SCALE

The Four Vital Signs of SaaS

Metrics Matter

Metrics accelerate the journey from founder-led sales to a repeatable sales motion. SaaS metrics should be to a management team what patient vital signs are to an emergency room doctor: a simple set of universally understood numbers that allow a doctor to quickly know how ill a patient is and what needs fixing first.

A company culture that emphasizes measurement and optimization is a critical success factor for software businesses as they scale. In the early stages of growth, they're also the best indicators of where founders should direct their attention — they tell you what needs fixing next.

This guide looks at why, when, and how to use metrics in scaling a software startup. Each section progresses from a broad overview to more nuanced and precise views of what we call the "Four Vital Signs of SaaS" covering growth, sales efficiency, customer retention, and cash burn.

We hope to provide founders a handbook for creating a metrics culture throughout their company. As well as a reference guide for getting the most out of Scale Studio, the benchmarking tool that Scale portfolio companies use to measure their performance.

Table of contents

Growth

Growth is the central purpose of a startup. Growth rates are the most important measures of whether your startup is creating value.

There's More to Growth than Growth Rate

There is way more to growth — and managing for growth — than growth rate. Yes, your company's growth rate matters in the process of securing top-tier venture investors. But the analysis that those investors perform on your data room goes much deeper than simply cross-checking how fast your ARR has grown and projecting the rate it will grow in the future.

In this section, we cover three big-picture concepts that are foundational for how the Scale investing team measures, predicts, and benchmarks a start-up's growth.

iCAGR: Growth Right Now

Back in 2016, Scale's Jeremy Kaufmann published The Growth Rate Mirage, an analysis of how different growth rate calculations yield different results (and tell different stories) for early-stage, high-growth startups.

The article introduced Scale's "instantaneous compound annual growth rate" or iCAGR. The name might be a little unwieldy but it's value is pretty straightforward: it is a calculation that we use to understand how fast a company is growing right now. Why that's useful takes some explaining.

The most common measure of growth is the quarter-over-quarter / year-over-year growth rate. It's the rate most everyone in SaaS is using when they say "We expect to grow 150% this year." The formula is the same whether you're talking ARR or GAAP revenue:

$$\frac{\text{Ending ARR}_{Q} - \text{Ending ARR}_{Q\text{-}4}}{\text{Ending ARR}_{Q\text{-}4}}$$

So what's the problem? There's not a problem, but there is a subtle shortcoming: the year-over-year growth rate can mask a recent slowdown in growth. Here's how. Follow the ARR Growth Rate line in the following table of a hypothetical early-stage hypergrowth startup.

($000's)	1Q 2015	2Q 2015	3Q 2015	4Q 2015	1Q 2016	2Q 2016
Run-Rate Data						
Beginning ARR	$2,091	$2,502	$3,402	$4,416	$5,978	$6,886
Ending ARR	$2,502	$3,402	$4,416	$5,978	$6,886	$7,711
ARR Growth Rate (Ending ARR YoY)	156%	196%	184%	186%	175%	127%
ARR iCAGR (Net New ARR/ Closing Balance)	84%	156%	129%	153%	64%	50%
Net New ARR	$411	$900	$1.015	$1,562	$908	$825

Now look at what's going on with Net New ARR: it's rapidly slowing on a sequential basis. There can be perfectly understandable reasons for why, but we would see that as a red flag. The point is that year over year growth looks strong despite slowing in Net New ARR, a key leading indicator of future growth.

That's where iCAGR comes in. By tweaking the standard compound annual growth rate formula, iCAGR tells us more about recent changes in growth. The formula looks like this:

$$\left(\left(1+\left(\frac{\text{Net New ARR}_Q}{\text{Ending ARR}_Q}\right)\right)^4\right)-1$$

We've learned that iCAGR is best used alongside other metrics (it is a weather vane, not the weather) and that a "good" iCAGR varies considerably by revenue level (it's best used to compare companies at similar stages of growth). You'll see iCAGR front and center on our Scale Studio Reference Benchmarks page.

Definitely spend some time with the original article for the full analysis and commentary. It's a great primer on a super useful SaaS growth metric.

Growth Persistence: A Reality Check on Growth Projections

Way back in 2012, Andy Vitus led an effort to answer a tricky question. The question arises at every company at the end of every year, as the annual planning process kicks off: How fast can, or should, we grow? It's tricky because "can" and "should" will often have two very different answers.

"How fast can we grow?"

This can be calculated because, at any given point in time, a company's resources (people, processes, capital) will determine a theoretical maximum to how much revenue can be generated next year. And thus next year's growth rate. Think of the jet engines hanging from a plane's wings: they can power the airplane up to a certain speed and no faster.

"How fast should we grow?"

This is a more complicated question. First, it can't be answered without a defined goal. Is your company racing to cash flow breakeven as quickly as possible to reduce reliance on outside capital? Are you maximizing growth to keep your company on the shortest possible path to IPO? Or something in between?

The best answer to these questions is often "it depends". But we wanted to give our portfolio companies better advice than "it depends" and landed on the idea of "growth persistence". We won't repeat the original growth decay article, but take the time to read it. As we'll see in a second, it remains incredibly useful and accurate even eight years later.

The first big takeaway is this:

For SaaS companies from $1M to $1B in ARR and 10% to 120% current-year growth rates, next year's growth rate tends to be 85% of this year's growth rate.

That right there is a powerful reality check on your annual planning assumptions. If your ARR falls into the $1M to $1B range, and you expect to match or exceed your current year growth rate next year — stop immediately and dig deep into your assumptions. Growth acceleration at scale is pretty rare.

To see just how rare, here's the original data set. The points far above the trendline represent companies that greatly accelerate year over year growth. There aren't that many.

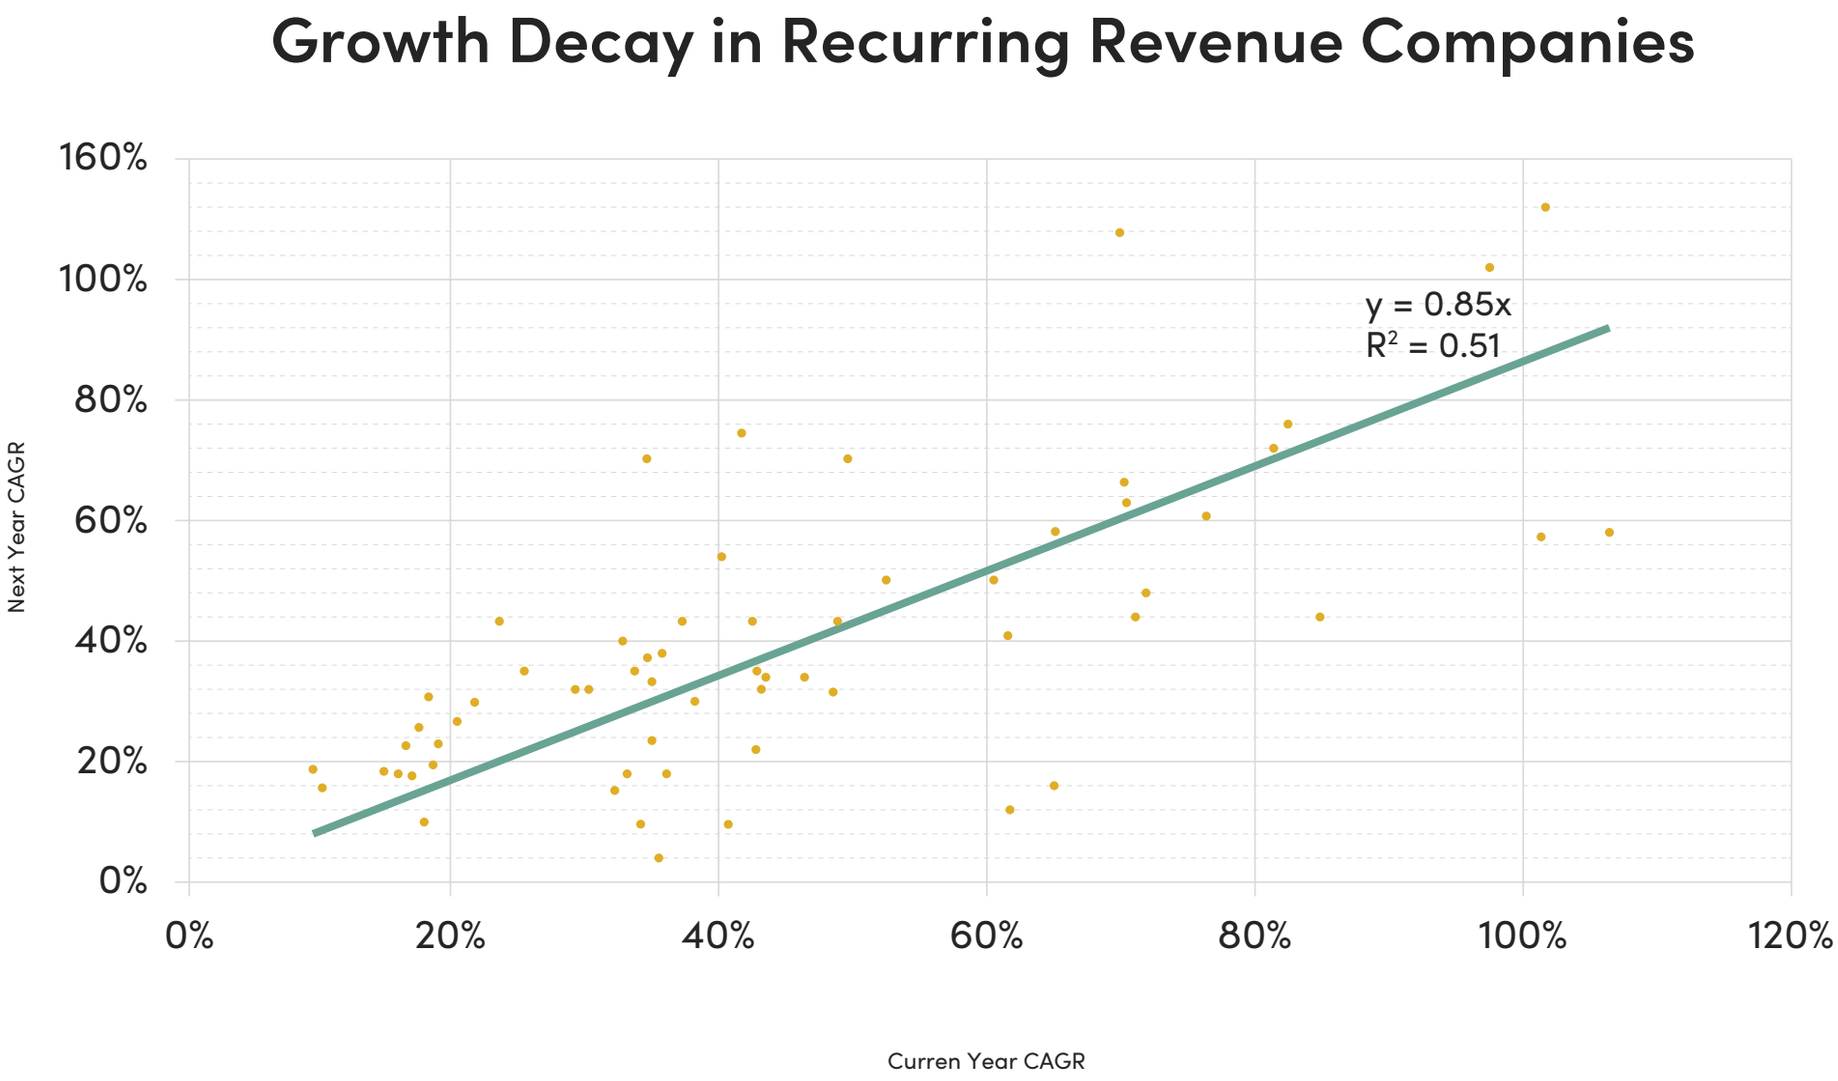

Important note: the analysis omitted companies in the very early, very high growth rate category. Out of the gate, early-stage startups can and should have accelerating growth rates. But even at 120% year-on-year growth, growth persistence is at work.

In case you're wondering whether that 85% figure has held up: it has. We recently re-ran the analysis on the public SaaS companies.

Year	Growth Persistance Rate
2017	77%
2018	95%
2019	91%
2020E	70% (COVID shock)
2021E	83%

** Note that this data derives from a smaller set of companies than the original analysis. We ran this analysis on the readily available public SaaS comps to verify that the original 85% holds true.*

We see that annual growth persistence rates jump around but cluster around the original 85% figure.

As a rule of thumb, then, 85% holds up — keep it in mind every year during annual planning.

The bottom-line takeaway: Increasing next year's growth rate is exceptionally difficult for recurring revenue businesses that have scaled beyond the early hyper-growth period. Even maintaining growth takes flawless execution.

The Mendoza Line: Defining the IPO Growth Track

> "How fast do I need to be growing to be interesting to a venture investor?"

Here is another question that often gets answered with an "it depends" or a rule of thumb like the Triple-Triple-Double-Double. Not satisfied with either answer, Rory O'Driscoll developed a model called the Mendoza Line to define the minimum growth rate at any revenue level necessary to be on track for an eventual IPO (and see the original article for background on the baseball reference).

We defined "IPO ready" as follows:

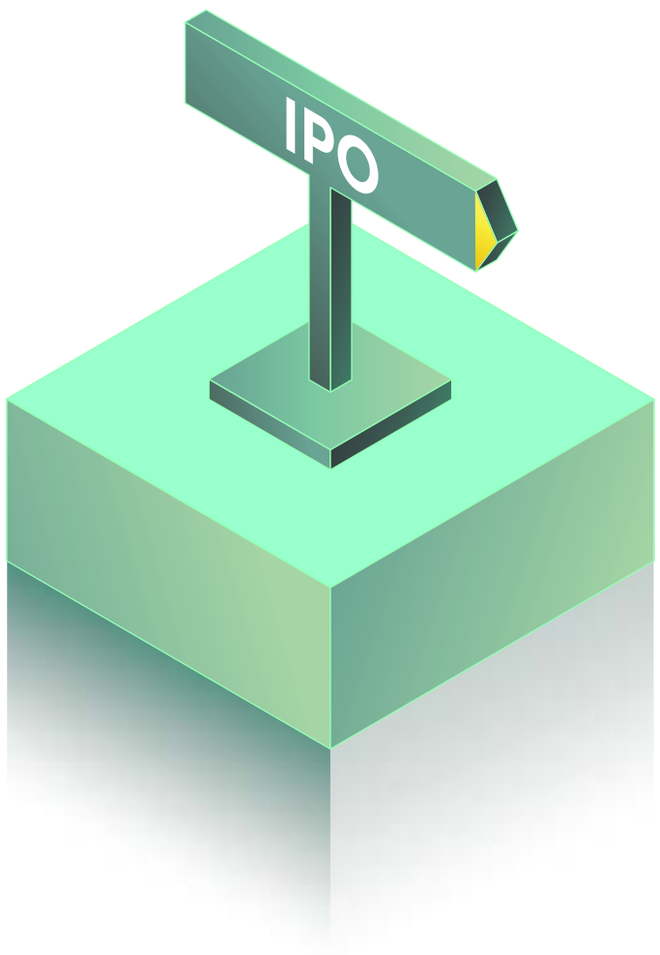

Looking at the realistic low bar of what it takes to be a public company, this implies being at run-rate ARR of $100MM at the time of IPO with a forward (next year) growth rate of 25%+.

Factoring in the reality of growth persistence, we derived the line that defines IPO viability at any point in time.

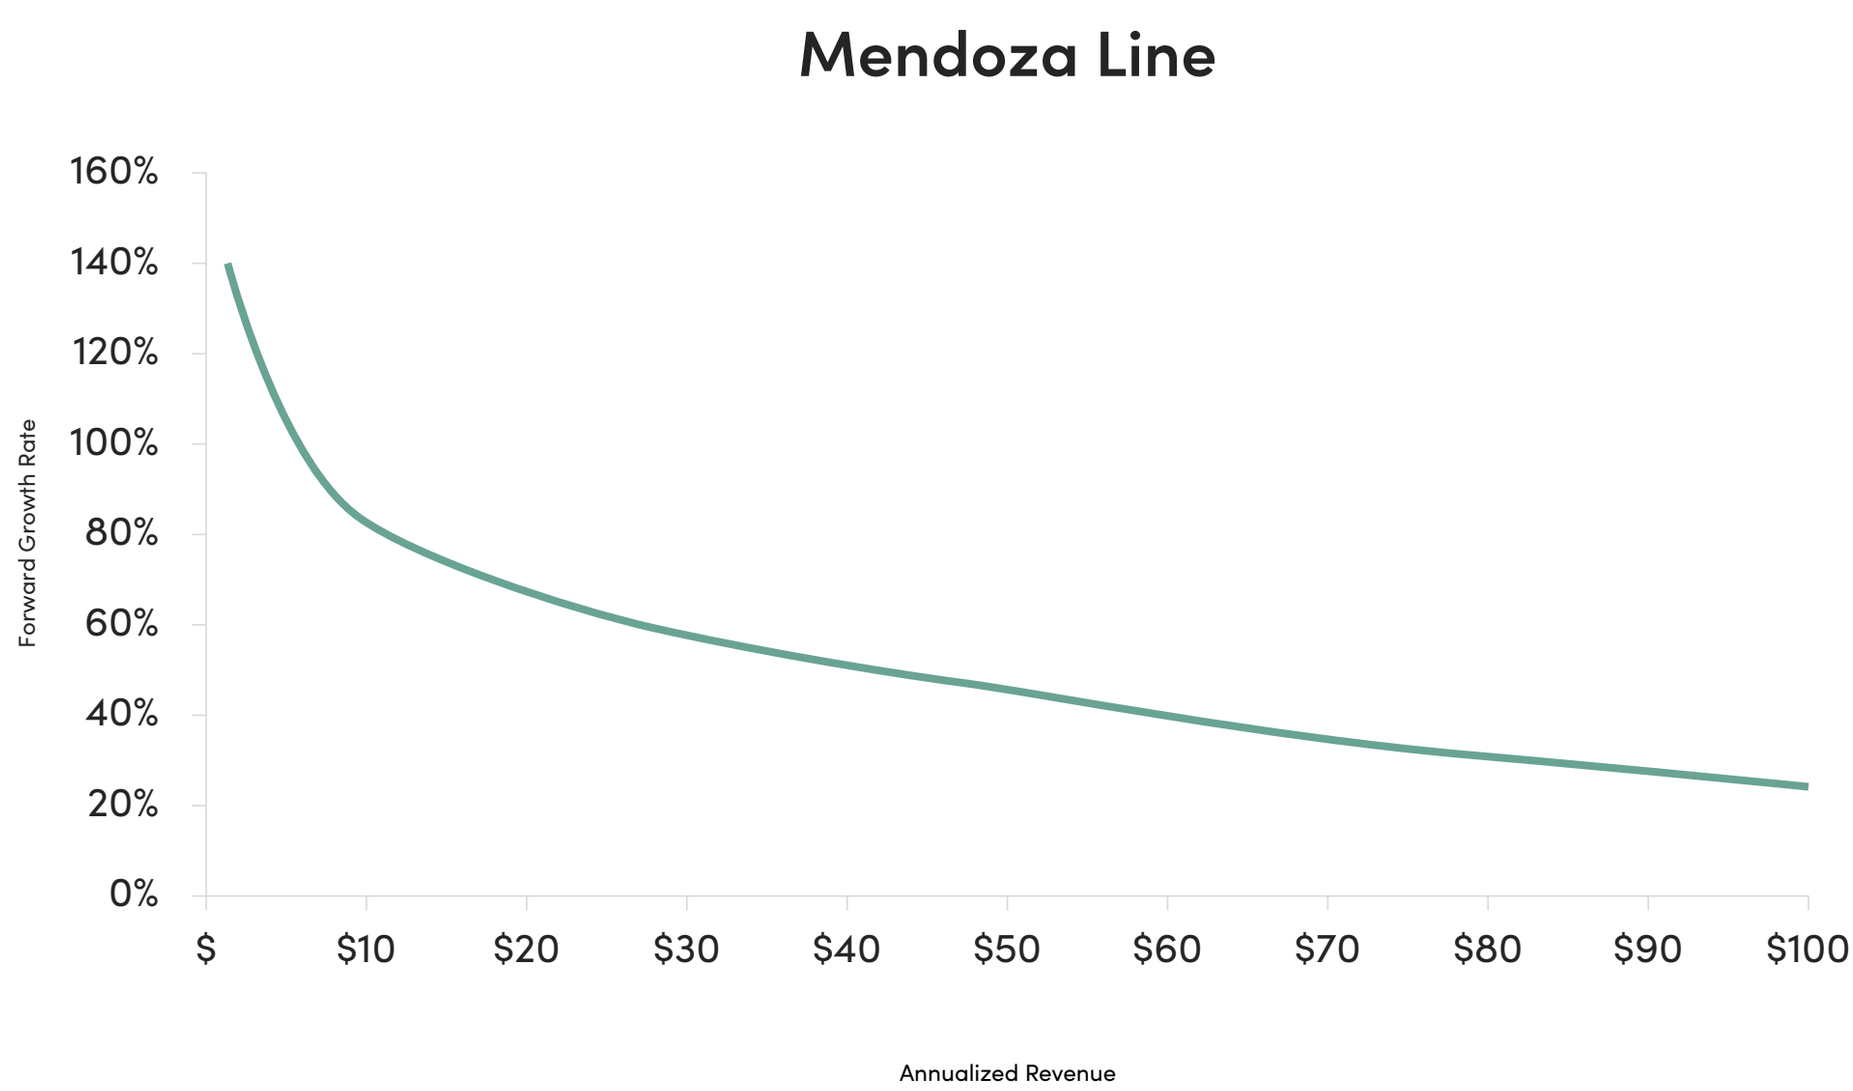

Next Year's ARR Growth Rate %

Current ARR	$1	$5	$10	$20	$30	$40	$50	$60	$70	$80	$90	$100
Mendoza Line Growth Rate	140%	94%	77%	62%	51%	46%	40%	36%	33%	30%	28%	25%
Years to 100M ARR	8.5	7.0	6.0	4.5	4.0	3.0	2.5	2.0	1.5	1.0	0.5	

The original Mendoza Line of SaaS article does a great job of explaining the model's strengths and shortcomings. Just know that your company's early growth rates are critically important because of the compounding effect of growth. The higher above the Mendoza Line early on, the better positioned you are to weather the persistent decay in forward growth as your company scales revenue.

Finally, keep in mind that "growth" isn't some feature of your company that can be modified at will. Your growth rate can be thought of as a score that tells a comprehensive story of how well your company is performing. Improvements to business processes, go-to-market strategies, or enterprise sales outcomes can contribute to improved growth rates.

As we stated in the original Mendoza analysis, "diagnosis isn't death". There are countless examples of companies that retool and re-accelerate growth.

Big picture:

Your current growth rate, times an estimate of growth persistence, is your most likely future. To change that future and raise that trajectory requires an act of will to overcome the inertia that all businesses face. It may mean new executives, new products, or changing business models.

Focus on the Growth of Your Growth

The previous section, There's More to Growth than Growth Rate, covered various measures of growth and two related Scale concepts, "growth persistence" (the predictable decline in growth rates year to year) and the "Mendoza Line" (the growth rate necessary at any revenue level to be on track for an eventual IPO).

Below we're going to drill down into another of our top-of-mind metrics, the Net New ARR growth rate. In a venture environment where few companies raise pre-revenue Series A rounds, investors are focused on rapid Net New ARR growth sooner than ever.

After all, the rough early quarters of COVID spotlighted an important fact about growth: it's possible to have negative Net New ARR growth and still show positive ARR growth. In SaaS, "growth" typically means ARR or GAAP revenue. It's an easy shorthand.

Yet it might be too much of a shorthand. The measure of growth that you really want to optimize for from the beginning — and put in your fundraising slide deck — is the growth rate of your growth rate. That is, your Net New ARR growth rate.

How Is Net New ARR Growth Rate Calculated?

First a quick reminder that Net New ARR (NNARR) is simply New ARR for a period (say, a quarter) plus Upsell and Expansion minus Churn. NNARR therefore reflects new sales (and thus your sales motion), the total amount of upsell and expansion (which speaks to both sales and customer success), and your churn situation (covering both product and customer success). It's a key metric to track organization-wide because it touches on pretty much every corner of your operations.

The formula for Net New ARR growth rate looks like this:

$$\left(\frac{\text{Net New ARR}_{Q}}{\text{Net New ARR}_{Q-4}} \right) - 1$$

Basically you're comparing how much the current quarter's Net New ARR grew compared to the same growth a year prior. More growth is always better — and more NNARR growth is even better than that.

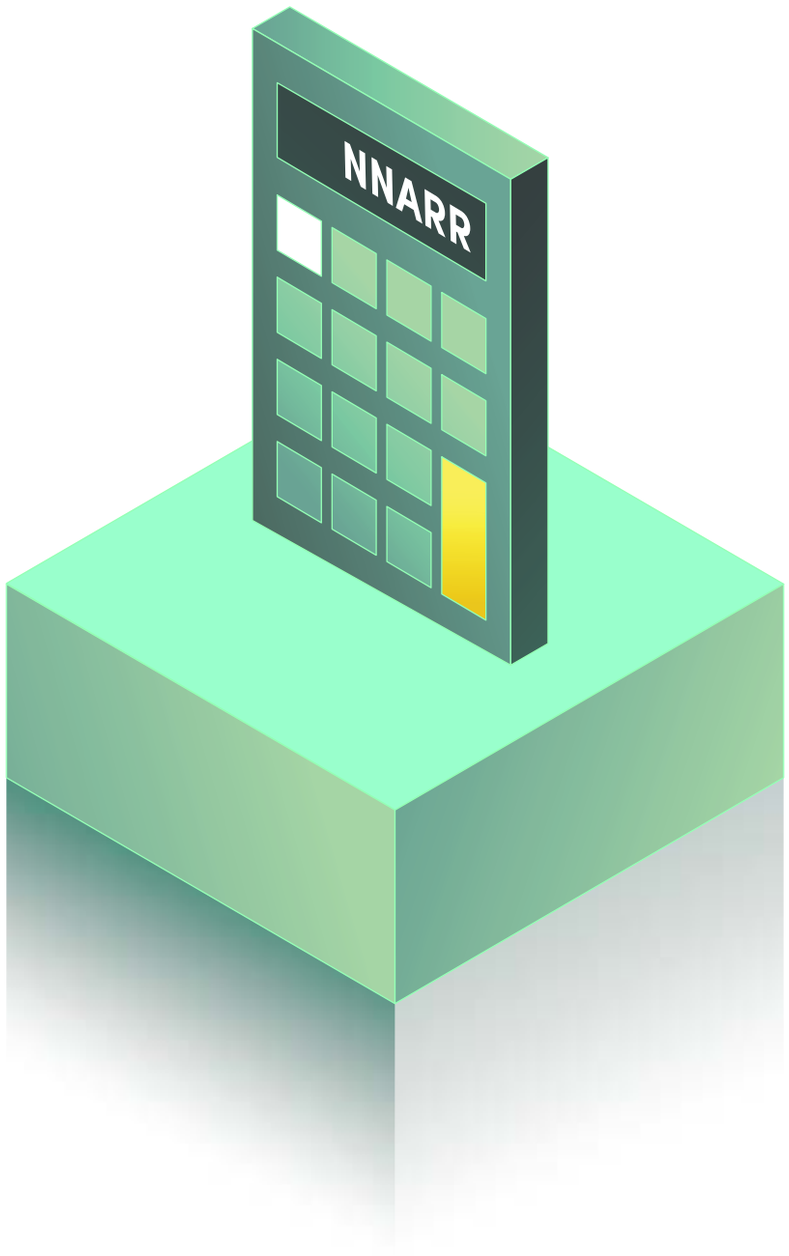

Why Does NNARR Matter?

There is a progression of importance from ARR to ARR growth to NNARR growth. ARR says how big you are. ARR growth is how fast you're getting bigger — and fast is worth more to investors than slow. NNARR growth is the same, but times ten.

In a late-2020 article, Scale's Rory O'Driscoll zeroed in on the importance of Net New ARR in annual plans for 2021, saying:

> Net New ARR is a very leveraged number and the growth rate of NNARR is an even more leveraged metric. Net New ARR captures all the "work" taking place. New customer additions, upsells, and churn are all reflected in this number without the smoothing out denominator effect that you see in growth rates of Total ARR. NNARR is like an amplifier in a power circuit, making painfully clear the impact of deceleration in a way that a Closing ARR growth rate does not.

Let's unpack this a little. Basically, NNARR is important because it rolls up a lot about your operations into a single number. It is also a leading indicator of future ARR growth. In a healthy business, NNARR growth accelerates as much or more than ARR itself.

For investors, strong NNARR growth signals a bunch of good things like go-to-market maturity, product-market fit, and happy customers all rolled up in one.

Good, Better, Best NNARR Growth Benchmarks

We'll close with NNARR growth rate benchmarks, always available and up-to-date on the Scale Studio Reference Benchmarks page.

Here are NNARR Growth Rate benchmarks by revenue:

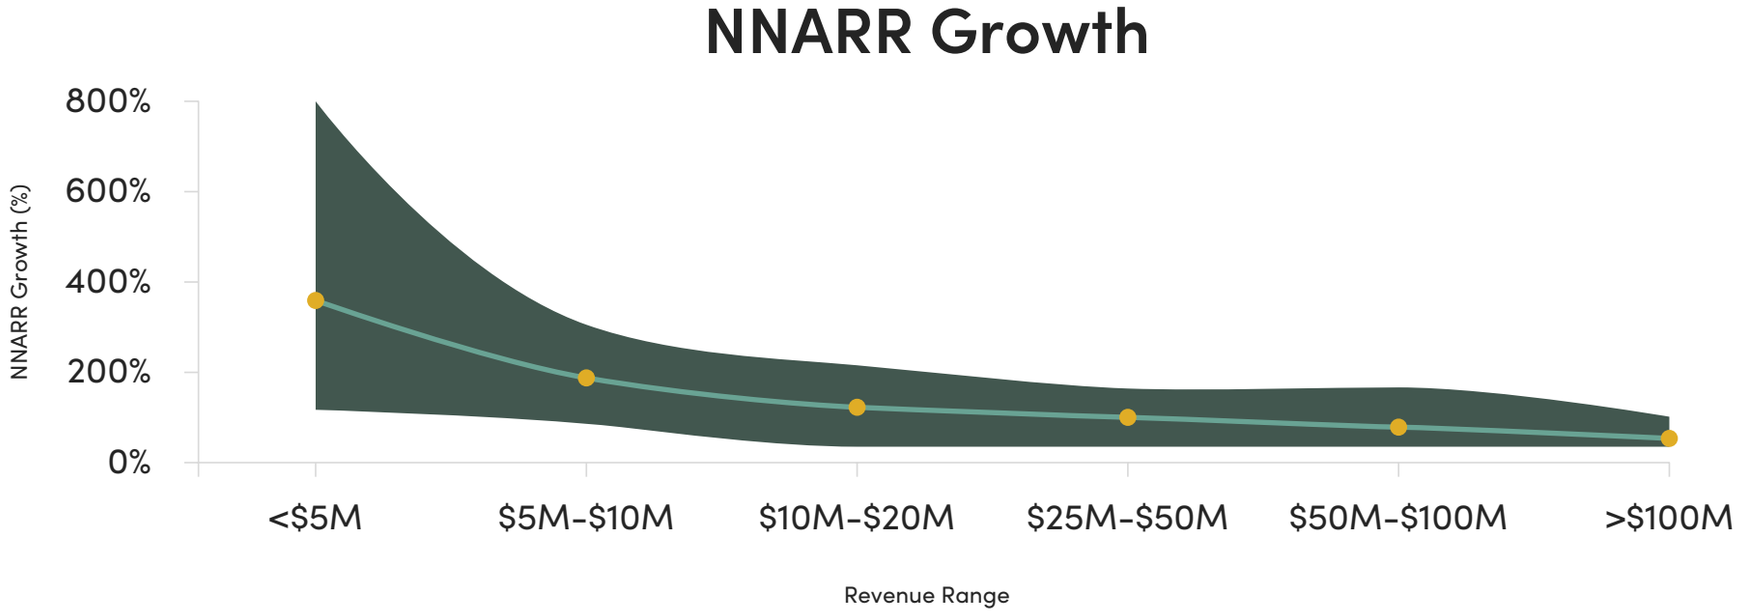

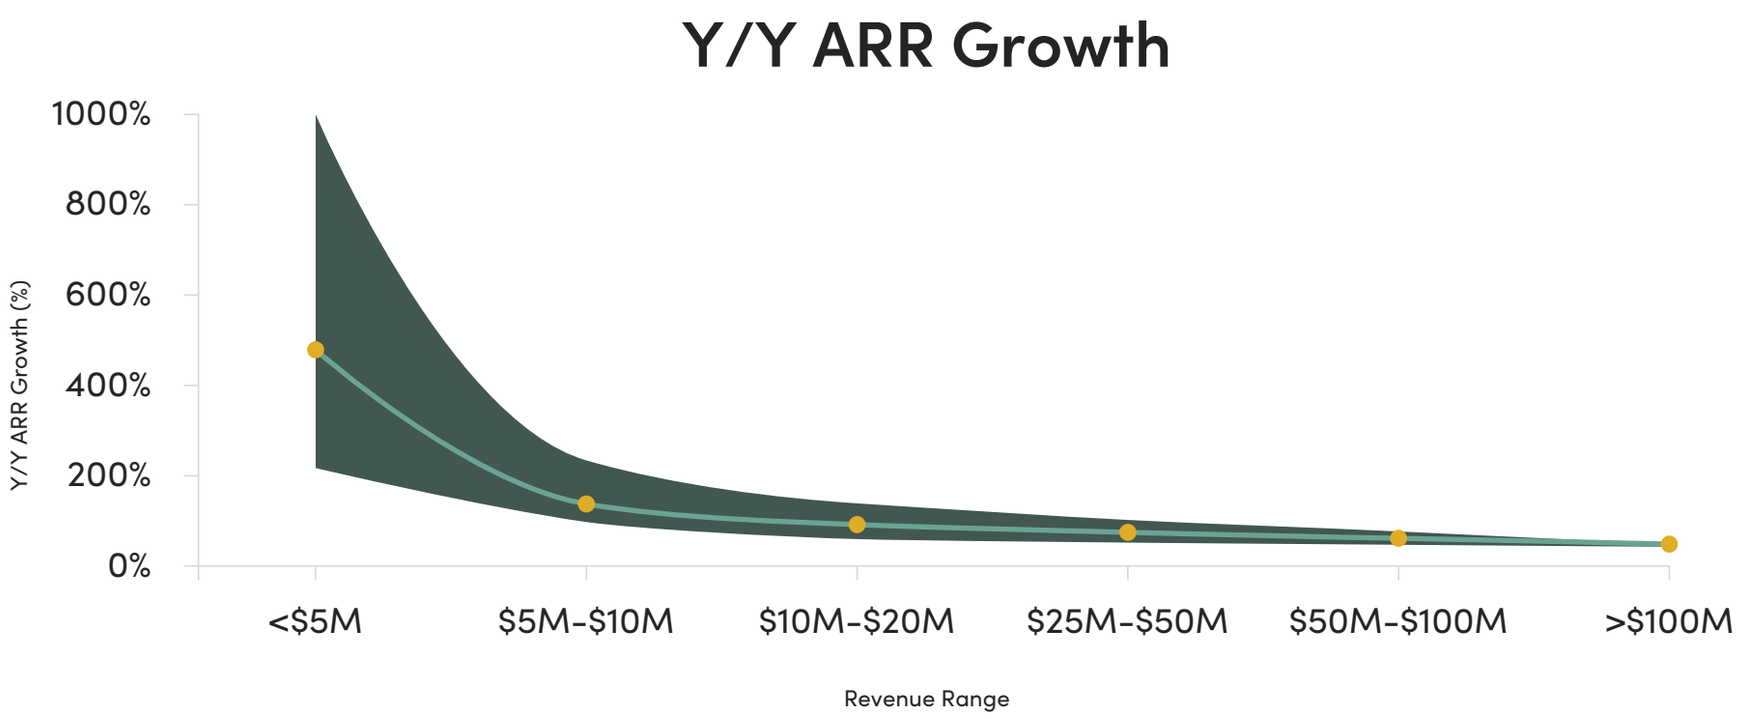

Notice that from $5M in ARR onwards, the highest performing companies have NNARR higher than their ARR growth rates. NNARR is the "engine" of future growth — and success. Investors know this, and these companies fetch the premium valuations round after round.

There are many ways to express growth. But you can't go wrong with a laser focus on maximizing Net New ARR growth rate from your first dollars in sales.

Efficiency

SaaS software at scale is all about distribution. Sales efficiency metrics tell you how much value your startup creates per dollar invested.

A Primer on SaaS Sales Efficiency

Let’s start with an observation: a lot has been written about SaaS Sales Efficiency (SE). So much that there are different measurement techniques, special cases, even conflicting schools of thought. It can be distracting and even a little confusing — especially because the concept itself is straightforward and easy to wrap your head around. And it’s worth it, because Sales Efficiency is a powerful tool for measuring how your company performs at one critical core function: selling software.

What Is SaaS Sales Efficiency?

Sales Efficiency tells you how efficiently your company creates value. It does this by comparing new revenue for a given period to what was spent on Sales & Marketing in the same period. It gives you a single, difficult-to-argue-with data point:

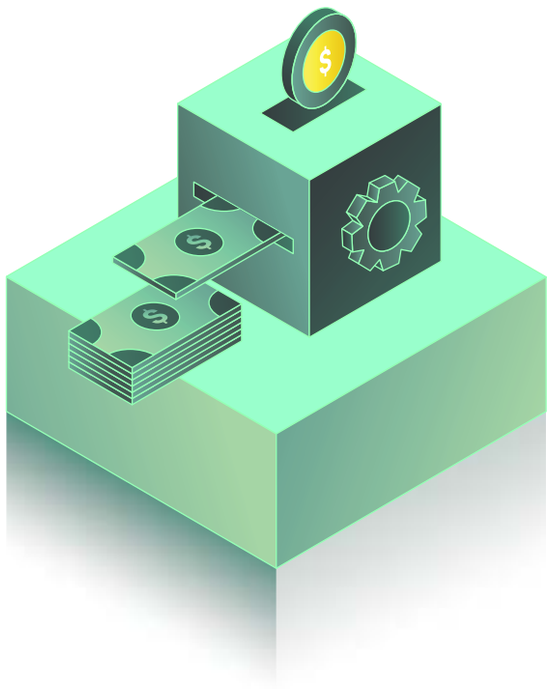

Sales Efficiency is the amount of new revenue generated for every dollar invested in selling and promoting.

Here's some context. In the Scale Studio dataset of 1,000+ growth-stage SaaS and cloud businesses, the long-term median Sales Efficiency hoovers around 0.7.* That means that a typical SaaS startup generates $0.70 in new ARR for every $1.00 it spends on Sales & Marketing. That is, "For every $1 this company spent on Sales & Marketing, you (Sales Leader) gave me $0.70 in sales."

* To be precise, the 0.7 figure is the median Magic Number (MN) for the dataset. MN is a GAAP-based version of Net Sales Efficiency, which we'll explain in more detail below.

Companies with highly efficient sales models and flawless go-to-market execution might see sales efficiency go above 1 but it's rare and speaks to unique and attractive aspects of their go-to-market model. Show a room of investors a data point like that in your pitch deck and your inbox will be overflowing.

Note that we have been using the generic "sales" and "revenue" and "sales and marketing" on purpose to keep things simple. But at this point we can get more specific. Sales Efficiency is most often calculated using either Gross New ARR or Net New ARR for a given period (like a month or quarter). Your finance team is already tracking these numbers, so it doesn't take extra work to feed it into the SE calculation. The denominator uses the entire "Sales & Marketing" line item from your P&L. (Generally, we prefer to exclude stock- based compensation (SBC) because it is a non-cash expense that gets included in GAAP financials.)

As we've said, there are endless exceptions and variations and schools of thought about the Sales Efficiency metric. But the bare bones, just-the-facts methods we share below provide a lot of insight for a lot less work, yielding metrics that anyone on your team can understand and make use of.

Gross Sales Efficiency: Focus on the Positives

Gross Sales Efficiency focuses on the positives: new ARR generated during the period with no adjustments for churn or downsell. The denominator is simply the total Sales & Marketing expense during the same period.

Gross Sales Efficiency

$$\frac{\text{Gross New ARR Q1}}{\text{S\&M Spend Q1}}$$

Because it uses unadjusted Gross New ARR, Gross SE gives you a clean look at how efficient your sales and GTM processes were during that period.

Net Sales Efficiency: Includes Wins and Losses

The numerator for Net Sales Efficiency is Net New ARR, which means we're accounting for both new sales and lost business. The denominator is the same as with Gross SE, total Sales & Marketing spending for the period.

Net Sales Efficiency

$$\frac{\text{Net New ARR Q1}}{\text{S\&M Spend Q1}}$$

One of the strengths of Net SE is how much information is captured in a single data point. Even early on in a company's growth path, there are an awful lot of people and processes involved in selling, marketing, and retaining customers. Net SE gives everyone in your company a target that's easy to understand and explain.

Magic Number: Benchmark across Companies

There is one important limitation in the Gross and Net Sales Efficiency calculations we just looked at. There's often a fair amount of variation from company to company on how ARR is recognized, which means that you can't easily formulate a comparison with other companies.

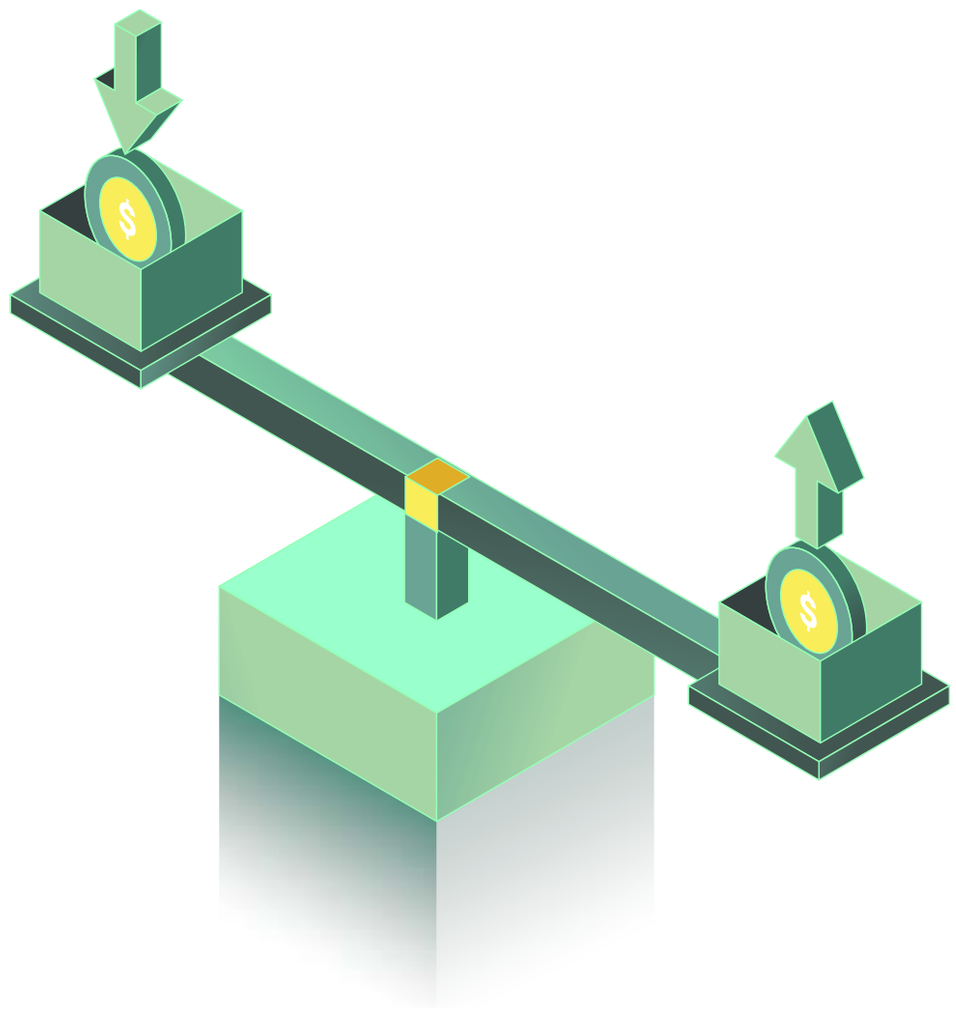

Scale developed the "Magic Number" (MN) as a workaround, replacing ARR with GAAP revenue. Because GAAP standardizes how and when sales are recognized, the Magic Number allows comparison between companies and, even better, benchmarking of one company against other comparable companies.

Here's how that looks:

Magic Number
$\frac{\text{(GAAP Rev Q2 - GAAP Rev Q1)*4}}{\text{S\&M Spend Q1}}$

MN replaces Net New ARR in the numerator with the difference in GAAP revenue between the most recent periods. It then annualizes that figure (*4) to bring it closer to an ARR value that is by definition an annual figure. With those adjustments, Magic Number becomes a great stand-in for Net Sales Efficiency that also allows you to compare apples-to-apples the efficiency of two or more different companies.

The Magic Number is the focus of the next three articles in this series. So we'll leave it at that for now.

Sales Efficiency and the Go-to-Market Model

Sales Efficiency behaves differently from, say, sales growth or churn rate as a company develops. Growth rates decline over time as a company scales. Likewise, burn rates tend to improve over time as companies mature and approach cash flow breakeven.

Sales Efficiency, however, tends to be pretty sticky. It often persists over time for a given company, especially after $10M in sales. This is important, because it makes developing an efficient go-to-market model critically important early on. That model — and its inherent efficiency — is most likely going to be with you for some time.

There are exceptions, of course. Some companies do in fact rewire their sales and go-to-market motions. But don't count on it. Instead, give everything you've got to developing an efficient sales engine early in the game. It will pay dividends down the line.

A History of the Magic Number

Roll back the tape to 2005, when SaaS represented only 10% of all software-related venture capital investment. At the time, many of the nuances of a subscription-based revenue model were not nearly as well understood as they have become today. During the process of evaluating Omniture for an eventual investment, Scale's Rory O'Driscoll was analyzing the company's revenue traction relative to its investment in Sales & Marketing. Seeing that they were able to generate more than $2 in first-year revenue for every $1 invested in their go-to-market engine, he exclaimed, "It's Magic!" The rest, as they say, is history.

Enterprise SaaS: Magic Number vs. EV/Revenue

Source: Scale Venture Partners / Scale Studio

SCALE STUDIO

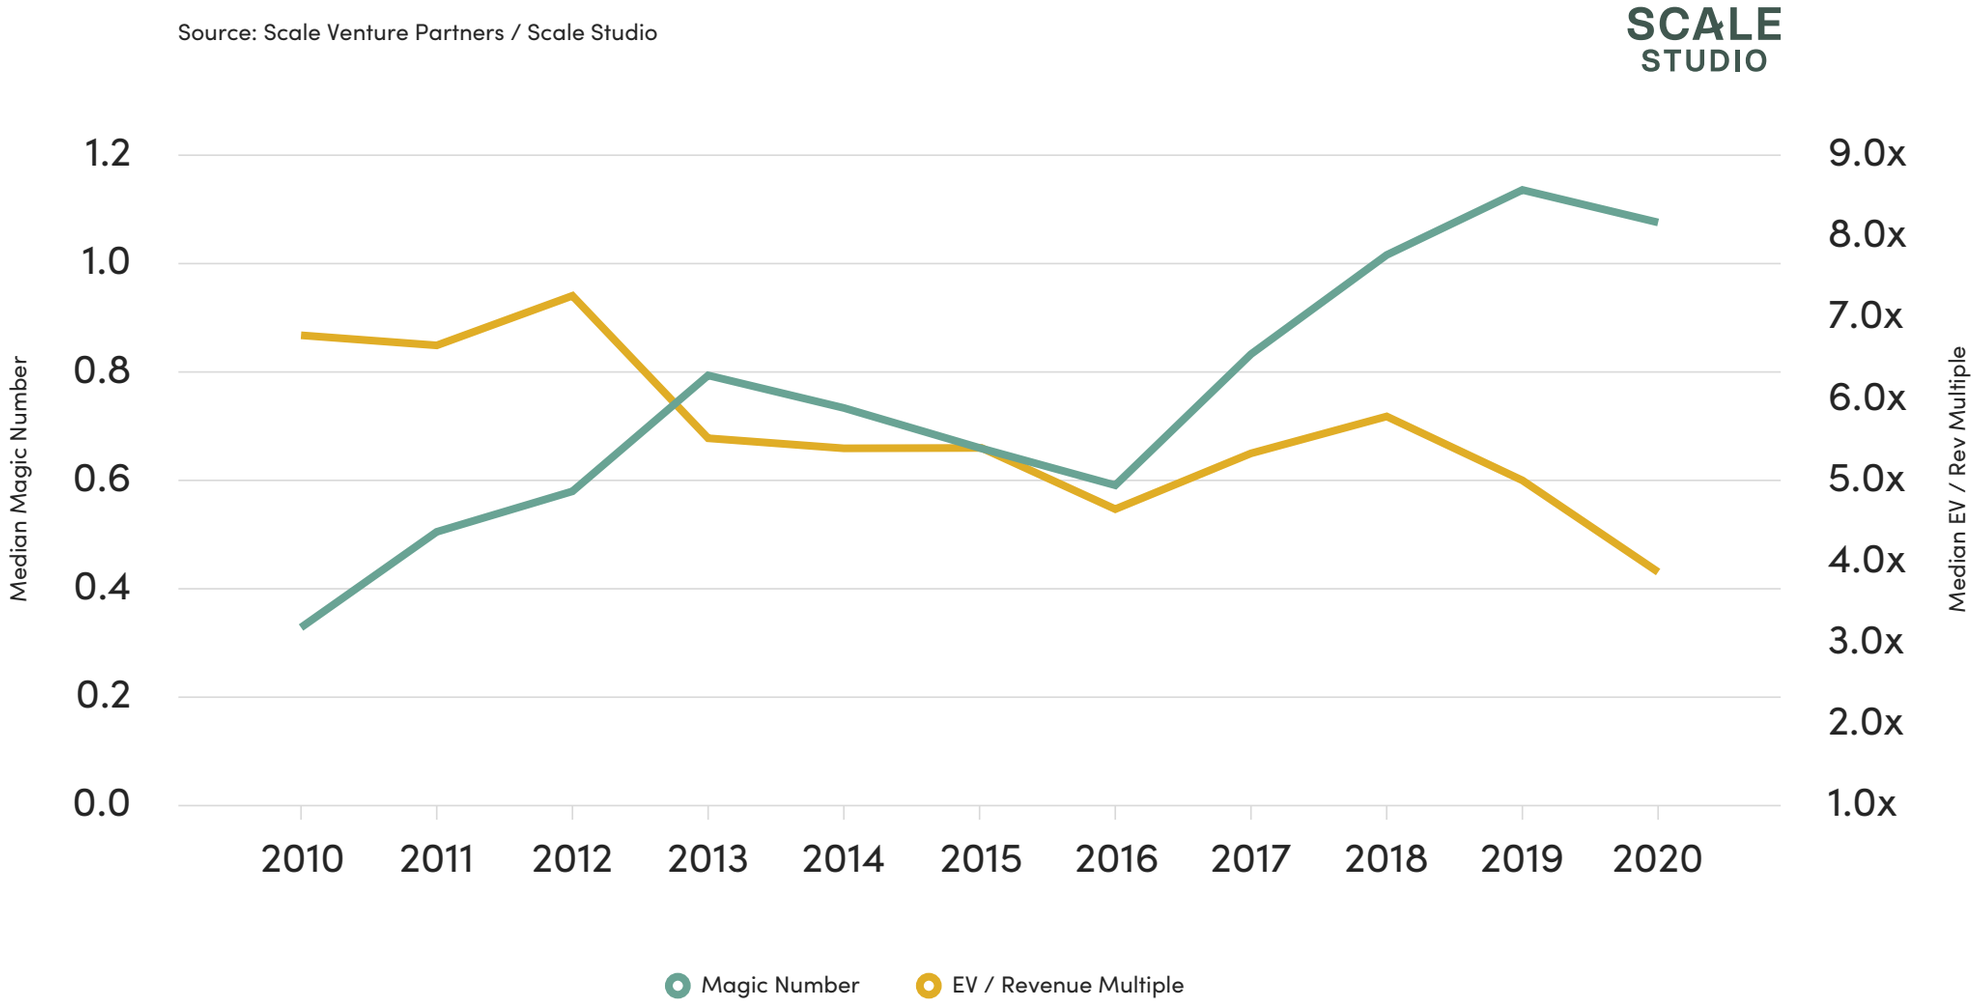

Analyzing Ten Years of Data on Private and Public SaaS

In the early 2000's, SaaS and cloud-based computing were still nascent concepts and poorly understood by most of the business world. Salesforce and Amazon Web Services (AWS), which have become the two dominant players in the SaaS application and cloud computing universes, were not launched until 1999 and 2002, respectively. Sure, there had been prior attempts at subscription-based models in other areas within the wider technology ecosystem (e.g. telecommunications and cable providers), but the advent of players like Salesforce and AWS marked the first real shift from CapEx to OpEx spend and forced the broader market to understand how these business models should be evaluated.

Stepping back, it is remarkable to see how far we have come in the past two decades. What this progress has given us is more than 20 years worth of data on how this ecosystem has performed and evolved over time, all the way down to the company level. And while progress within SaaS and cloud has been lightning fast as compared to other corners of the economy, these advancements have not happened without immense amounts of capital investment. Many of today's SaaS and cloud businesses have raised hundreds of millions of dollars on the path to IPO.

As early stage SaaS investors, we are always trying to understand how the ecosystem is changing. How are things different now versus last quarter, versus last year, and five years, and a decade ago? What does it take to build an category-leading SaaS business today compared to the past? Do any patterns emerge?

To better understand this dynamic, we wanted to start by looking closely at how capital investment within the SaaS universe has changed over time. To do so, we created a data set of private company data from Scale Studio and our Scale SaaS Index of public companies in order to track two data points:

Magic Number

A measure of go-to-market efficiency that broadly reflects "return" on capital investment.

EV / Revenue multiple

A measure of point-in-time valuations over the past decade.

It's important to note that for this analysis, we specifically took aim at how Magic Numbers in the private universe are affected by revenue multiples in the public universe as we believe there is some pull-through effect from publics to privates.

The Past Decade of Private Company Magic Numbers

Scale's rule of thumb is that a Magic Number of 0.7x is a fairly healthy efficiency baseline for most SaaS businesses. That means for every dollar invested in Sales & Marketing, the company generates $0.70 of revenue after the first year. Within these parameters, so long as a customer's lifetime value exceeds the amount of time it takes the company to recoup its investment in S&M, this dynamic tends to lead to extremely profitable unit economics in the long run.

As we look at our SaaS database over the past decade, the long-term median Magic Number has hovered around 0.7x. But to say that this benchmark has held steady over time isn't necessarily true among private SaaS companies. While the median does hold around 0.7x, it has bounced around quite a bit and most recently has been trending down.

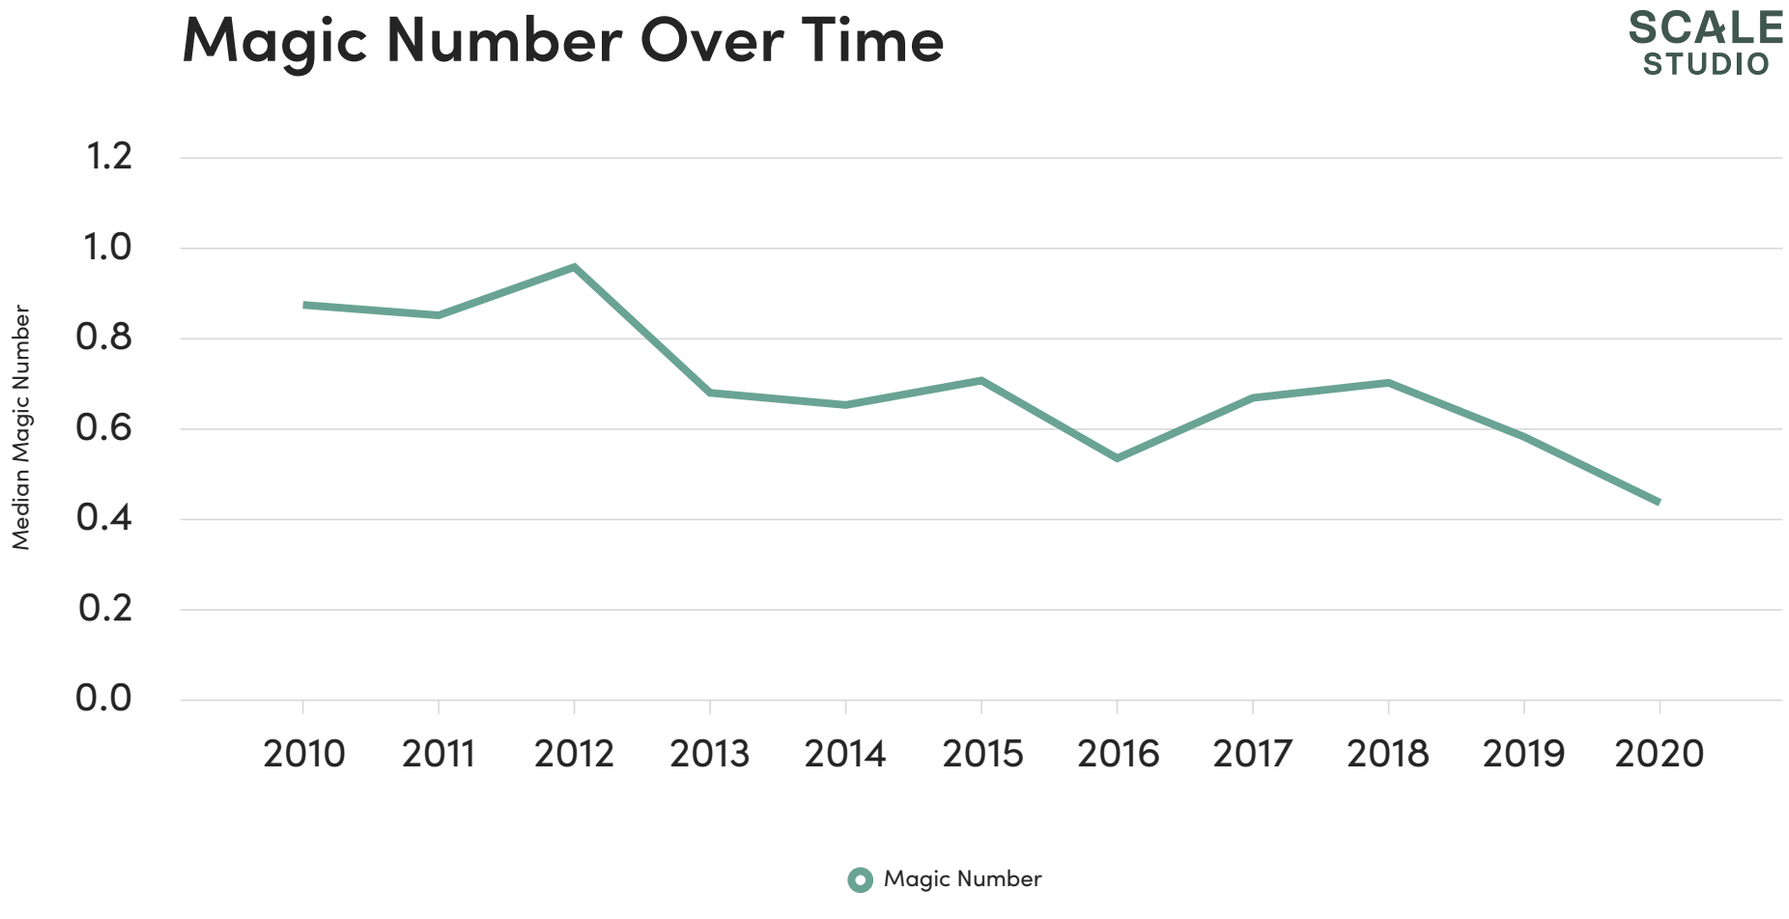

Bookmark that thought for now, we'll get back to it in a moment.

The Last Decade of Public Company Revenue Multiples

A few years ago, Scale partner Alex Niehenke wrote a piece on long-term SaaS valuations, in which he plotted how EV / Revenue multiples have changed over time. We took this same chart (shown below) and updated the revenue multiples to match the same time period as our Magic Number chart above.

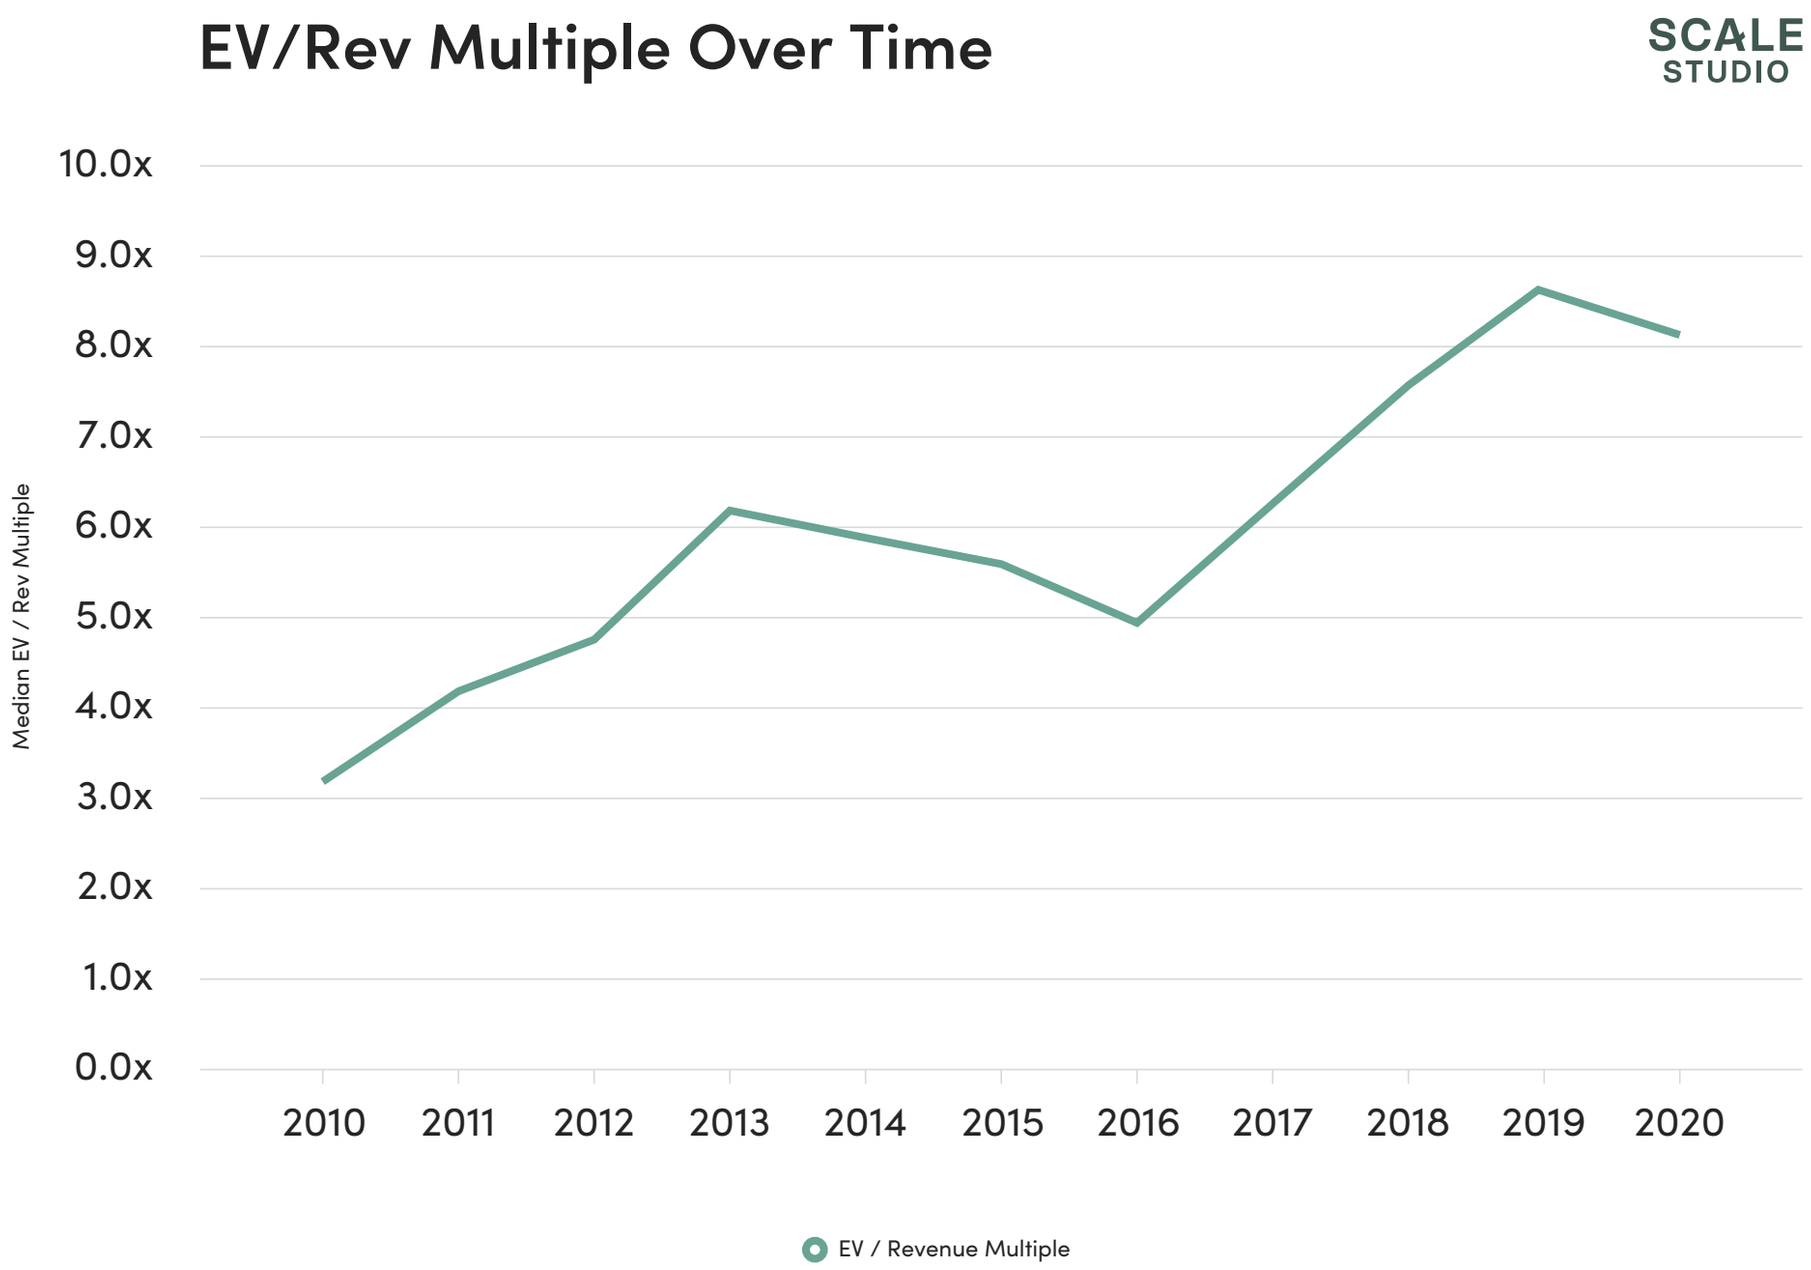

It's safe to say, as of this writing, we're well above the long-term median SaaS revenue multiple of 5.0x. In fact, as of the close of the most recent quarter, we hit a high of 11.7x EV/revenue — which points to the market's continued exuberance for SaaS and cloud. Investors are betting on the future growth of this segment of the market, in no small part because of its decades-long streak of consistent top-line growth. Looking ahead, that which has not digitized soon will. And COVID-19 has only accelerated that trend.

Magic Number x Revenue Multiple

We've seen fluctuations in both public SaaS revenue multiples and private SaaS sales efficiency. The big surprise about these metrics comes when looking at these two trend lines plotted on the same chart. There is a correlation between the two, and it's negative.

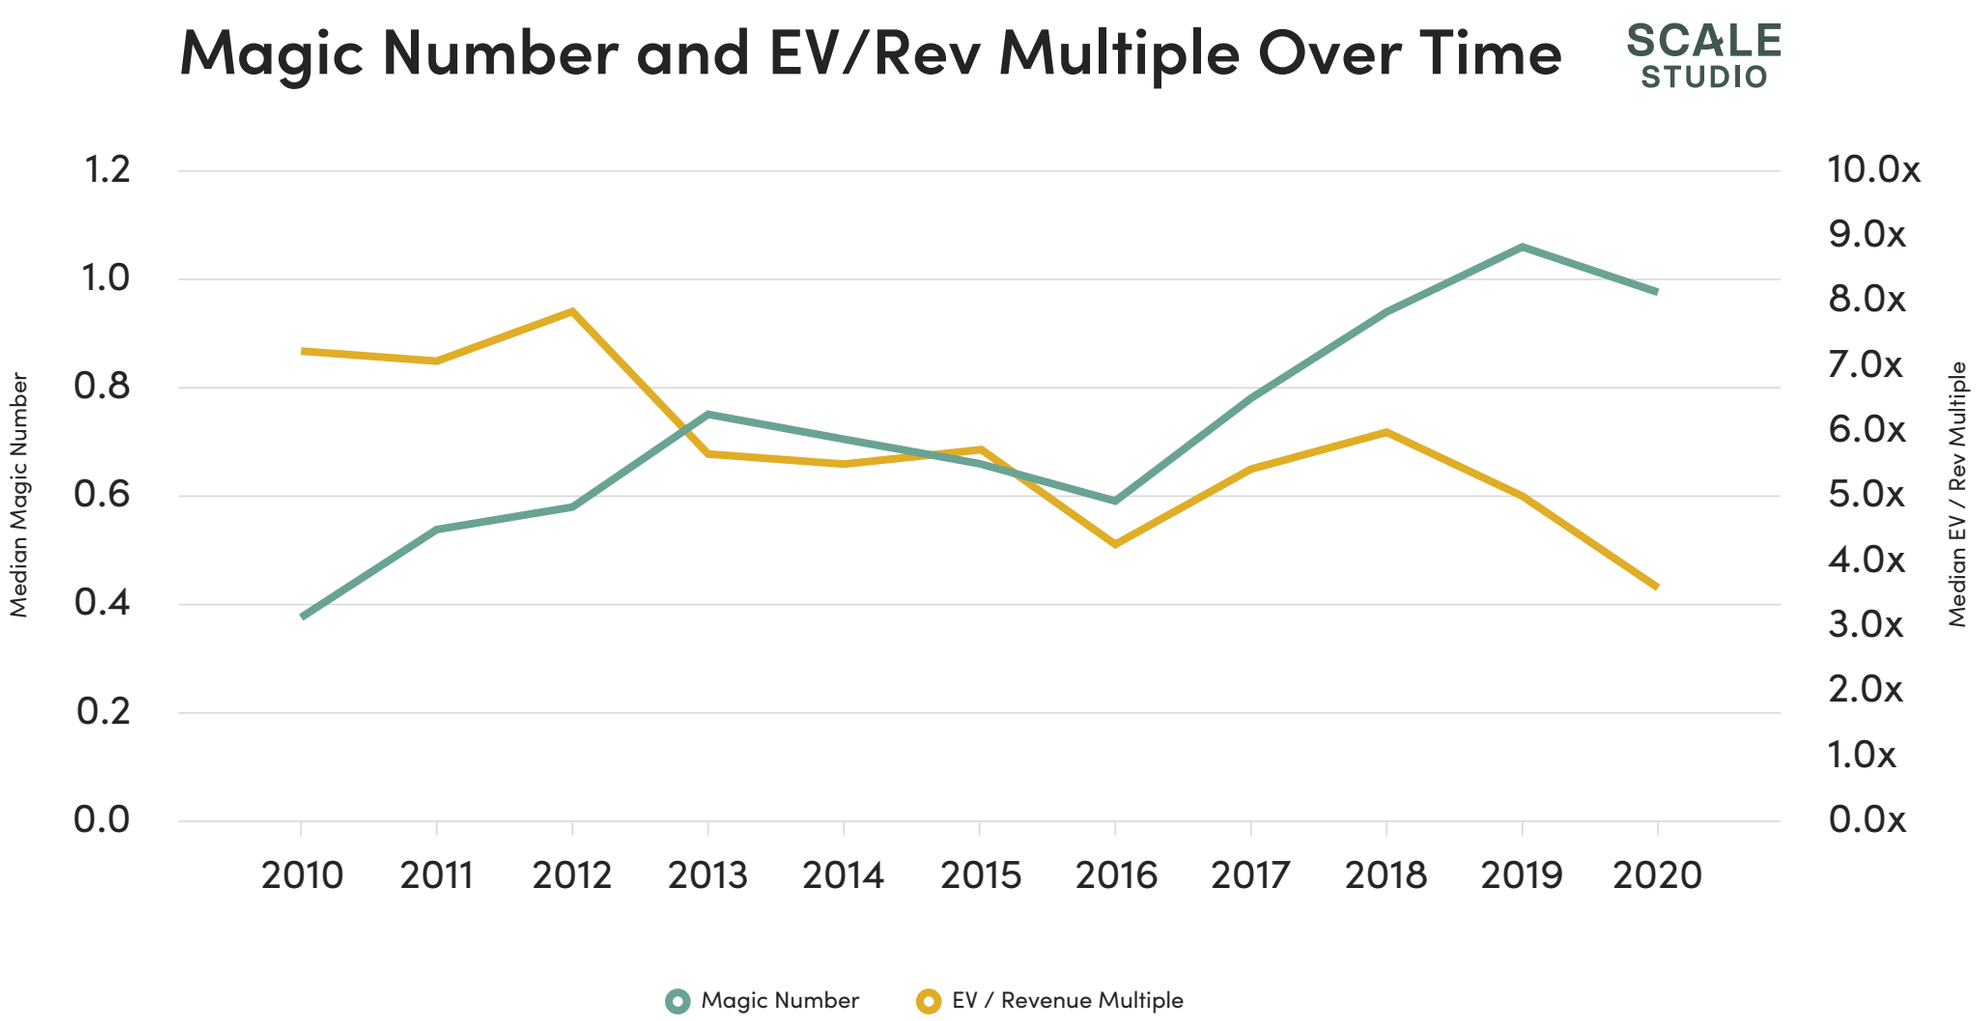

In general, the higher the multiple in the public markets, the lower the median sales efficiency in the private markets.

What to Make of This?

Why does the data show an inverse relationship? Here is one theory.

For many venture-backed startups, the siren song of an open IPO window with high valuation multiples pushes pre-IPO companies to step on the gas to achieve critical mass as quickly as possible. Some are ready and can maintain (or even improve) their go-to-market efficiency as they continue to scale. These are the success stories that you hear about, but they're few and far between.

It is unfortunately more common that early stage SaaS businesses try to scale prematurely only to see efficiency quickly degrade due to stresses on shaky processes, lack of real product- market fit, or heavy competitive pressure. These companies may achieve the growth needed to reach escape velocity, but they will burn through quite a bit of cash on their journey to get there.

The lesson here for founders is this:

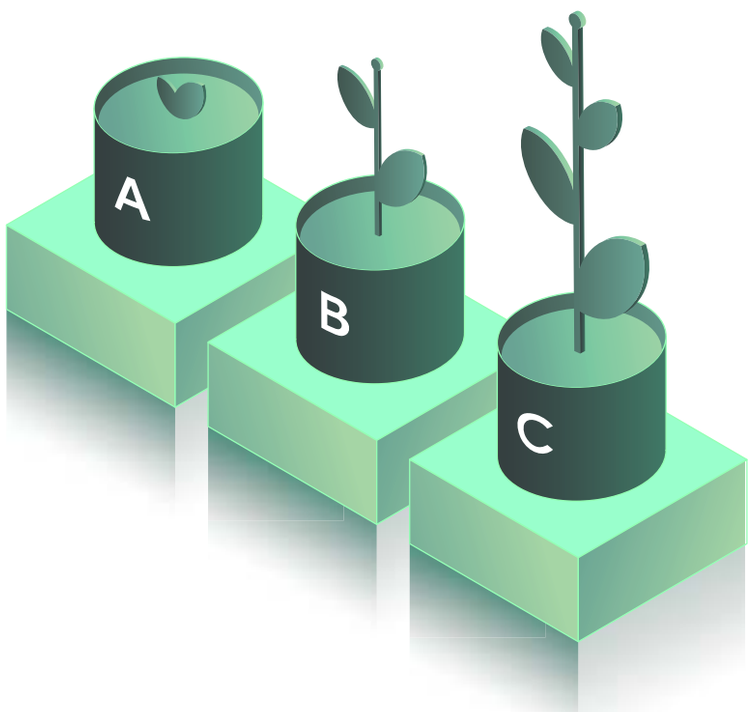

In an environment where capital is cheap and easily accessible, SaaS businesses can continue to finance their growth round after round.

Early- and mid-stage startups can tap a growing pool of venture investors and, more recently, a myriad of cross-over and corporate venture investors trying to take advantage of the SaaS market's momentum.

The risk is that when capital gets more expensive or a company starts preparing for a public market debut, sloppy (inefficient) growth can carry consequences. In these instances, management needs to quickly implement a strategy to get back to a highly efficient go-to-market model, lest the company face the discount that late-stage or public investors apply to cash hungry, inefficient go-to-market models.

Churn

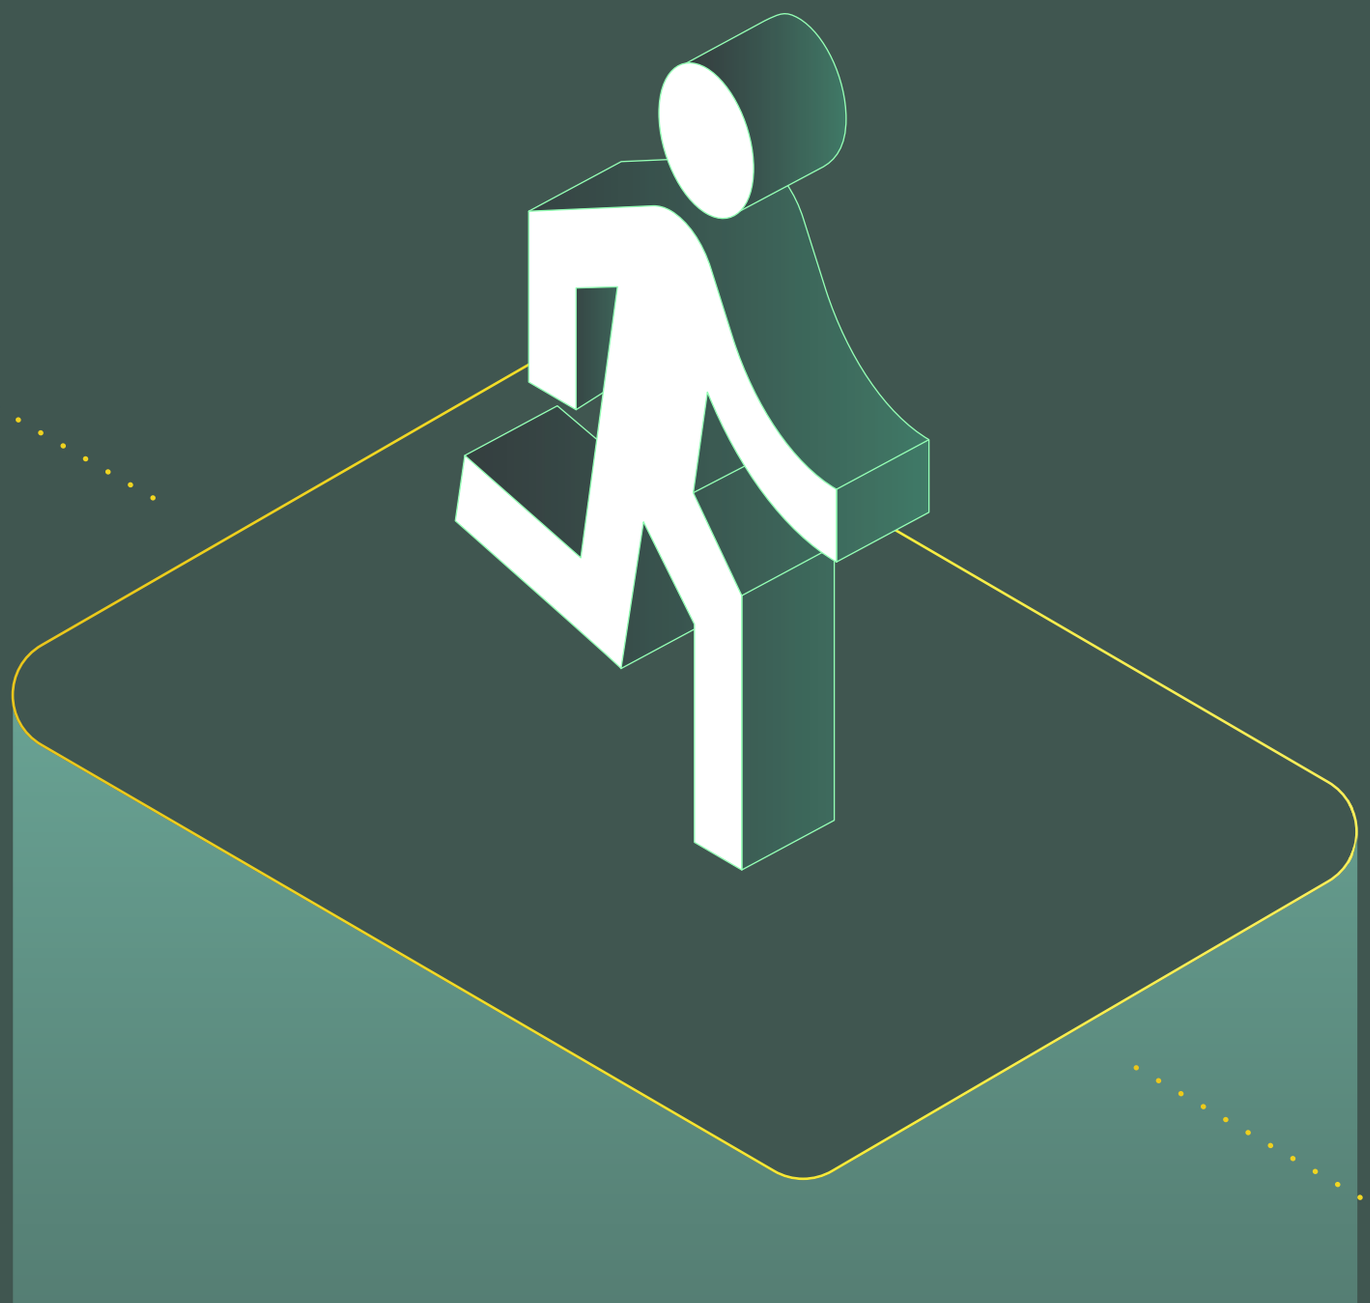

Less is more. As growth naturally slows with scale, churn sets the upper limit on how big a company can become.

To Understand Customer Retention, Start with Churn

You're at a networking event when someone shares that one of your competitors has "Churn of 5%. They're killing it." Your heart sinks because you know your company's churn is worse. Far worse.

Before you call an all hands meeting, though, ask yourself how that 5% was calculated. There's a good chance that the churn rate calculation the competition used is quite a bit different from what your company is using.

Churn is one of those concepts that everyone in SaaS understands, yet there's no real convention or common definition for how it is calculated and used inside a company. Measured correctly, churn is a key tool in your SaaS metric toolbox, effective at monitoring lost revenue and the first place you're likely to spot a problem with customer retention.

That "measured correctly" part is the focus of this article. We'll look at two ways that time impacts the accuracy of the churn rate calculation. It's the first in a series of deep dives into churn and retention, building on the work done on Scale's Four Vital Signs of SaaS framework.

What Churn Does and Does Not Measure

Existing customers are a growth engine. They tend to buy more over time (as long as you have more to sell them!), adding an important tailwind when your revenue base grows larger and your growth rate slows.

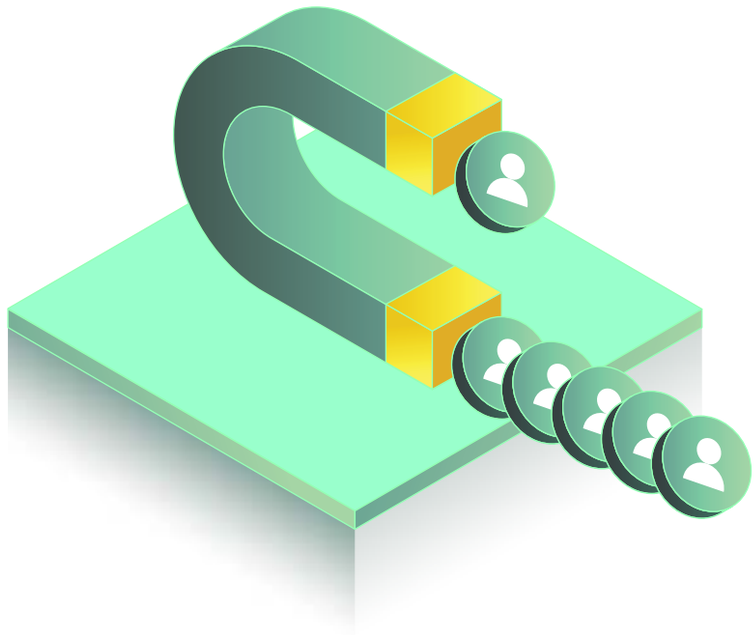

Looking at it another way, the cost to secure a new customer is 5 to 25 times the cost to retain a customer.

Scale partner Rory O'Driscoll published an article on how Scale evaluates SaaS company performance called the Four Vital Signs of SaaS. The Vital Signs are the four metrics that we've found to provide the most reliable signals about the health of a subscription revenue business. It's what our team uses internally during due diligence and to help our portfolio companies navigate go-to-market and beyond.

The Vital Signs include a single metric in each of four areas: growth, efficiency, retention, and burn rate. They're a starting point because:

- They aren't meant to tell you everything (but they will highlight when and where deeper analysis is needed)
- They are meant to facilitate company-wide alignment on performance (easy to measure, track, and explain)
- They allow comparison across companies (making meaningful benchmarking possible)

Churn is a bit of a special case among the Vital Signs:

> As growth slows the impact of churn escalates and provides an upper bound on HOW BIG a company can become. The simplest way to measure Gross Churn is by taking Churned ARR and dividing it by opening ARR for the period...[but] we acknowledge that this metric is a horrible oversimplification.

To see why selecting any one measurement standard for churn is an oversimplification, click through to this Key Bank research report on how public SaaS companies measure churn and retention. Of the nearly 100 companies listed, there are nearly 100 different measurement standards. By the time a company reaches that magic $100M ARR where an IPO is possible, they have (or should have) determined the best way to adapt the standard churn rate formula to their business model.

PagerDuty	**Total ARR Churn Rate**	**$ ARR Lost**	**Customers contributing revenue in t-1 but no revenue in t** **t**	**<5%**	Dollar-based churn excluding the benefit of upsells, based on ARR "Our ARR churn rate represents revenue from customers that contributed no revenu in the measurement period but did contribute revenue in the equivalent prior year period" (424B4 filed on 4/11/19, Page 55)
		$ ARR	**All customers t-1 end** **t-1 end**		

SaaS companies at scale will often use a company-specific formula for measuring churn and retention.

Churn is a tricky thing to measure because of two related aspects of time: measurement period and contract length. Let's look at how each of these complicates the picture.

Time to Recalculate Churn

Churn tells you about lost business during a specific time period. It is related to customer retention rate, which tells you how effectively you held onto the customers who were eligible to renew (or cancel) during a time period. As we'll see, churn and retention aren't precisely the inverse but it's ok to think of them that way as long as you understand the math behind how they're calculated. So let's get into that.

The Vital Signs Gross Churn calculation is straightforward:

$$\frac{\text{Churned ARR (quarter)}}{\text{Beginning ARR (quarter)}}$$

Here's a hypothetical ARR waterfall of an early-in-revenue SaaS company with its churn rate highlighted.

ARR Waterfall (Quarterly) : Hypothetical Saas Startup

($000's)	Q1:19	Q2:19	Q3:19	Q4:19
Beginning ARR	$ 1,900	$ 2,900	$ 3,925	$ 7,225
New Customer ARR	$ 600	$ 750	$ 3,000	$ 1,200
Upsell/Expansion ARR	$ 500	$ 500	$ 600	$ 900
Gross New ARR	$ 1,100	$ 1.250	$ 3,600	$ 2,100
Churned ARR	$ (100)	$ (225)	$ (300)	$ (500)
Net New ARR	$ 1,000	$ 1,025	$ 3,300	$ 1,600
Ending ARR	$ 2,900	$ 3,925	$ 7,225	$ 8,825
Gross Churn	-5.3%	-7.8%	-7.6%	-6.9%

This company started Q1 with $1.9M in ARR, then lost $100k of ARR during the quarter, so its Gross Churn rate for the quarter was -5.3% — it lost 5.3% of its ARR base during that time period. Right?

That depends. We don't know for sure because we don't know the contract period the company offers its customers. One common mistake with churn rate is mismatching measurement period (month, quarter, full year) and contract length. To bring it back to the sample company, assume they offer monthly contracts. That means a quarterly Gross Churn calculation jumbles together 3 separate cancellation windows, all without adjusting the ARR in the denominator.

In general, any combination of mismatched measurement periods and contract terms will give you a lot less signal from your churn metric. A quarterly calculation for a SaaS business with multi-year contracts, to give another example, will tend to understate customer defections during periods when they can't actually leave.

Let's go back to the model company's ARR Waterfall stated by month. Assume that it offers month-to-month contracts, which means the full ARR base is eligible to continue doing business or cancel every month.

ARR Waterfall (Monthly) : Hypothetical Saas Startup

($000's)	Jan 19	Feb 19	Mar 19	Apr 19	May 19	Jun 19
Beginning ARR	$ 1,900	$ 2,075	$ 2,400	$ 2,900	$ 3,400	$ 4,025
New Customer ARR	$ 100	$ 200	$ 300	$ 300	$ 450	$ -
Upsell/Expansion ARR	$ 100	$ 150	$ 250	$ 250	$ 250	$ -
Gross New ARR	$ 200	$ 350	$ 550	$ 550	$ 700	$ -
Churned ARR	$ (25)	$ (25)	$ (50)	$ (50)	$ (75)	$ (100)
Net New ARR	$ 175	$ 325	$ 500	$ 500	$ 625	$ (100)
Ending ARR	$ 2,075	$ 2,400	$ 2,900	$ 3,400	$ 4,025	$ 3,925
Gross Churn	-1.3%	-1.2%	-2.1%	-1.7%	-2.2%	-2.5%

The monthly Gross Churn data now accurately reflects the actual rate of customer churn. It also makes clear why the quarterly Gross Churn rate (from the first waterfall above) was misleading. The quarterly churn calculation was using an "outdated" Beginning ARR for 89 more days than it was accurate. By not accounting for new ARR, the denominator was smaller than reality and thus churn appeared worse than it was.

It's an easy fix: set your measurement period (monthly, quarterly, yearly) equal to your contract period (monthly, quarterly, yearly). By doing so, your Beginning ARR balance will always equal the amount of ARR up for renewal during the measurement period. This gets messy if you sell multiple contract lengths, but it's possible with a much more complicated spreadsheet.

Put another way:

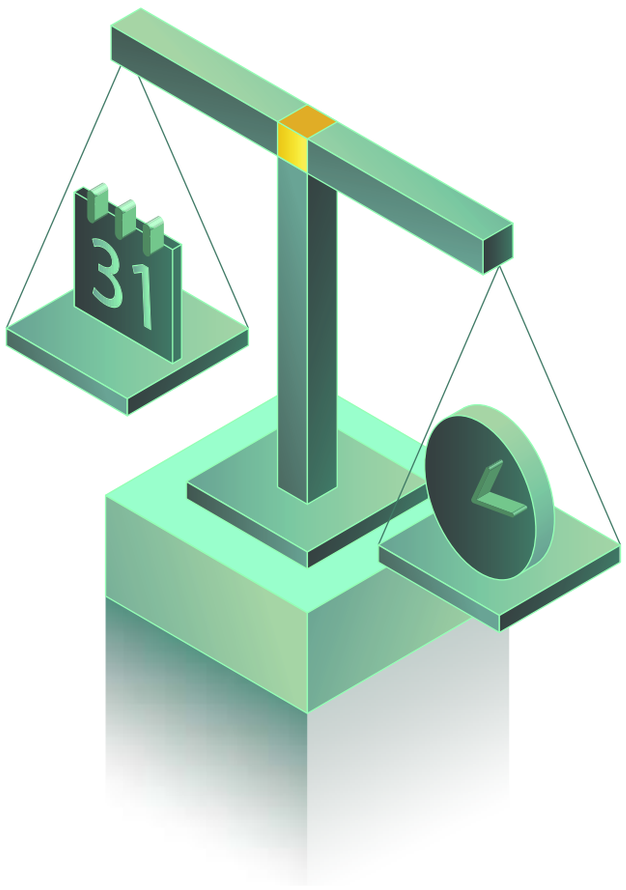

When your contract period equals your measurement period, your churn rate provides an accurate count of the $ disappearing from your ARR base during that period.

Getting More Signal from the Churn Metric

The previous chapter looked at what the Churn metric does really well and where its limitations lie.

This article looks at how to get the even more “signal” from the Churn metric. We’ll return to our hypothetical early-stage SaaS company, which has just posted a big quarter. By looking at the impact of revenue upside on Churn rate, we can launch into specific ways to smooth out the calculation and make it more reflective of the situation on the ground.

After all, the true value of tracking any given metric is how well it helps founders and teams make better decisions.

Adjusting Churn Rate for Big Quarters

You don't often hear about a SaaS startup having too big a quarter. But it's good to remember that quarters with big jumps in ARR will impact your Churn rate by inflating the denominator. It's obvious when you think of it, but in our experience not everyone remembers to think about it.

Why is this an issue? Often we see companies working with Churn data where the reporting period doesn't match up to the contract period. If the contracts are quarterly and they're measuring Churn quarterly, then this isn't an issue. But if the contracts are longer - say, a year - then the new ARR being added to the denominator can't churn for a year out. Refer back to the first article in the series for the discussion on aligning the Churn rate calculation to contract length.

A "big quarter" can mean a big jump in Beginning ARR balance at the start of the next quarter. This jump increases the Churn rate denominator (Beginning ARR) so that, all things being equal, the Churn rate declines.
A "big quarter" can also come from seasonality. A tax filing SaaS product, for example, might see a rush of new subscribers every 3Q as tax season kicks off. Software purchased by department heads who want to drain their discretionary budgets at year end can create a spike in 4Q sales.

There's another version of the "big quarter" problem that we see a lot in early-stage companies. There's a window of time after product-market fit — when both Product and Sales have hit their strides — but before nascent Customer Success teams have fully smoothed out their own processes. Companies see a lot of new customers sign up... then far too many leave soon after.

Back to our hypothetical SaaS startup for an illustration of how these variations on a “big quarter” impact Churn rate:

ARR Waterfall (Quarterly) : Hypothetical Saas Startup

($000’s)	Q1:19	Q2:19	Q3:19	Q4:19
Beginning ARR	$ 1,900	$ 2,900	$ 3,995	$ 7,385
New Customer ARR	$ 600	$ 750	$ 3,000	$ 1,200
Upsell/Expansion ARR	$ 500	$ 500	$ 600	$ 900
Gross New ARR	$ 1,100	$ 1.250	$ 3,600	$ 2,100
Churned ARR	$ (100)	$ (155)	$ (210)	$ (300)
Net New ARR	$ 1,000	$ 1,095	$ 3,390	$ 1,800
Ending ARR	$ 2,900	$ 3,995	$ 7,385	$ 9,185
Gross Churn	5.3%	5.3%	5.3%	4.1%

The quarterly view shows the strong 3Q with $3.4M in Net New ARR, bringing the 4Q Beginning ARR to $7.4M. At the same time, the Gross Churn rate improves from 5.3% to 4.1%. Things are really looking up!

The wrinkle here is this. The company’s 4Q Churned ARR could go as high as $(390) and still show a Churn rate in line with the first three quarters (5.3%). It would be easy to feel good about it. But that’s why you don’t rely on Churn alone to understand retention trends. We’ll leave this alone for now, but dive deeper into the solution when we look at cohort analysis in the next article in this series.

The takeaway here:

> Big changes in ARR balance have big impacts on Churn rate.
> All things being equal, a jump in ARR means a better Churn rate.
> But that may or may not actually be the situation on the ground.

Fortunately, there are two easy adjustments that can provide added precision in the Churn rate: Fixed-Based Churn rate and Annualized Churn.

Fixed-Base Churn Rate

Fixed-based Churn involves anchoring on a specific ARR balance at a specific point in time. It's a little counterintuitive: you get more precision by making fewer adjustments. The key is to look at the resulting periodic Churn rates cumulatively.

Here's an example that anchors on the Opening ARR balance at the start of the year and measures Churn rate quarterly. Each period's Churn correctly reflects the portion of the fixed Opening balance that churned. So by year end, the sum of all of the periods equals the annual amount of customer churn.

ARR Waterfall (Quarterly) : Hypothetical Saas Startup

($000's)	Q1:19	Q2:19	Q3:19	Q4:19	FY19
Beginning ARR	$ 1,900	$ 2,900	$ 3,995	$ 7,385	$ 1,900
New Customer ARR	$ 600	$ 750	$ 3,000	$ 1,200	$ 5,550
Upsell/Expansion ARR	$ 500	$ 500	$ 600	$ 900	$ 2,500
Gross New ARR	$ 1,100	$ 1.250	$ 3,600	$ 2,100	$ 8,050
Churned ARR	$ (100)	$ (155)	$ (210)	$ (300)	$ (765)
Net New ARR	$ 1,000	$ 1,095	$ 3,390	$ 1,800	$ 7,285
Ending ARR	$ 2,900	$ 3,995	$ 7,385	$ 9,185	$ 9,185
Gross Churn	5.3%	5.3%	5.3%	4.1%	40.3%
Fixed-Base Gross Churn	5.3%	8.2%	11.1%	15.8%	

This example uses a quarterly Churn rate but you can use the same approach to anchor on the Opening balance for a monthly calculation. Just reset the anchor point as you move into a new quarter and use that same anchor as the denominator for each of the quarter's three monthly Churn rate calculations.

Annualizing Single-Period Churn

You can easily create a projection of how the churn situation today will turn out if it continues for a full year by annualizing a single period's Churn rate.

The easiest way to annualize a single period's Churn rate is to multiply it by the number of periods in the year. So for a monthly Churn rate, multiply by 12. For quarterly, multiply by 4. The advantage here is this is a really easy calculation to make; the disadvantage is that you'll tend to overstate your Churn.

A slightly more involved approach works better: geometric compounding. The formulas here look like this:

For annualizing quarterly churn:

1-((1-Churn.Quarterly)^4)

For annualizing monthly churn:

1-((1-Churn.Monthly)^12)

Here is how that looks:

ARR Waterfall (Quarterly) : Hypothetical Saas Startup				
($000's)	Q1:19	Q2:19	Q3:19	Q4:19
Beginning ARR	$ 1,900	$ 2,900	$ 3,995	$ 7,385
New Customer ARR	$ 600	$ 750	$ 3,000	$ 1,200
Upsell/Expansion ARR	$ 500	$ 500	$ 600	$ 900
Gross New ARR	$ 1,100	$ 1.250	$ 3,600	$ 2,100
Churned ARR	$ (100)	$ (155)	$ (210)	$ (300)
Net New ARR	$ 1,000	$ 1,095	$ 3,390	$ 1,800
Ending ARR	$ 2,900	$ 3,995	$ 7,385	$ 9,185
Gross Churn	5.3%	5.3%	5.3%	4.1%
Annualized Churn Rate	19.4%	19.7%	19.4%	15.3%

Next Up: Cohort-Based Customer Retention

Tracking Churn rate — or any metric for that matter — isn't an academic exercise. The goal should always be to give your teams reliable information that helps them make better decisions. You're on the right track when your Customer Success and Marketing leads are working from the same data, speaking the same language, and working towards the same targets.

Diagnose Customer Loss with Renewal and Retention Rates

The first two chapters in this section looked at how to use the Churn rate metric to spot a customer retention problem, then presented adjustments to the Churn rate calculation that better capture your company's operational realities. This chapter looks at the metrics that you'll use to analyze churn issues: Renewal and Retention rates.

We'll start with an overview of Renewal and Retention rates, then looks at how to use each to examine different aspects of customer renewal behaviors. We'll end with some thoughts on how cohort analysis can help track customer expansion over time.

The Yin and Yang of Customer Retention: Churn and Renewal

Churn measures the volume of dollars or number of customers that your business lost during a given time period. Tracked over time, the Churn rate provides big-picture insight about customer loss. Go here to see churn benchmarks for SaaS startups as they grow from $1M to $100M.

The Renewal rate measures dollars or customers retained during a time period, limited to only those customers that were eligible to renew during the period. Renewal rate's superpower is that it gives you a snapshot of whether customer decisions to renew or cancel went for or against you.

Since we started this section focused on Churn rates, we can talk about Renewal rates in terms of Churn. When you set the period in which you are measuring churn equal to your contract length, your Non-Renewal Rate will equal your Churn rate. The Renewal rate, then, is simply the inverse, or (1-Non-Renewal Rate). So a Non-Renewal Rate of 7% equates to a Renewal rate of 93%.

From your first dollars of revenue, you should be tracking both Churn and Renewal rates. You'll share Churn data to coordinate Sales, Marketing, and Customer Success. And when (if!) Churn starts trending against you, you'll lean on Renewal rates to tell you a lot about "why?".

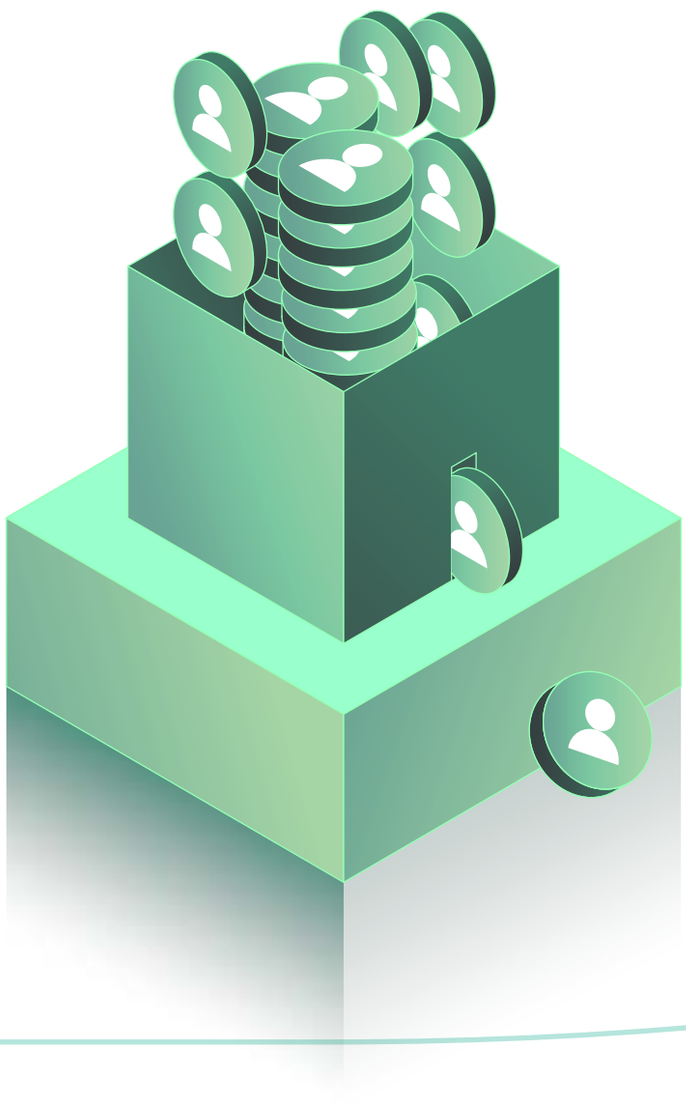

Defining Gross and Net Renewal Rates

Now for the definitions.

Gross Renewal Rate shows you the dollars that renewed as a percentage of all dollars that were eligible to renew during a period. It is calculated by adding the total renewed contract dollars then dividing by the total dollars that were eligible for renewal. That looks like:

Gross Renewal Rate = Dollars renewed/Dollars eligible to renew

Net Renewal Rate tells you about the growth of your existing customer base by including net expansion (upsell - churn - downsell). Eligibility is calculated using contracts that expired during the time period. That looks like:

Net Renewal Rate = (Dollars renewed + dollars expanded) /Dollars eligible to renew

Big picture: renewal rates are measured using the “end state” of a time period. It tells you how many customers/dollars renewed as a percentage of those that could have renewed. It’s your contract renewal batting average.

The Difference Between Renewal Rate and Retention Rate

If Renewal rate is a snapshot in time, Retention rate is a trend over time. Retention rates tell you about the renewal decisions of a group of customers that signed up at the same time and how many of those customers remained customers over time.

“Renewal” and “Retention” are often used interchangeably. They shouldn’t be. To analyze either requires you to group your customers into cohorts, or groups of customers that share specific traits like when they signed up or when they last renewed. How you define membership in those cohorts is what distinguishes the two:

- Renewal rate examines customer cohorts defined by the end date of their contracts
- Retention rate measures starting cohorts (customers that signed up during the same period) through the life of their contracts

That distinction has implications for what Renewal and Retention rates tell you and how to use them. Here’s how it shakes out:

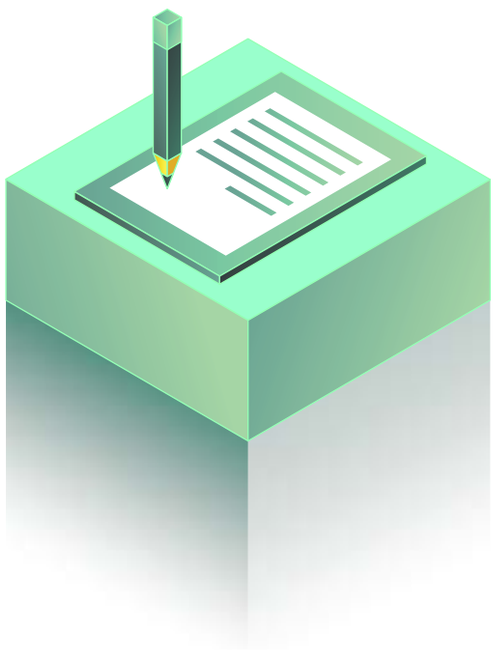

Renewal rate tells you how many customers re-signed at the end of their contract.

In practice that means listing out all the customers that renewed during a period, ignoring contract start date.

The Renewal rate is the best count (customer or $) of whether customers said yes or no when their contracts came up for renewal. Here's a look at how to set that up:

Renewal Rate: 1Q20					
Cohort = 1Q20 Renewals (ARR $K)	Renewal Date	Base ARR Eligible for Renewal	Base ARR Renewed	+Upsell -Downsell -Churn	Total net ARR Renewed
Customer 1	1/5/20	$ 750	$ 750	$ 75	$ 825
Customer 7	2/24/20	$ 500	$ 500	$ 100	$ 600
Customer 8	1/15/20	$ 1,250	$ 1,250	$ 50	$ 1,300
Customer 14	n/a	$ 650	$ –	$ (650)	$ –
Customer 17	3/3/20	$ 900	$ 900	$ 50	$ 950
Customer 22	1/29/20	$ 825	$ 825	$ (25)	$ 800
Total 1Q:20 Renewals		$4 4,875	$(,225	$4 400)	$,475
Renewal Rate (customer)	83.3%				
Gross Renewal Rate ($)	86.7% By definition, 100% max				
Net Renewal Rate ($)9	1.8% > 100% shows strong customer expansion				

Retention rate shows you how well you're keeping customers that signed up during the same period. The Retention rate shows how customers that started at the same time buy more and/or increase adoption. On a dollar basis, it shows how much an initial $1 in sales has grown after a specific amount of time. So, for example, how long it takes customers that signed up in the same period to grow from $100K to $150K in ARR.

Here's a simplified view of retention tracking. It illustrates the common early-stage challenge of retention dollar decline caused the impact of customer loss on a small base:

Retention Rate: FY19				
Cohort = 1Q20 Signups (ARR $K)*	1Q19	2Q19	3Q19	4Q19
Customer 1	$700	$725	$750	$900
Customer 2	$400	$450	$500	$700
Customer 3	$1,250	$1,250	$1,250	$1,250
Customer 4	$300	$300	$300	$0
Customer 5	$600	$700	$800	$1,000
Customer 6	$800	$825	$900	$1,100
Retention (customer)	100.0%	100.0%	100.0%	83.3%
Gross Retention ($)	100.0%	100.0%	100.0%	92.9%
New Retention ($)	100.0%	104.9%	105.9%	110.0%

A few notes on the table. These six customers signed annual contracts in 1Q19; that column shows the initial contract value. Expansion during a quarter represents upsell like adding seats or services to the existing contract. Finally, we're measuring dollar retention using the starting 1Q19 contract value rather than adjusting to the beginning ARR each quarter (though that calculation too can give you a useful view of the same data).

SaaS Customer Retention Tracking

For SaaS businesses, it's the trend that often matters most. And, by definition, trends become clearer with time and data.

Enterprise-focused SaaS companies typically sell annual contracts, which means they don't see much signal from renewal tracking until lapping the first batch of renewals. And even then, the renewal decisions of those first customers tend to be outliers that don't predict the longer-term trend.

On the other hand, don't get lost in the data until you have accumulated enough of it to reach reliable conclusions. We see a lot of early-stage companies adding complexity to their renewals tracking too early.
It's better to start simple and build.

In the early innings, you can get by with a spreadsheet consisting of columns for each customer (logo) and it's first quarter of ARR contribution. Then add columns each quarter to track each customer's ARR progression.

But you have to start somewhere, so start with the following two approaches to get initial insight into your renewal situation. Consider them placeholders until enough time has passed for you to accumulate more data.

CURRENT QUARTER PROXY

When you don't have much historical data, use your current quarter as the proxy for your forward projections. That looks like this:

Renewal Rate: Current Quarter Proxy				
	Current Qtr (Actual)	Q+1 (Estimate)	Q+2 (Estimate)	Q+3 (Estimate)
Eligible to renew	100	120	135	150
Renewals	85	102	115	128
Renewal rate	85.0%	85.0%	85.0%	85.0%

Just don't project out too far — at even the best performing startups, Renewal rates bounce around from quarter to quarter in the early innings.

Once you have a few quarters of renewal rate data, you can start using a weighted average in your planning and projections.

WEIGHTED AVERAGE

Grab as many periods of renewal data as you have, then average them to generate an estimate for your current Renewal rate.

Renewal Rate: Weighted Average				
Last year's actuals	Current Q (Estimate)	Q+1 (Estimate)	Q+2 (Estimate)	Q+3 (Estimate)
Eligible to renew	100	122	135	150
Renewals	87	107	119	133
Renewal rate	87.0%	87.7%	88.1%	88.7%
This year's estimate	Q4 (Actual)	Q3 (Actual)	Q2 (Actual)	Q1 (Actual)
Eligible to renew	165	175	190	200
Renewals	145	154	167	176
Renewal rate	88.0%	88.0%	88.0%	88.0%

Weighted averages won't provide insight into the impact from things like new product features (which may or may not meet customer needs), new Marketing programs like content campaigns (which may or may not provide value to customers), or new Customer Success initiatives (which may or may not provide value to your customer segments). But it's a good stopgap until you accumulate more data.

Cohort Analysis

What you're working towards with Renewal and Retention tracking is cohort analysis. A cohort analysis has two steps: grouping customers into cohorts defined by beginning or ending traits, then tracking how those cohorts behave over time. The analytical gold is seeing how well your company is expanding its customer relationships by, for example, looking at how $1 of initial ARR grows over time.

This is important because great cohort expansion relieves some pressure from your sales efficiency. When customers consistently become more valuable over time, the cost to acquire those customers is no longer the sole metric that determines ROI on selling. Net dollar retention growth captures not just sales efficiency, but customer success upselling and even market-ing's impact on increasing usage and expansion.

Burn

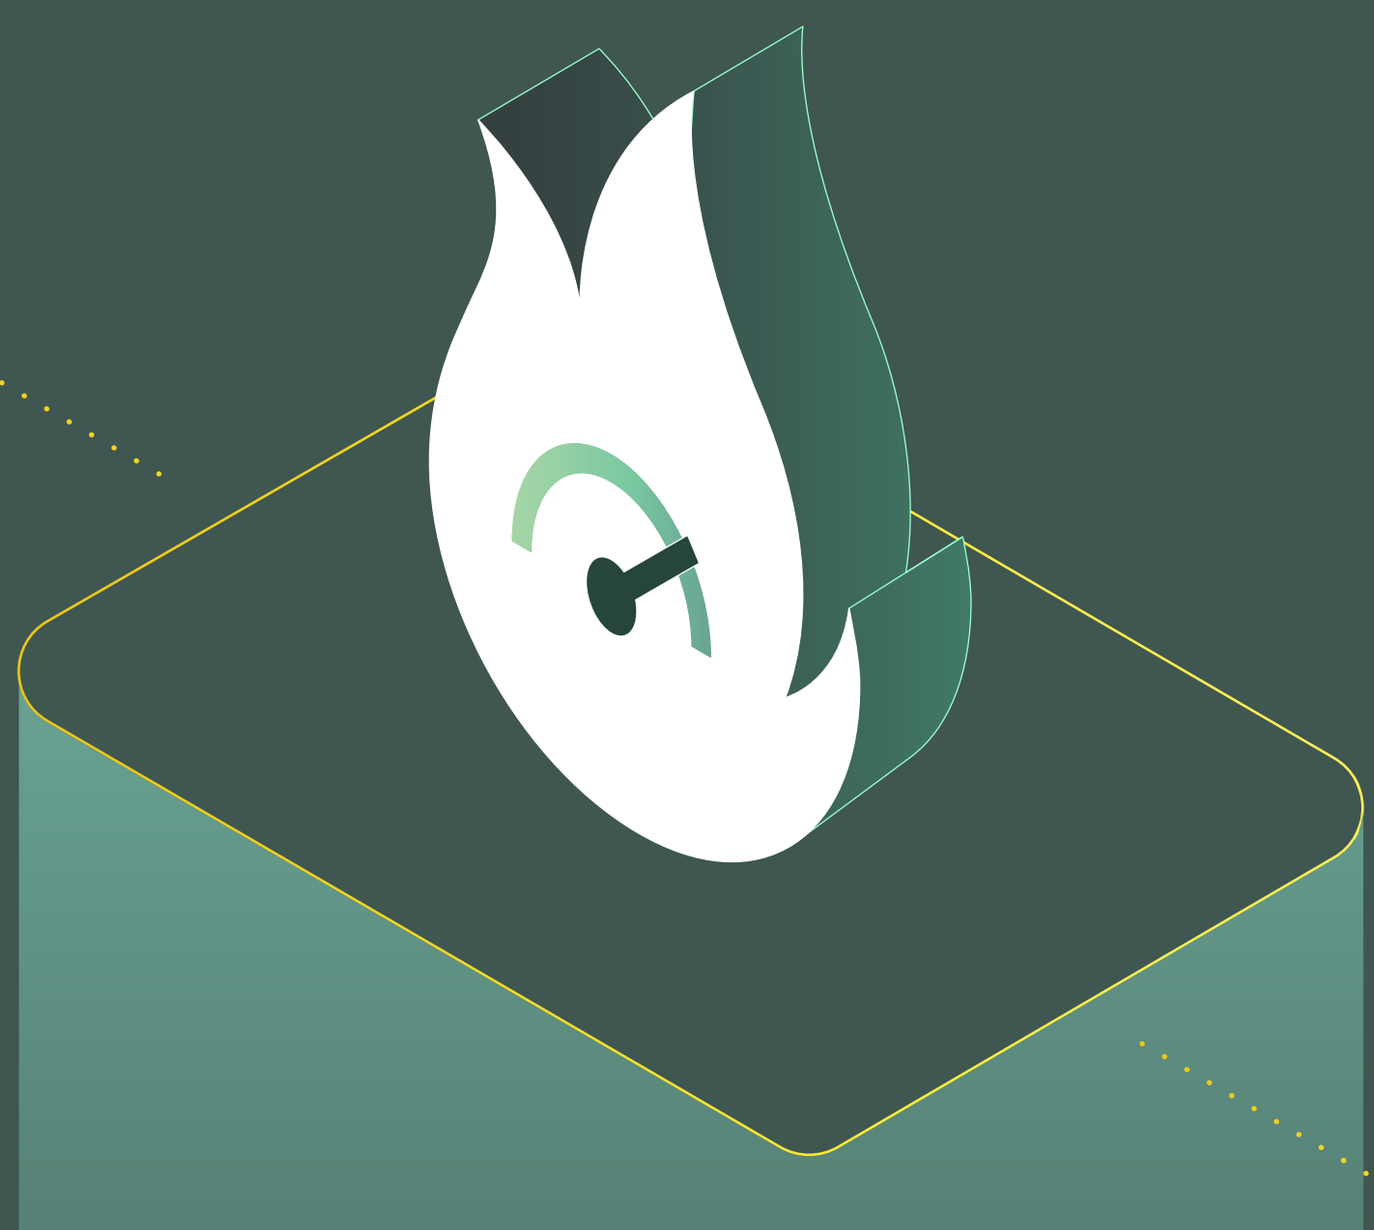

You can't scale without cash. Managing burn is critical for ensuring your runway is long enough to meet your goals.

A Founders Refresher on SaaS P&L

Building a SaaS product and founding a SaaS company do not require mastery of the Profit & Loss (P&L) statement. Leading a successful venture-backed SaaS startup does.

For subscription software businesses, Annual Recurring Revenue (ARR) is the leading indicator for growth — but building a company to go all the way to IPO also requires fluency in Generally Accepted Accounting Principles (GAAP). GAAP accounting is the basis for the financial data that public companies release every quarter.

The following is a founder-focused overview of the GAAP P&L and the special requirements it has for subscription-based software companies. It kicks off a multi-part look at measuring and improving cash burn rates. We'll start with the fundamentals then go very deep into how to use this information to run a data-driven company.

SaaS Profit & Loss Overview

The Profit & Loss statement is one of three financial statements that a typical corporation generates on a periodic basis (monthly, quarterly, annually). It is more formally known as the Income Statement in the trio of core financial statements that also includes the Balance Sheet (the table of assets, liabilities, and shareholder equity) and the Statement of Cash Flows (a cash-oriented view of the business).

The basic SaaS P&L has a handful of key categories:

REVENUE

Sales for the period, often broken into two or more sub-categories that track different products and/or services. Note that this is revenue as defined by GAAP and not ARR/MRR.

COST OF REVENUE

The costs of serving a company's SaaS product(s) and servicing its customers, primarily employee and technology infrastructure costs. Also called Cost of Sales (COS) or Cost of Goods Sold (COGS).

GROSS PROFIT

Revenue minus Cost of Revenue. Gross Margin % is calculated using this number.

OPERATING EXPENSES

Costs related to ongoing operations, including three main main sub-categories:

- Research & Development (product development)
- Sales & Marketing
- General & Administrative (e.g., HR, IT, Finance, real estate, and other administrative costs not otherwise attributed to R&D/S&M)

OPERATING PROFIT/(LOSS)

Gross Profit minus Operating Expenses.

OTHER INCOME/(EXPENSE)

Non-operating expenses like interest income or expense, income taxes, expenses associated with discontinued operations, etc.

NET PROFIT/(LOSS)

Operating Profit (Loss) minus Other Costs.

The basic reporting structure doesn't change much from startup to public company. As you can see in Zoom's first quarter 2020 income statement:

Zoom Video Communications, Inc.
Condensed Consolidated Statements of Operations
(Unaudited, in thousands, except share and per share amounts)

	Three Months Ended April 30	
	2019	2018
Revenue	$ 121,988	$ 60,070
Cost of revenue	24,104	11,660
Gross profit	97,884	48,410
Operating expenses:		
Research and delevelopment	13,783	6,264
Sales and marketing	64,041	36,261
General and administrative	18,503	7,569
Total operating expenses	96,327	50,094
Income (loss) from operations	1,557	(1,684)
Interest income, net	658	436
Other income, net	513	5
Net income (loss) before provision for income taxes	2,530	(1,243)
Provision for income taxes	(316)	(97)
Net income (loss)	2,214	(1,340)
Undistributed earnings attributable to participating securities	(2,016)	—
Net income (loss) attributable to common stockholders	$ 198	$ (1,340)
Net income (loss) per share attrubutable to common stockholders:		
Basic	$ 0.00	$ (0.02)
Diluted	$ 0.00	$ (0.02)
Weighted-average shares used in computing net income (loss) per share attributable to common stockholder:		
Basic	109,708,898	80,623,861
Diluted	136,428,379	80,623,861

Revenue vs. ARR and MRR

You'll notice that ARR and MRR are nowhere to be found in the P&L. Though everywhere in SaaSland, ARR bookings are an industry metric and not officially recognized under GAAP. GAAP is revenue focused, and requires that most SaaS companies recognize annual contracts in monthly chunks ("ratably over time"). ARR / MRR bookings are absolutely critical measures of growth and the basis for the four key operating metrics. But when it comes to P&L accounting, GAAP is the focus for its utility showing you line-by-line how you are performing as a business.

Also keep in mind that there is a lot of variation in accounting practices from company to company, even among traditional, annual contract enterprise SaaS and especially for SaaS pricing models with service offerings or usage-based / metered pricing. GAAP provides guidance on how to broadly account for and recognize revenue.

What's most important for a founder to know is that your finance team should account for and report each category of income and expense in accordance with prescribed guidelines under GAAP.

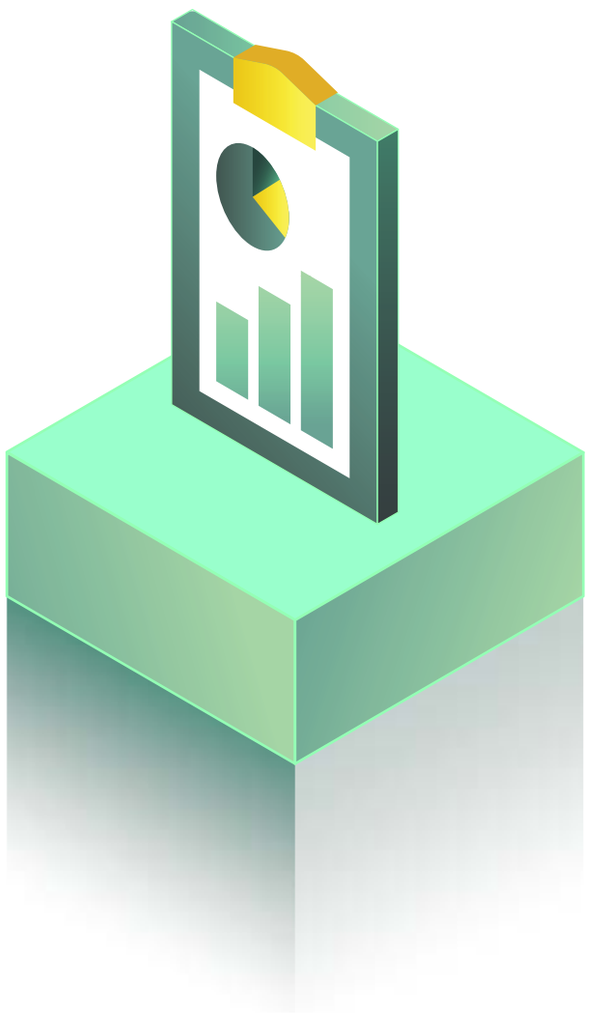

The big picture advice here is to emphasize good bookkeeping practices from the beginning then remain consistent over time so your performance metrics have as much historical perspective as possible.

The following chart shows how a single $100,000 annual SaaS contract could be recognized in different financial contexts: MRR, ARR, GAAP revenue recognition, and cash. We've assumed that the customer makes bi-annual cash payments of $50,000, and shown GAAP revenue recognition quarterly. ARR and MRR are shown as one-time additions to the company's ARR/MRR bookings pool.

Gross Margins

You'll often hear that a well-run SaaS company should have 80%+ Gross Margins. That's a fine rule of thumb, but there's much more to the story.

First, the expenses that go into that 20% of revenues can vary widely. Think of your Cost of Revenue as what you spend to keep your software running and your customers happy. That will generally includes a few big categories:

- Technology costs like cloud hosting, servers, and related infrastructure spending.
- People costs like the salary/wages, benefits, and payroll taxes of employees that maintain your software. That does not include product development costs, which go into R&D operating expenses (as described below).
- Other technology and people costs like the cost to support your product, including technical customer support, onboarding costs, and retention-focused customer success.

Note that customer success team members focused on renewals, up-sell, and cross-sell are considered Sales & Marketing expense (in OpEx, not COR). In practice, most SaaS Customer Success programs have CSMs devoting some portion of their time to renewals and upsells; even a small portion of time spent on these activities means that those expenses should be allocated to Sales & Marking in OpEx.

Not every SaaS company can or should hit the “80% rule.” Consider the costs required to maintain a SaaS infrastructure product that promises Five-Nines uptime, has a global customer base, and always-on functionality versus an SMB product where uptime isn’t a critical factor and usage is periodic throughout a week. Or the front-loaded investment of the human-in-the-loop AI business model. Both have very different cost structures from traditial SaaS. Venture investors will take those differences into account and factor in the many other factors that matter — market size, technology advantage, company leadership, and so forth.

Often what is most important with Gross Margin is the trend: Can your company demonstrate that it has sufficient operational leverage to improve margins over time? If GM% is low early on, can you articulate a path to higher margins as you scale?

Operating Expenses

Expenses that do not fall into the Cost of Revenue bucket are considered operating expenses (OpEx). Here the three major categories are:

Sales & Marketing

All of the costs incurred in selling and promoting your product. For a typical early-stage SaaS company, S&M expense accounts for around half of total OpEx. These expenses include:

- Payroll for sales team (salary+commissions+benefits+payroll taxes)
- Payroll for marketing team (salary+benefits+payroll taxes)
- Advertising and conferences (e.g., digital offline ads; events that team members attend to promote your product/company)
- Software and tools used by sales and marketing teams
- Cloud costs associated with trial acquisition, as it's generally a lead generation platform and, thus, marketing.
- Outside vendors who support S&M, including consultants and contractors

Research & Development

Product development costs including dev teams, the software and tools they use, and other direct costs for building your SaaS product. "Product development" is what we're talking about here, distinct from the product costs in COR tied to running and maintaining the software.

- Developers and product teams, including salary + benefits + payroll tax per employee
- Software and tools used by these teams in their work

General & Administrative

Overhead not accounted for in the other categories. Often a catch-all for all day-to-day expenses related to running the business like:

- Rent on office space, utilities payments, communications services and usage
- Executive salaries+benefits+taxes
- IT/Finance/Accounting/HR teams and other support staff salaries+wages+benefits+payroll taxes

There's an important best practice related to how expenses are attributed to the different categories: be scrupulous about fully allocating departmental expenses to the department — and not to a G&A catchall bucket. This means that your S&M and R&D lines should reflect the department's portion of costs like legal, rent, and recruiting. In some ways, it's like treating the two keep departments as operating entities, setting up your P&L to accurately reflect what it costs to run them. Doing so will improve your internal tracking of KPIs like, for example, marketing ROI. And of course allow you to benchmark those functions against other similar companies.

Other Costs

Expenses that don’t fall into any of the above categories are grouped into an Other Costs bucket that sits between Operating Profit (Loss) and Net Profit (Loss). These include items like Interest Expense (payments on debt and lines of credit), interest income (earned on idle funds) and income taxes. Non-cash charges, or expenses required by accounting rules that do not correspond to a cash outlay such as depreciation expense on computer equipment and software (depreciation is the accounting concept that assets decline in value over time as they are used) may also be reported in Other Costs, or may be considered an Operating Expense.

This lower section of the P&L statement isn’t generally a focus of early stage investors. The key signals of a company’s performance and potential lie higher up at the Sales, Gross Margin, and OpEx lines. Other Costs become more important as a company approaches the $100M ARR threshold where an IPO becomes possible. At that point, Net Profit comes into focus, often in the form of measuring a company’s Rule of 40 metric.

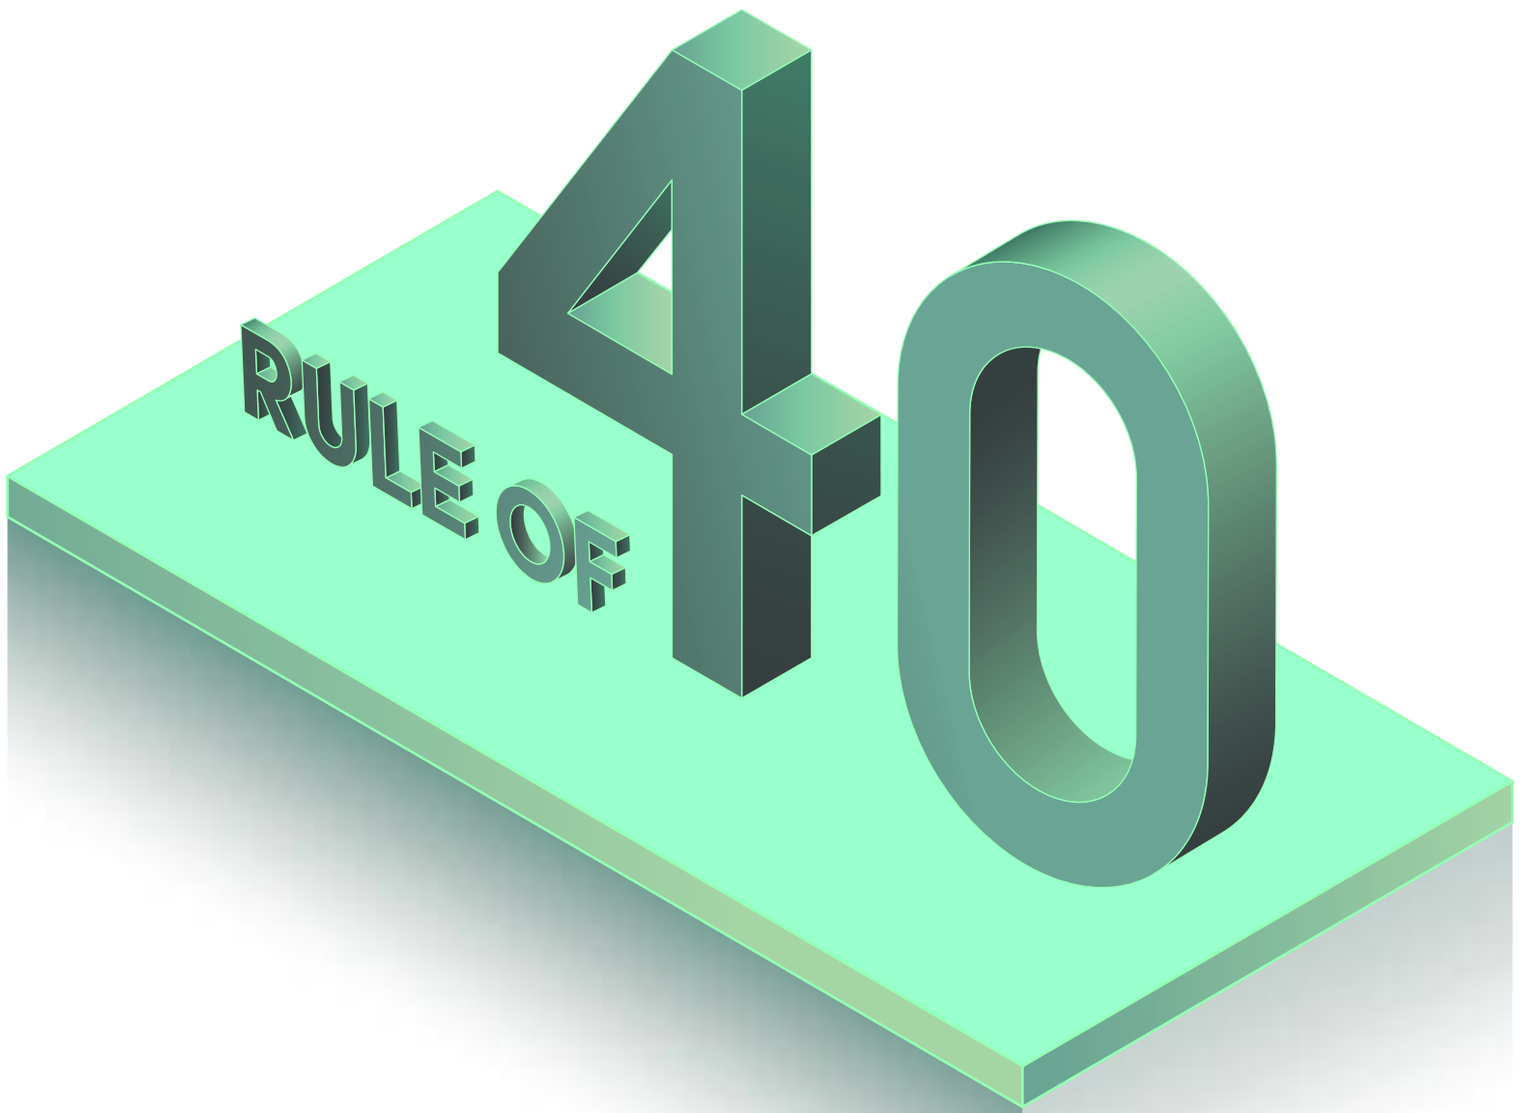

Performance Benchmarking with Common Size

So far we've looked at the P&L on a dollar basis: each line item displaying the actual dollars taken in as sales and spent as expenses. But there is another view of the P&L that provides a lot of value: Common Size. With Common Size, you convert each expense line into a percentage of revenue. It looks like this:

Profit & Loss (Common Size)		% Revenue
Revenue	$10,000,000	100.0%
Cost of Revenue (COR)	$2,000,000	20.0%
Gross Profit/(Loss)	$8,000,000	80.0%
Sales & Marketing	$3,500,000	35.0%
Research & Development	$2,500,000	25.0%
General & Administrative	$1,500,000	15.0%
Total Operating Expenses	$7,500,000	75.0%
Operating Profit/(Loss)	$500,000	5.0%
Other Income/(Expense)	$0	0.0%
Net Profit/(Loss)	$500,000	5.0%

From a metrics perspective, investors often use common size to normalize the cost structures of two or more companies, calculating OpEx line items as a percentage of total Operative Expenses. So, R&D as % of Total OpEx or S&M as a % of Total OpEx.

As we dive deeper into burn rate reduction, Operating Expenses is an area filled with opportunities to reduce or eliminate spending — or, more formally, to bring OpEx to a level better aligned with growth rates and return on capital.

Operational Efficiency Isn't a P&L Line Item... But It Matters

How streamlined is your business? One key performance measure that P&L alone can't tell you about is the efficiency of your teams and operating processes.

> *Take this example: a company spends $1M on S&M during a quarter and generates $2.5M in revenue. What happens if it doubles S&M spending to $2M — should the company expect to generate $5M in revenue? It's rarely that simple. There are a lot of operating processes between S&M spend and new sales. Operational efficiency is an x-factor in a company's money out versus money in.*

That's where SaaS metrics come into play. By thoughtfully tracking and analyzing key operating metrics over time, you can learn about the sensitivities between spending and new sales.

Calibrating SaaS Cash Burn

Leadership teams across the startup ecosystem are rightly focused on cash burn. On-demand access to venture capital can no longer be assumed in an environment where venture firms are thinking far more carefully about where to allocate capital. Startups that are spending to grow but have yet to reach cash independence are being forced to protect the money that they have and reduce unnecessary overhead. Burn must now be more deliberate, as cash is one of the non-negotiable items that is crucial for survival.

What, precisely, is “cash burn” and why is there more to it than just what a company spends? For such an important concept, cash burn is often loosely defined or confused with other terms such as operating loss or cash from operations.

What Is Cash Burn?

Cash burn refers to the change in a company's cash position from the beginning to end of a period. It can be calculated by looking at a company's cash flow statement, which outlines how a Company's cash position changes during a period of reference. Cash flows are typically broken down into three distinct buckets:

Cash from operations

This outlines how a Company's cash inflows from customer collections less general business expenses (e.g. salaries, technology spend, hosting costs) increase (or decrease) one's cash position. This number is close to a Company's net income (or loss) - with which "cash burn" is often confused - but with non-cash expenses (such as stock-based comp, depreciation, or amortization) added back.

Cash from investing

It shows how a Company's cash is affected by internal investment activities, such as the purchase or sale of stocks, bonds, or physical assets. Because most early stage SaaS startups tend to be asset-light businesses (e.g. do not purchase heavy equipment or real estate and have relatively small amounts of investible cash on the balance sheet, activities in this bucket tend to have a minimal impact on cash changes from period to period.

Cash from financing

This bucket outlines how external financing affects Company's balance, that is an equity financing from an outside investor like a venture fund that purchases company stock in exchange for cash. At more mature stages, startups may also conduct debt financings, whereby an outside investor loans money to a company that is repaid at a specified schedule along with interest payments. These types of payments would fall into this bucket.

Factors Impacting Cash Burn

Outside of external events - which we'll get into below - the most important factors that directly affect a startup's burn rate are a company's cash collections (from invoices), gross margin profile, and the operating costs used to grow the business.

In the world of early-stage SaaS, growth and burn are often closely related, as it's unusual (but not unheard of) for startups to find growth levers that translate into disproportionate levels of efficiency gains. Churn is also the enemy of growth and, if severe, can have material effects on a startup's cash burn.

Some investors encourage startups to make the distinction between "gross" and "net" cash burn to delineate how much cash a company is consuming independent of cash collections. The two measure related aspects of your cost structure and operations:

"Gross" cash burn equals the amount of cash that the company spends on operations before cash collections from revenue are added back.

"Net" cash burn is measured by a company's cash consumption during a period after accounting for any cash inflows resulting from invoices.

We often advise tracking both measures to understand how cash (and total runway) differ across various planning scenarios.

The Less Obvious Factors Affecting Cash Burn

There are a few other factors that affect a startup's cash burn that may be less obvious. At first glance, some of these areas may feel like they're on the margin, but as startups begin to scale, they can start to have a material impact on a company's burn profile. Let's jump into a few of these areas below:

CUSTOMER PREPAYMENTS

Part of the beauty of most SaaS businesses is that customers agree to pay for services on a subscription basis. Most SaaS customer agreements have monthly or annual terms where the customer has the choice to renew their contract at the end of each period.

Occasionally, a customer may elect to prepay for services, opting to pay in advance for the entire year instead of, say, four quarterly installments. Not surprisingly, this has a positive impact on a startup's cash position. In accounting terms, a prepayment is recognized as a liability (deferred/unearned revenue) on the company's balance sheet.

In cases where customers agree to prepay for multi-year contracts, the positive effect on a startup's cash position will be even more pronounced — which is why sales compensation is often structured to incentivize prepayment terms.

INTEREST AND OTHER PAYMENTS ASSOCIATED WITH EXTERNAL FINANCING

We discussed the two major types of external financings above, which are equity and debt offerings. Equity financings involve the exchange of cash for company stock, which dilutes founders and other earlier investors. A debt financing is effectively a loan to the business that the startup agrees to repay at a future point in time. For the startup, debt's advantage is that it is non-dilutive. From an investor's perspective, one advantage is that it will be senior to any equity in the event of a sale.

Debt offerings can become complex, but for simplicity's sake, debt usually consists of two forms of payments: principal, the amount of the loan, and interest, the cost of the loan. Interest payments are often amortized over the life of the loan, but in some cases (take balloon payments, for example), can happen all at once. Not surprisingly, debt repayment can have a material effect on a startup's cash position and should be carefully planned for.

DAYS SALES OUTSTANDING (DSO)

This is the average number of days a startup takes to collect payment from a customer. If a company, for example, requires two months from the date of an invoice to collect payment, their DSO would be 60 days. Moving this number up or down can affect a company's cash position, as receivables convert into cash earlier or later depending on which direction DSO fluctuates. As an aside, this is often a point of contention during vendor procurement as startups tend to have opposite motivations from their customers with respect to cash collections.

CLAWBACKS AND OTHER ADJUSTMENTS

Occasionally service agreements don't work out and startups are forced to adjust a contract's terms of or terminate it all together. In cases where a customer has prepaid for their service, this situation could result in a clawback, whereby the startup is required to repay some amount of cash back to the customer.

Tying It All Together

Cash burn is one of the Four Vital Signs of SaaS and ultimately the one asset without which a startup cannot survive. Run out of money? The party's over.

We have entered a period in which cash burn has become a focal point, especially for early-stage startups where growth has traditionally been prioritized over spend and cash independence is usually not achieved until later stages. In this environment, it's critical to understand and measure the many factors that go into burn rate. Doing so is a key to ensuring access to your next round of venture investment.

The next chapter in our Burn Rate section will present Scale's research into just how much cash SaaS startups use to scale their businesses and benchmarks on operating expenses at various revenue levels.

A Quick Primer on the Rule of 40

There is an expansive universe of SaaS metrics out there. Which is great for tracking performance internally, as it allows a company to customize metrics to reflect its operating realities. But what if you want to compare performance between companies? That's when the lack of standard definitions becomes a problem. Just look at how many ways customer retention is calculated.

To solve the apples-to-oranges problem, Scale developed a framework called the Four Vital Signs of SaaS, which examines four critical operating metrics - revenue growth, sales efficiency, customer retention, and burn rate - to help entrepreneurs better understand the health of their business. As vital signs quickly help medical professionals diagnose the health of a patient, these four metrics can do the same for quickly understanding the health of your SaaS business and provide a basis for comparison and benchmarking.

We wanted to apply the spirit of the Vital Signs standardization to another SaaS metric that suffers from a lack of common definition: the Rule of 40 ("RO40"). The RO40 measures the balance of a SaaS company's growth and profitability. While this sounds simple in theory, what makes this more complicated in practice is the lack of a standard measure for profitability. In this post, we'll take a closer look at some of these nuances.

A quick disclaimer before we get into it: as investors in early-stage enterprise software companies, Scale doesn't use the Rule of 40 to evaluate prospective investment opportunities. We understand that most early-in revenue, venture-backed startups tend to favor growth vs. profitability. In this context, looking too closely at the Rule of 40 may drive entrepreneurs toward misleading conclusions. The RO40 is, however, much more meaningful for more mature SaaS businesses, especially as they begin to enter the public realm.

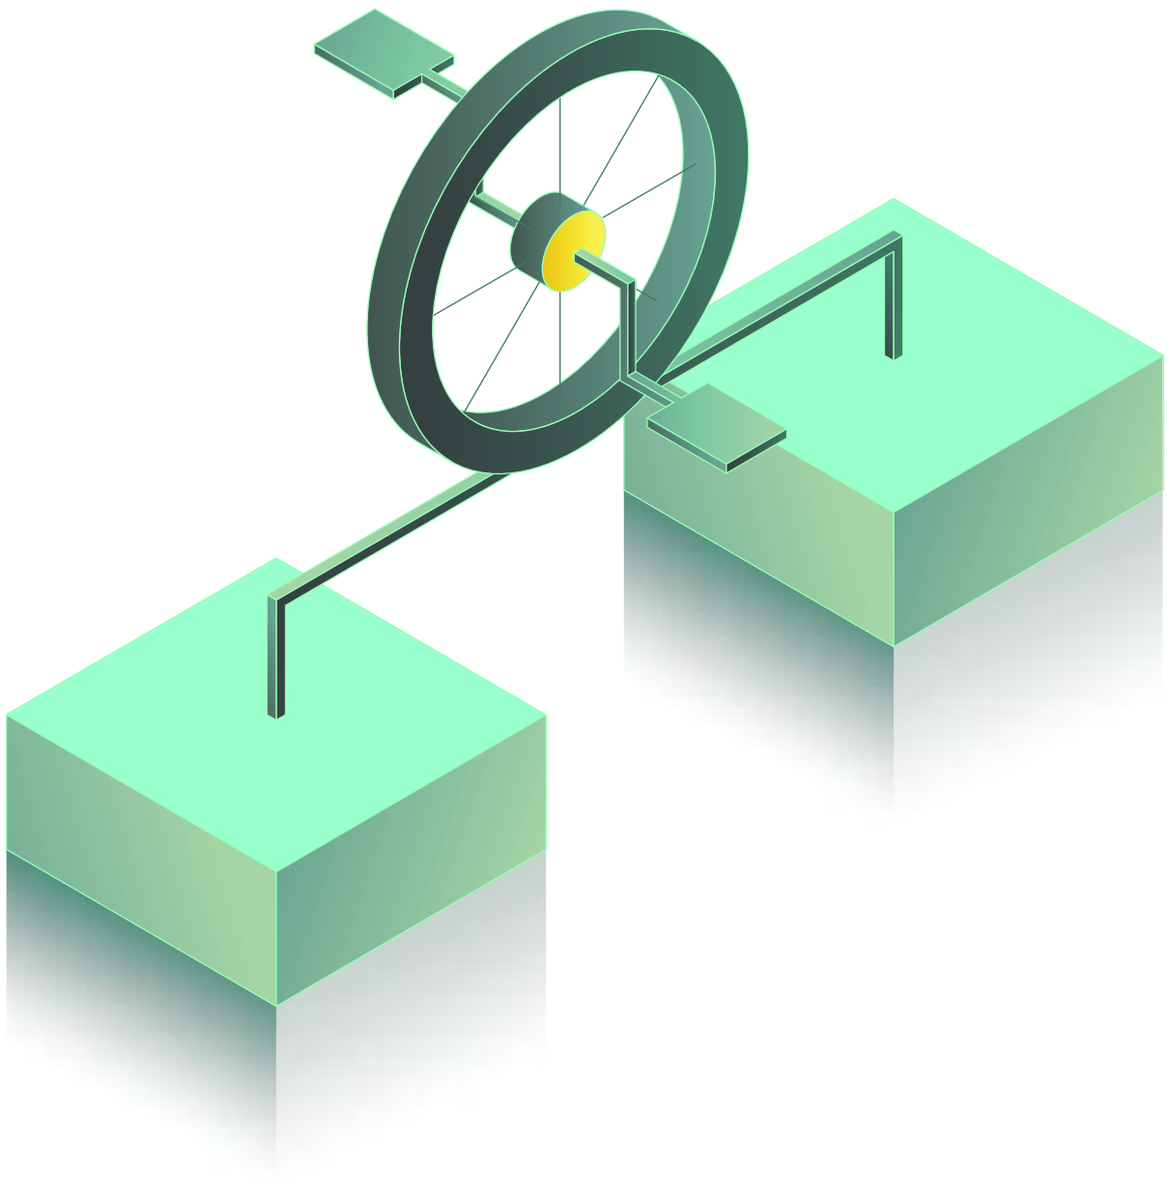

What Is “The Rule of 40”?

The Rule of 40 states that, at scale, a company’s revenue growth rate plus profitability margin should be equal to or greater than 40%. SaaS management teams are often driving towards either rapid growth or increased profitability, and the Rule of 40 has become a construct for framing the balance of these two phenomena. Given that increased investment (whether from external or internal sources) is usually required to drive growth, rapid expansion and strong profitability are usually at odds with each other, and finding the right mix between the two can be tricky.

It is worth noting that the Rule of 40 will not help answer whether an early-stage company is growing fast enough or is profitable enough. There are other metrics and benchmarks to help us understand each of these questions. Rather, the RO40 is a measure of the balance between growth and profitability, and by extension, the sustainability of the business more broadly.

The Rule of 40’s Key Drivers: Growth and Profitability

The Rule of 40 equation simply sums two numbers:

1. Growth rate %

2. Profitability margin %

Let’s take a closer look at both.

GROWTH RATE %

While there are many different ways to measure growth rate, the easiest and most commonly used from company to company is year-over-year growth % based on GAAP revenue. GAAP revenue doesn't lie (or it certainly shouldn't) and provides the most consistent basis for comparison. While this is especially true for more mature SaaS businesses, early stage SaaS companies may be more focused on MRR or ARR growth. To be clear, these top line growth measures are perfectly fine to use when calculating your Rule of 40, with the one caveat that this may create an inconsistent basis for comparison between companies. In a perfect world, the way in which MRR and ARR are recognized would be as consistent as the methods for recognizing GAAP. Unfortunately, and especially in the early days, this is often not the case for many SaaS businesses.

PROFITABILITY MARGIN %

This part of the equation is a bit more tricky because there is no generally agreed upon measure of profitability. Unlevered Free Cash Flow, Cash from Operations, Net Change In Cash, Operating Income, and EBITDA are all different measures of profitability (the margin of which is calculated as a % of GAAP revenue) and legitimate candidates for use in the RO40 calculation (as the legendary Brad Feld wrote about several years ago). And what about stock-based compensation? Should these costs be included or excluded when calculating the RO40?

To be clear, there is no "correct" answer as to which profitability measure should be used when calculating and tracking the RO40. All of the above are valid. What's important here is understanding that each measure of profitability will yield different results — and sometimes these differences are quite meaningful. For simplicity and comparability across companies, we use EBITDA excluding stock-based compensation (SBC) costs and will get into why a bit later in this post.

Let's Jump Into an Example

The table below calculates the Rule of 40 for DocuSign's (NASDAQ: DOCU) recently reported quarter, CQ4 2019:

GAAP Revenue Growth (%)		
Q4 2018	$	178.4
Q4 2019	$	249.5
YoY GAAP Growth Rate (%)		39.9%
Profitability Margin (%)		
Cash From Operations		13.7%
Unlevered Free Cash Flow		8.6%
EBITDA ex. SBC		3.5%
Operating Income ex. SBC		(0.7%)
EBITDA		(12.6%)
Operating Income		(17.6%)
Rule of 40 (Growth Rate + Profitability Margin %)		
Growth + Cash From Operations		53.6%
Growth + Unlevered Free Cash Flow		48.4%
Growth + EBITDA ex. SBC		43.4%
Growth + Operating Income ex. SBC		39.1%
Growth + EBITDA		27.3%
Growth + Operating Income		22.2%

During the quarter, the Company posted YoY GAAP revenue growth of 39.9%. The Company's profitability margin, however, varies considerably depending on which way you look at it. Of the six measures shown above, half point to net profitability (where the profitability margin % was > 0%), while the other half suggest that the Company operated at a loss during the period (where the Company's profitability margin % was < 0%). Combining DocuSign's YoY GAAP revenue growth % with the first three profitability margins would indicate that the Company would clear the 40% threshold, while the other three methods would point to the Company's missing the mark — and in some cases by a wide margin.

To reiterate, there is no “right” or “wrong” answer (or measure of profitability) to use here. All of the above are valid measures, but each tells a different story. Some measures of profitability are more conservative than others and can make a meaningful difference in determining whether or not a company’s performance would land over or under the 40% benchmark.

As we see above with DocuSign, the choice here can shift the RO40 profitability contribution from +13.7% (Cash from Operations) to -17.6% (Operating Income).

The point here is to be clear about what you’re saying.

Using EBITDA excluding SBC costs as a profitability measure is, not unlike using GAAP when calculating growth rate, another way of creating a consistent basis for comparison from company to company. First, measuring profitability by using EBITDA smooths out differences in interest payments or depreciation of intangible assets when comparing companies. While these payments are generally trivial for early stage SaaS businesses that lack the cash flow to service outstanding debt, for more mature companies, these kinds of payments can have a meaningful impact on the bottom line.

Second, and perhaps more importantly, issuing stock to employees is a common practice at most software companies. Why? This is usually the most readily available currency that companies in this segment have at their disposal (especially those with exceedingly high cash burn). As such, at certain times of the year (e.g. bonus season), SBC expenses can be meaningful (and also inconsistent), which is why netting out SBC costs from EBITDA helps create a more consistent basis for comparison from company to company over time.

How Has the Rule of 40 Played Out In The Market Over Time?

Scale maintains a database of key metrics for – at the time of this writing – 68 publicly traded SaaS businesses. One of these metrics is the Rule of 40, calculated by using the method above: adding YoY GAAP revenue growth % + EBITDA margin % ex. SBC costs. The chart below shows the median Rule of 40 across all companies in our database dating back through 2011.

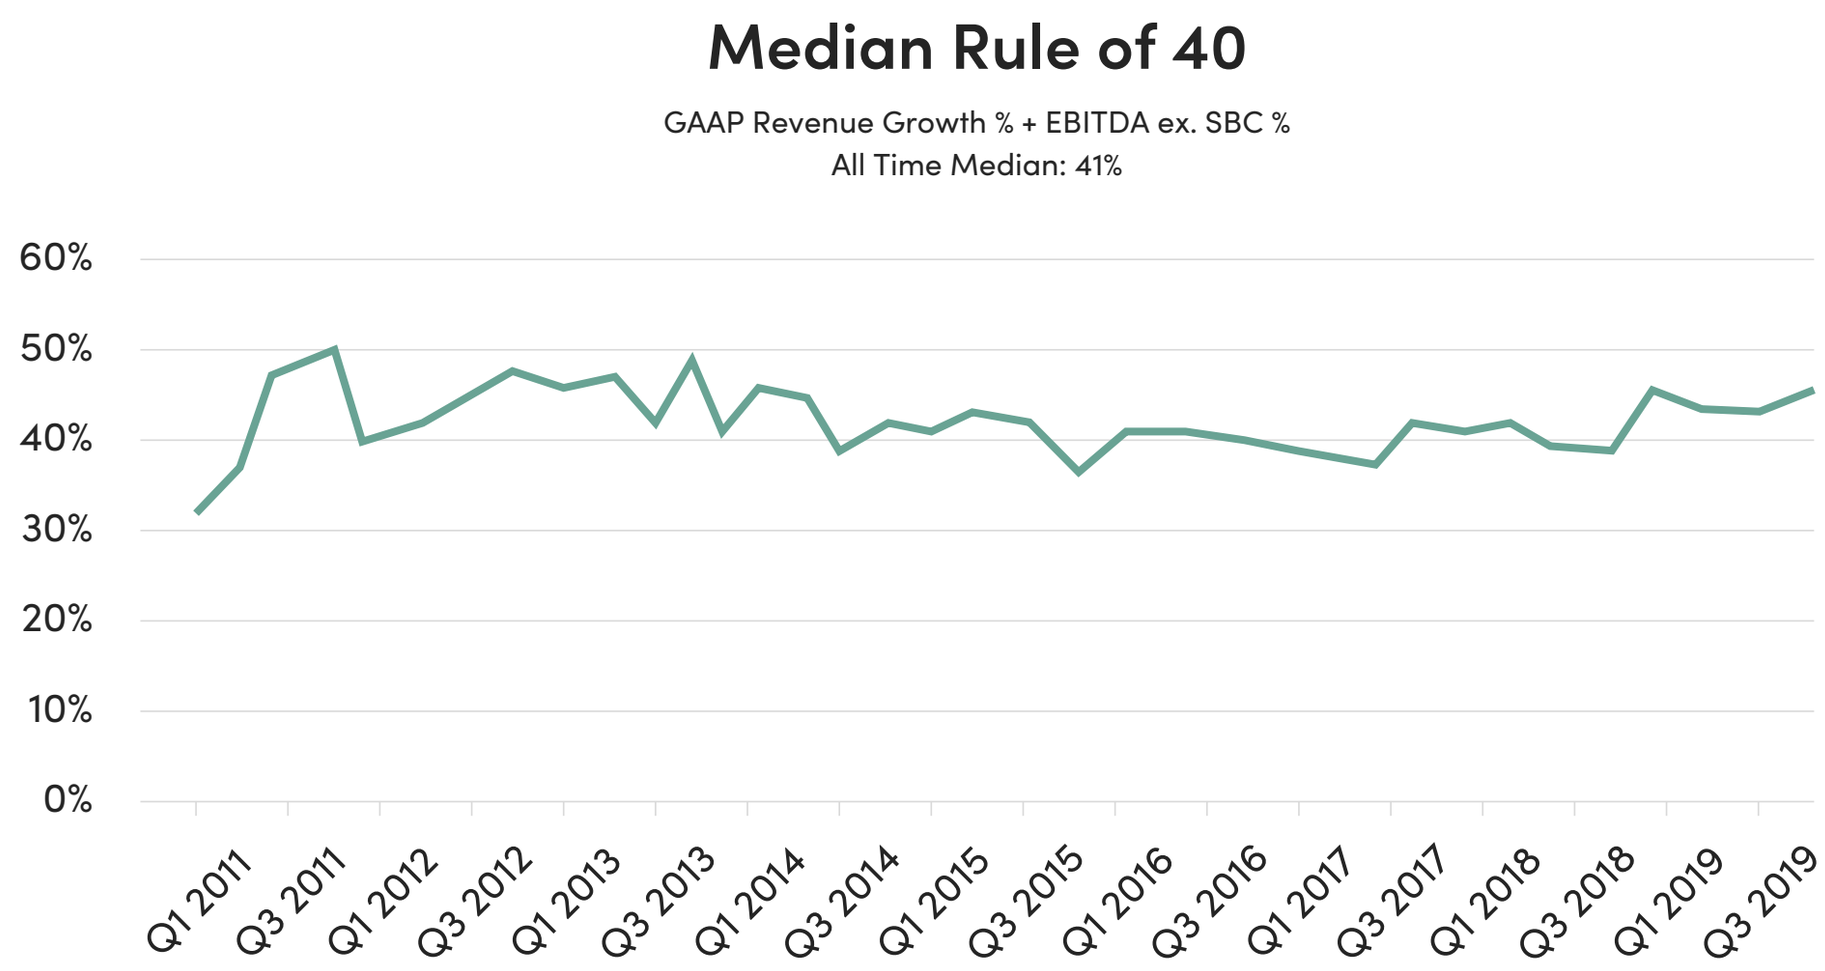

What becomes clear is that this measure has stayed rather consistent for nearly a decade, and not surprisingly, has converged around 40% (+/- a few percentage points) over this time horizon. The median across the last ten years has been 41%, which suggests that the Rule of 40 is in fact a good proxy for measuring the balance of growth and profitability for more mature SaaS businesses.

It is worth noting that the businesses in this comp set have all reached meaningful scale. The median annual GAAP revenue of this basket of companies during the stated time period is $185M, which points to considerable company maturity. The chart above would imply that the Rule of 40 is very fitting for this set of comps, though going back to our earlier disclaimer, one may argue that at different stages (especially in early development) and for different objectives, over-indexing towards growth or profitability may be the right thing to do for the business. In these cases, anchoring too closely to the Rule of 40 may draw management towards misleading conclusions. Nothing is valid without context.

Tying It All Together

The Rule of 40 has become a widely used unit of measure for understanding the relationship between a SaaS company's growth and profitability.
A ten-year look back of the data shows that the RO40 has remained quite consistent among public SaaS companies, suggesting that the measure is a useful barometer of the balance between a business's expansion and profitability, and by extension, the general sustainability of company performance over longer intervals of time. We hope this post has helped bring some clarity to understanding some of the specifics around this measurement, which often go overlooked.

Ready to Scale?

Scale Studio benchmarks let you compare your growth and operations to hundreds of successful SaaS startups.

SCALE STUDIO

SCAN ME

Get your report at
www.scalestudio.vc

SCALE

scalevp.com