2023 to 2025 re-rating from growth toward durable, profitable growth, margins expanded across the public set, and profitability became a structural rather than tactical trait of how these companies are valued.[2] The implication for sellers is that the easy multiple gains from simply turning profitable are largely captured. The base is unlikely to expand much from here on profitability alone; differentiation now comes from growth quality and defensibility, the things that were always the real drivers.

Fintech: the software premium is widening the spread

Within fintech, the gap between compounding and processing models has widened, not narrowed, as the market has grown more discriminating about revenue quality.[3] Software-led fintech is increasingly priced alongside SaaS, while volume-linked models are increasingly priced as the financial businesses they resemble. For founders, this raises the stakes on positioning: the cost of being categorized as a processor when a software framing is available has gone up, and the reward for genuinely building compounding revenue has gone up with it.

Cyber: consolidation is inflating the strategic premium

Cyber's M&A premium is being driven by an active consolidation cycle, as platform vendors absorb point solutions to close suite gaps, with the AI-security niche the most contested.[5,6] As long as a handful of large strategics are competing to own emerging categories, the spread between public multiples and strategic deal multiples will stay wide, and well-positioned assets in contested subsectors will continue to clear far above their standalone value. The risk is concentration: a premium that depends on a small number of acquirers is vulnerable if any of them steps back.

AI: the base is the question the whole market is asking

The AI base is the least stable number in this report and the one whose movement matters most. It rests on a collective bet on the size and defensibility of the opportunity, paid in advance, and the central debate in the market is whether that bet is justified.[4] Two paths are worth holding in mind. In one, AI revenue proves as durable and defensible as the multiples imply, and the base holds or rises. In the other, defensibility disappoints, and the base compresses toward software, collapsing the most extreme gaps in Exhibit 1. Windsor Drake takes no position on which arrives. The point for a seller is that an AI-adjacent company carrying a narrative premium is holding the most valuable and the most fragile base in the market, and timing matters more there than anywhere else.

Every base in this report is moving. The software re-rating is spent, the fintech spread is widening, the cyber premium depends on a few buyers, and the AI base is the open question of the cycle.

EXHIBIT 10

What to watch, by vertical

VERTICAL	DIRECTION OF THE BASE	THE SIGNAL THAT MATTERS
Software / SaaS	Stable	Profitability re-rating largely complete; watch growth quality and net retention for further differentiation
Fintech	Spread widening	The gap between software-led and volume-led models; the cost of being framed as a processor is rising
Cybersecurity	Premium elevated	Pace of strategic consolidation and the number of active acquirers in contested subsectors
AI	Most uncertain	Whether realized defensibility and retention justify the base; a re-rating here moves everything

Source: Windsor Drake analysis, mid-2026. Directions are qualitative judgments, not forecasts.

The common thread is that the durable sources of value, compounding revenue and genuine defensibility, are being rewarded more sharply over time, while premiums that rest on profitability alone or on narrative alone are either spent or at risk. A seller positioning for the next twelve to eighteen months should build toward the durable drivers and treat the fragile premiums as timing opportunities rather than permanent features.

PART V

Why identical quality is priced differently

If two companies post the same Rule of 40 and trade at multiples five times apart, the difference is not in the score. It is in four things the score cannot see. Each one moves the base multiple, and together they explain the dispersion in Exhibit 1.

One. Revenue that compounds versus revenue that processes

The single largest discriminator inside a vertical is whether a unit of revenue deepens over time or has to be regenerated. Compounding revenue, the kind a customer expands as it stays, carries higher gross margins, higher retention, and a longer tail, and the market capitalizes that tail into a higher multiple. Processing revenue, tied to transaction volume or credit issued, grows only by doing more, often at the cost of more risk or more expense. This is why wealthtech and software-led fintech trade in the high teens while payments and lending trade low, even at the same Rule of 40.[3] The score counts both dollars equally; the market does not.

The mechanism is net revenue retention. A compounding business keeps and expands its customers, so a cohort acquired this year is worth more next year without additional sales spend, and the market capitalizes that expanding annuity into a higher multiple. A processing business has to win the next dollar more or less from scratch, often by taking on more transaction risk or thinner spreads, so the same headline growth rate is worth less because it is more expensive and less certain to repeat. When two companies post the same Rule of 40 but one retains and expands revenue while the other churns and replaces it, the rule sees a tie and the market sees a rout.

Two. Defensibility and the cost of displacement

The market pays for revenue that is hard to take away. Switching costs, data advantages, network effects, and mission-critical positioning all raise the multiple because they lower the risk that growth reverses. This is most visible in cyber, where mission-critical, high-retention cloud-security platforms trade at three to four times the multiple of commoditized hardware with similar financials.[5] In AI, defensibility has explicitly overtaken growth as the primary valuation driver in many investor frameworks, which is why a moderately growing company with proprietary data can out-multiply a faster-growing one without it.[4]

Three. Category narrative and the size of the prize

Multiples embed a story about the market, not just the company. A business in a category the market believes is large and early will carry a higher base than an identical business in a category seen as mature. AI is the clearest case: the base multiple reflects a collective bet on the size and durability of the opportunity, paid in advance of the cash flows that would justify it.[4] This is also the most fragile of the four drivers, because narrative can re-rate quickly, and a company that lives on category premium alone is exposed when the story cools.

Four. The buyer pool and the strategic premium

Price is set by who is bidding. A company with many credible strategic acquirers, each for whom it would be mission-critical, will clear a higher multiple than one with a thin or purely financial buyer pool, regardless of score. Cyber again shows the effect at its starkest: the same assets that trade near eight times in public markets change hands near nineteen times in M&A, because a strategic buyer is pricing the asset's value inside its own platform, not as a standalone.[5] Where a company sells, and how competitive the process, can move the outcome more than a year of operating improvement.

Four forces the score cannot see, compounding, defensibility, narrative, and the buyer pool, set the base. The Rule of 40 only fine-tunes what they decide.

Five. The venue: public, private, and strategic

The same company is worth different amounts depending on where it is valued. Public markets, private financings, and strategic acquisitions price the identical business on different scales, and the gaps are large. In fintech, private companies have carried multiples several times those of their public peers; in cyber, strategic acquisitions have cleared at more than twice public levels.[3,5] Part of this is liquidity and part is the buyer's perspective, a strategic acquirer pricing the asset inside its own platform rather than as a standalone. For a seller, the venue is a choice, and choosing the wrong one, running a financial process for an asset that two strategics would fight over, can cost more than any operating shortfall. The Rule of 40 is silent on venue, which is precisely why a score-led view of value misses one of the largest levers a seller controls.

These drivers are not independent of the Rule of 40; a compounding, defensible business often posts a strong score precisely because its economics are good. But the causation runs through the economics, not the score. Improving the score without improving the underlying revenue architecture or competitive position raises the smaller multiplier and leaves the larger one untouched. That observation leads directly to the traps in Part V.

PART VI

The traps the rule sets for founders

A rule that is widely trusted and partly wrong is more dangerous than one that is ignored, because it directs real effort toward the wrong target. Five traps recur in the companies Windsor Drake advises.

Optimizing the score in a low-base category

The most common error is pouring effort into clearing the Rule of 40 in a vertical where the base multiple is structurally low. A payments business that grinds its score from 35 to 45 may earn a point or two of multiple. The same effort spent making the revenue more software-like, more recurring, more expanding, can move the company toward a different and higher base entirely. Founders should ask which dial they are turning before they turn it.

Letting the company be categorized as the lower-multiple thing

Many lower-middle-market companies straddle categories: a fintech with a real software product, a services business with a recurring platform, a security tool with an AI core. How the company is framed to the market decides which comp set it is priced against, and the gap between the two can be several turns of revenue. Allowing a buyer or banker to default the company into its lowest-multiple description, when a defensible higher-multiple framing exists, is value left on the table.

Starving growth or positioning to hit the number

Because the rule rewards margin, a company under pressure can hit 40 by cutting, trimming the investment in growth, product, or go-to-market that builds the very compounding and defensibility that set the base. Hitting the score while eroding the franchise is a poor trade, and a sophisticated buyer will see through a number propped up by underinvestment.

Benchmarking against the wrong comparables

A founder who reads that software trades at five times and assumes that is the ceiling, or one who reads that AI trades at twenty-four times and assumes that is the floor, is anchoring on the wrong set. The relevant comp is the company's true category and revenue architecture, not the broad average of an adjacent one. Mis-anchored expectations produce both disappointment and, occasionally, deals declined that should have been taken.

Believing the score travels to the negotiation

The Rule of 40 is a screening heuristic, not a price a buyer will pay. In a live process, value is set by the buyer pool, the competitive dynamic, and the strategic fit, the things in Part V. A strong score helps a company into the room; it does not set the number on the page. Treating the rule as a valuation rather than a filter leads founders to under-prepare the parts of the process that actually move price.

Treating the benchmark as fixed across the cycle

The final trap is to take any single year's multiples, including the ones in this report, as permanent. Every base in Part IV is moving, and the most extreme, the AI premium, is also the most fragile. A founder who builds a plan around today's AI multiple, or who assumes the software re-rating of recent years will repeat, is anchoring to a moment rather than a trend. The durable signal is structural: compounding revenue and genuine defensibility are rewarded, and that holds across cycles. The specific multiples are weather, not climate, and a seller should plan around the climate while timing around the weather.

The five traps share a root. Each comes from treating a single, simple number as if it carried more information than it does. The Rule of 40 is a good first question and a poor last one. A founder who keeps it in its place, as a quick read on operating health rather than a verdict on value, avoids all five, and arrives at a process with expectations calibrated to how the market actually prices the company rather than to how a heuristic says it should.

PART VII

Using the benchmark

The benchmark is only useful if it changes a decision. Below, what it means for the two sides of a transaction, followed by a short way to locate any company within it.

FOR FOUNDERS AND SELLERS

Find your true base before you fix your score. Identify the category and revenue architecture the market will actually price you on, then decide whether your effort is better spent improving the score or moving the base.

Frame the company as its highest defensible category. Where the business genuinely straddles, prepare the evidence that supports the higher-multiple description, and take it to buyers who price on that scale.

Build the buyer pool deliberately. The strategic premium is earned in the process, not on the income statement. A competitive set of strategic acquirers is worth more than another few points of margin.

Choose the venue with intent. A financial process and a strategic process can value the same company very differently. Decide early which buyers will pay the most and design the process around reaching them, rather than defaulting to the easiest path to a deal.

FOR BUYERS AND INVESTORS

Do not pay the AI base for SaaS economics, or accept the payments base for compounding revenue. Underwrite the revenue architecture and defensibility directly rather than trusting the category label or the headline score.

Treat a strong Rule of 40 as a filter, not a thesis. The score confirms operating health; it does not establish that the multiple is justified. The base, and what supports it, is where diligence belongs.

Watch for narrative-only bases. A multiple sustained by category story rather than economics is the one most exposed to a re-rating, and the place where disciplined buyers find or avoid risk.

Use the cross-vertical gap as a sourcing lens. Companies mispriced as the lower-multiple category, when their economics support a higher one, are where value is created. The reverse, paying a high base for economics that do not support it, is where it is destroyed.

Locating a company in the benchmark

The sequence below is how Windsor Drake positions a company before a process. It treats the Rule of 40 as the last input, not the first.

WORKSHEET

Four questions, in order, before the score

1. What is the base?	Which category and comp set will the market actually price this on? Start from Part II, not from the broad software average.
2. Does the revenue compound?	Does a unit of revenue deepen with the customer, or must it be re-won? This moves the base more than the score does.
3. How defensible is it?	What makes the revenue hard to take away? Retention, switching costs, data, mission-criticality. This is what a strategic buyer pays up for.
4. Then, what is the Rule of 40?	With the base established, the score fine-tunes the multiple within the range, roughly 1.1x per ten points. It is the adjustment, not the answer.

Source: Windsor Drake. The ordering is deliberate: the score is the last question, because it is the smallest lever.

Used this way, the Rule of 40 keeps its proper role. It is a fast, honest read on operating health, and a company that clears it is usually a better business than one that does not. It is simply not the thing that sets the price. The price is set by the category, the architecture of the revenue, the defensibility of the franchise, and the buyers in the room. A founder who gets those right and posts a middling score will out-sell one who gets them wrong and posts a perfect one.

PART VIII

Reading private-company multiples with care

Most lower-middle-market sellers are private, and private multiples do not behave like public ones. A founder benchmarking against either set without understanding the difference will misjudge what the company is worth, usually in the optimistic direction.

The headline gap is large and runs opposite to intuition. Private companies in these markets often carry multiples well above their public peers; private fintech has averaged in the mid-teens against a public average near six.[3] A founder who hears that private companies trade richer than public ones, and stops there, will anchor far too high. The gap is real but it is not free money, and three things explain it.

The numbers are set by the marginal investor, not the market

A public multiple reflects the price at which the marginal share changes hands every day, across the whole float. A private round reflects the price a single lead investor agreed to pay for a slice of new shares, at a moment chosen to be favorable. The private number is a point, not a distribution, and it is struck under conditions, competitive financing, a hot category, that flatter it. Treating that point as the value of the whole company, or as a price an acquirer would pay for all of it, overstates what is actually available.

The headline valuation hides the structure beneath it

Private valuations are frequently inflated by terms that the headline number omits. Liquidation preferences, ratchets, and other downside protections let an investor agree to a high nominal price in exchange for guarantees that lower their real risk. The result is a headline valuation that overstates the value of the common equity, sometimes substantially. A founder reading a private comp, or celebrating their own last round, should ask what was promised to get that number, because the structure can matter as much as the multiple.

The exit reprices against a different buyer

A financing values a minority stake bought by an investor who expects further upside. An acquisition values the whole company bought by an owner who has to operate it. These are different transactions priced by different parties, and the financing multiple does not carry over to the exit. This is the same lesson as the public-versus-strategic gap in Part V, viewed from the private side: the venue, and the buyer's perspective, can move the number more than the fundamentals do, and a private round is the least reliable guide to what a sale will clear.

> *A private round is a single price struck on a good day, dressed up by terms the headline hides. It is the least reliable benchmark a seller has.*

There is a further wrinkle specific to the current market. In AI and AI-adjacent categories, private financings have been struck at multiples far above anything the public market would support, and those numbers now circulate as reference points for companies that are not remotely comparable. A founder in software or fintech who benchmarks against an AI private round is anchoring to the least relevant comp available, set in the most favorable conditions, for the most narrative-driven category. The gap between that number and what an acquirer would actually pay for a different kind of business is not a discount to negotiate; it is a category error to avoid.

The practical guidance is to benchmark a private company against the economics, not against another private headline. Locate the business in the public scale established in Part II, adjust for its real revenue architecture and defensibility, and treat private comps as context rather than evidence. The Rule of 40 has the same standing here as everywhere else in this report: a useful read on health, and a poor guide to price. In the private market, where the price signals are noisier and more flattering, that discipline matters most.

CONCLUSION

The score is the smallest number on the page

The Rule of 40 endures because it is honest about one thing and silent about everything else. It tells a founder whether the business is operating in healthy balance. It does not tell anyone what the business is worth.

The evidence in this report points to a single, useful correction. Value in software and its neighboring markets is set first by category and revenue architecture, second by defensibility and the buyer pool, and only last by the growth-plus-margin score that gets the most attention. The same financials support a valuation that varies seven-fold depending on how the market reads the company. A payments business that passes the test trades below an AI business that fails it. A lending platform and a wealthtech platform, neighbors on the org chart, trade an order of magnitude apart. None of this is visible in the score, and all of it decides the outcome.

For a founder, the implication is liberating rather than discouraging. The largest levers on value, what the company is, how its revenue is built, how defensible it is, and who is brought to the table, are also the levers a seller can actually pull, in the years before a process and in the design of the process itself. The score is the one number that is mostly fixed by the time it matters. The premium this report is named for is real, and worth earning. It is simply not where the money is. The money is in being the right kind of company, in the right market, sold to the right buyers, and in refusing to be priced as anything less than what the business actually is.

> *A founder who gets the category, the revenue architecture, the defensibility, and the buyer pool right, and posts a middling score, will out-sell one who gets them wrong and posts a perfect one.*

This does not diminish the Rule of 40. A business should still want to clear it, because the discipline it enforces, growing without burning or earning while still growing, is the discipline of a healthy company. The argument of this report is only about hierarchy. The score belongs near the end of the valuation conversation, not the start, and a seller who treats it as the headline will consistently misjudge both what the company is worth and where the effort to raise that worth should go. Put it in its place, behind category, revenue architecture, defensibility, and the buyer pool, and it becomes what it was always meant to be: a fast, honest check on operating health, and the last small adjustment to a number that four larger forces have already set.

APPENDIX

Multiples by sub-segment

The reference tables below collect the sub-segment multiples used throughout this report in one place. They are central tendencies drawn from public market data and disclosed transactions as of mid-2026; dispersion within each line is wide.

APPENDIX A

Fintech, by sub-segment

SUB-SEGMENT	EV/REVENUE	REVENUE CHARACTER
Wealthtech	~18–25x	Compounds
SMB / enterprise fintech software	~15–17x	Compounds
Public fintech (overall avg)	~5.9x	Blended
Payments	~4–6x	Processes
Lending platforms	~2.5x	Processes
Fintech M&A (2025 avg)	~4.4x	Down from ~7.7x at 2021 peak

Source: Finro, Multiples.vc, Windsor Drake fintech research.[3]

APPENDIX B

Cybersecurity, by segment

SEGMENT	PUBLIC EV/REVENUE	NOTE
Cloud security	~18–22x	Highest growth
Application security	~13–15x	High growth
Data security / GRC	~6–8x	Middle band
Legacy hardware / services	~4–6x	Discounted
Public median	~7.8x	Average ~9.2x
Strategic M&A (avg)	~18.8x	Deals observed ~10–45x

Source: Finro cybersecurity data, Multiples.vc, Windsor Drake cyber research.[5]

APPENDIX C

Software and AI, reference points

COHORT	MULTIPLE	BASIS
Public software (median)	~5.1x	EV/revenue, end-2025
Mature, profitable software	~26.6x	EV/EBITDA
AI (median)	~24x	EV/revenue
AI (range)	~10–50x	By stage and moat
Anthropic (2025 round)	~37x	Implied revenue multiple
Databricks	~28x	ARR multiple

Source: Aventis Advisors, Meritech, Value Add VC.[1,2,4]

THE EARLY READ

Into 2026: the premium widens into a discount

The central finding of this report, that the same Rule of 40 score is priced very differently by category, hardened in 2026 into something sharper: the dispersion is no longer only a premium paid for the favored, it is now a discount applied to the rest, as AI-native economics pull the two ends of the software market further apart.

The early-2026 data widened the spread the report measures. AI-native companies drew premium multiples, with the AI-native share of large strategic technology deal value approaching half, while median software M&A multiples held in a narrow band near four to five times revenue and legacy software valuations came under renewed pressure as buyers reassessed the durability of traditional subscription models. The same Rule of 40 score that earns a rich multiple inside an AI-native story earns a guarded one inside a legacy one, and the gap between those two readings grew.

> *In 2026 the dispersion stopped being only a premium for the favored and became a discount for the rest. Category sets the base multiple; the Rule of 40 only adjusts it.*

Category over score

The 2026 evidence reinforces the report's core argument and sharpens it. A company's category, and specifically which side of the AI divide the market places it on, now sets the base multiple before the Rule of 40 score adjusts it, and that base moved further apart across categories during the year. A strong score cannot rescue a company the market reads as legacy; a weaker score is forgiven in a company the market reads as AI-native. The metric still matters, but the category matters more, and in 2026 it matters more than it did.

What it means for founders

For a founder, the practical lesson is to understand the multiple the market will start from before optimizing the score that adjusts it. Improving growth and margin is worthwhile, but positioning the company on the right side of the AI-native divide, through product, data, and narrative, moves the base from which all of it is valued. In a year when the premium widened into a discount, the highest-leverage work is establishing which multiple a company is measured against in the first place.

Sources for this section: Software Equity Group (SaaS M&A multiples, 4–6x band); PwC, “How AI is reshaping software valuations”; FE International, “AI M&A Trends 2026”; Windsor Drake valuation research.

Methodology and sources

This report compares valuation across four software-adjacent verticals using a common lens, the Rule of 40, defined here as revenue growth rate plus profit margin, most often computed on a forward basis. Source definitions vary, particularly between free cash flow margin and operating or EBITDA margin, and between trailing and forward revenue; where figures are compared across sources, those differences widen the error bars and the central tendencies should be read as directional.

Multiples are EV/revenue unless stated otherwise and are drawn from public market data and disclosed transactions as of mid-2026. Every category discussed here contains wide internal dispersion; the medians and ranges reported describe the sample and are not forecasts for any individual company. The "base multiple" and "at Rule of 40" figures in Exhibits 1 and 6 are Windsor Drake analytical constructs built by applying the observed within-software quality premium to each vertical's prevailing base multiple. They are illustrative of the structure of the finding, not transaction quotes, and any company can fall well outside them.

The central claim of this report, that category and revenue architecture set the base multiple while the Rule of 40 adjusts it, is a structural observation that holds across the cycle. The specific magnitudes, especially for AI, reflect a mid-2026 market and would change, in some cases sharply, with a re-rating of any vertical. Nothing here is investment, legal, tax, or accounting advice.

Endnotes and sources

1. Rule of 40 definition, predictive power, and the within-software premium (median score ~28; ~20% above 40; ~1.1x EV/revenue per 10 points; above-line roughly double below-line): Meritech Capital benchmarks and Meritech Software Pulse; Benchmarkit, "Rule of 40, Analysis of Meritech Capital Benchmarks"; Aventis Advisors, "Rule of 40 in SaaS: 2026."
2. Software / SaaS multiples (~5.1x EV/revenue end-2025; ~26.6x EV/EBITDA for mature names; profitability re-rating): Aventis Advisors, "SaaS Valuation Multiples: 2015–2026" and "Software Valuation Multiples"; Windsor Drake, "SaaS Valuation Multiples 2026."
3. Fintech multiples and the compounding-versus-processing split (public avg ~5.9x; M&A ~4.4x; wealthtech ~18–25x; payments ~4–6x; lending ~2.5x): Finro, "Fintech Valuation Multiples Q1 2026"; Multiples.vc, "Public Fintech & Payments Valuation Multiples, June 2026"; Windsor Drake, "Fintech Valuation Multiples 2026."
4. AI multiples (10–50x; median ~24x; Anthropic ~37x; Databricks ~28x ARR; defensibility over growth): Aventis Advisors, "AI Valuation Multiples in 2025"; Qubit Capital, "AI Startup Valuation Multiples"; Value Add VC, "AI Company Valuations 2026."
5. Cybersecurity multiples (public median ~7.8x; M&A avg ~18.8x; cloud security ~18–22x; legacy ~4–6x): Finro, "Cybersecurity Valuation Multiples" and database editions; Multiples.vc, "Cyber Security Cloud"; Windsor Drake, "Cybersecurity Valuation Report Q4 2025."
6. AI-security acquisitions (Palo Alto Networks / Protect AI; SentinelOne / Observo AI; Check Point / Lakera): Finro and contemporaneous press reporting, 2025.

ABOUT

About Windsor Drake Research

Windsor Drake is a sell-side M&A advisory firm for founders and boards of software, financial technology, artificial intelligence, and cybersecurity companies in the lower middle market.

The firm advises technology founders through the full arc of a sale, from positioning the company and building the buyer universe to running a competitive process and negotiating terms. Its research program exists to make the dynamics that decide outcomes legible to the people on the selling side, who face these transactions once or twice in a career against buyers who do them constantly. Each report is built from public sources and the firm's transaction experience, and is written to be useful before a process begins, when the decisions that determine price are still open.

This report is part of the Windsor Drake Research series, a set of independent studies on valuation, deal structure, and consolidation across the firm's sectors. Other titles examine the talent-led acquisition of AI companies and the consolidation of the cybersecurity industry, alongside sector-specific valuation benchmarks. The series is intended to be read by founders, boards, investors, and the journalists and analysts who cover these markets.

Contact

Windsor Drake maintains offices in Toronto and New York. Founders considering a transaction, and members of the press citing this research, can reach the firm through windsordrake.com.

Windsor Drake Research · Toronto · New York · windsordrake.com