

Desk Notes On SaaS Metrics

evp_

If you're a SaaS founder who's looking to raise a Series A or are post Series A, you probably have enough data points to track your SaaS metrics and use them to refine your decision making. One of the things we help our companies with is cohort analysis. Here are our desk notes to help you think through cohort analysis, why it's useful and some things to look out for.

Not all cohorts are made equal

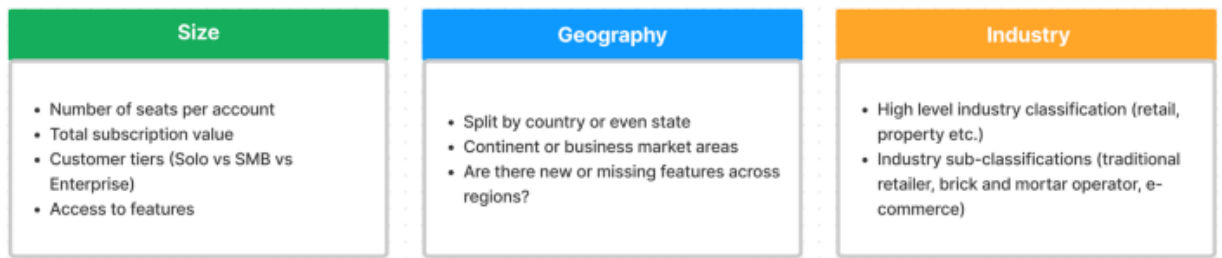
At EVP we love analysing cohorts. Cohort analysis is essentially a way to tell how a select group of users are 'behaving' over a period of time. This behaviour becomes a proxy for us to understand the performance of a select group of users over time, and how that has translated to revenue and how 'sticky' a product is.

Here are a few of the things we consider when analysing cohorts...

- **Time parameters.** Cohort analysis almost always involves a time-bounded approach to grouping users. The most straightforward example is taking all users that joined on a particular month and grouping them together as one cohort. That cohort ages consistently so that you know at 6 months, 8 months and 12 months later the factors impacting that cohort are essentially the same. You can also group into

quarterly or annual cohorts for simplicity and larger sample sizes.

- **Performance metric.** SaaS cohort analysis requires a metric to track, the most common being subscription revenue (Revenue Retention) or number of accounts (Logo Retention). Subscription revenue can be analysed on a gross or net basis which includes revenue increases from upgrades, cross-sells and increasing usage. Ultimately, the point of tracking these two is that they provide the strongest signal of product stickiness and fit. Other performance metrics can include stats like user engagement (logins, shares, usage) or unit economics (payback, LTV)... anything that can signal stickiness, fit or value created.
- **Segment investigation.** Your SaaS business might include multiple sub-products or feature sets bundled into one subscription with different use-cases depending on the target customer. There might also be multiple tiers of features or country specific content so it's often interesting to isolate the behaviour of user groups by further segmentation. You want to consider how user groups might engage with the product differently and how certain features may or may not be available or applicable to these groups.



Examples of segments we like to look at

Presenting cohorts

[These notes assume an understanding of how to create cohort analysis]

Cohort analysis is typically presented in one of three ways shown below. These visualisations are useful in making sense of the data and eliciting patterns in the numbers. They're arguably the all-star visual communication tools for understanding business performance.

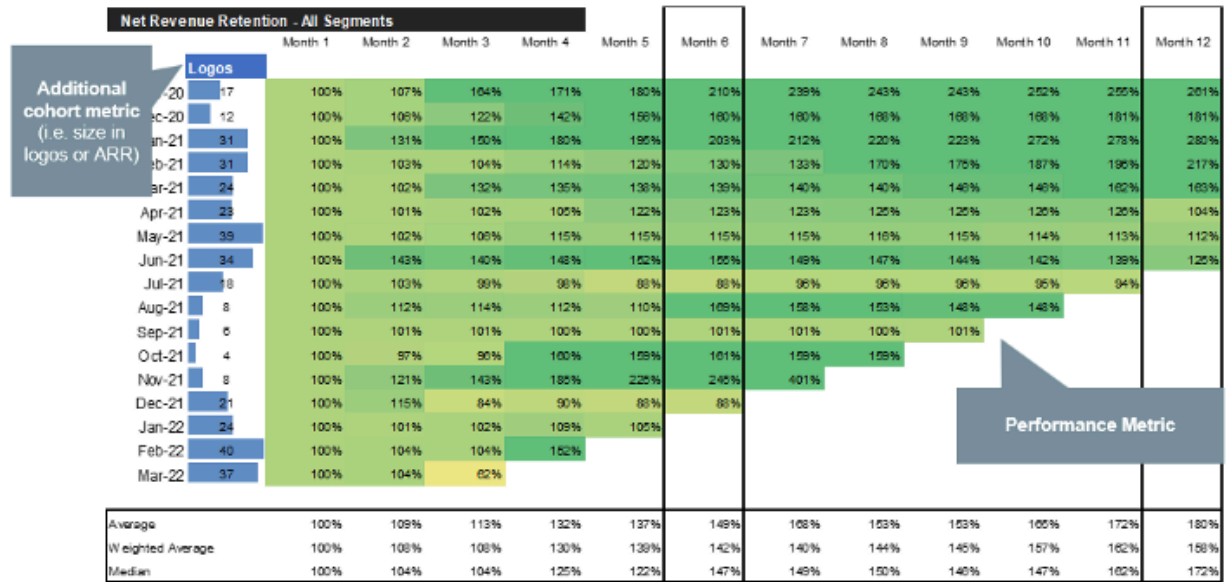
The trusted data table

The most common way SaaS investors present cohorts for analysis is in a data table. Cohort data tables generally group on the Y axis by the month each cohort joined. Lifetime is displayed on the top X-axis, representing their first, second and so-on month with the product.



Net Revenue Retention Table—EVP analysis

Performance metrics are shown through the body of the table and colour scales indicate good or bad. The size of each cohort is often referenced on the side (blue bar charts next to the date). Reading down the column helps assess improvement across cohorts over time and across a row shows retention within a cohort.



Net Revenue Retention Table—EVP analysis


The layered cake

The layered cake is a stacked area chart normally used to show ARR over time. Layered cake charts are incredible for showing composition of revenue by cohort and the level of revenue expansion or contraction across the board. Ideally the layered cake shows each cohort is stacking and continually expanding such that overall ARR growth is exponential.

Layered cake also highlights the proportion each cohort is contributing to overall revenue at each point in time and the overall age of your revenue base.

—Mark Velik


Annual Recurring Revenue (ARR) by Annual Cohort through January 31, 2019



ARR layered cake chart—from Slack Prospectus

The party streamer 🎉

The party streamer is a type of line chart used to display the time it takes for each cohort to hit a certain milestone. It's a popular way to monitor performance in direct to consumer business where time to fully payback acquisition cost is typically tracked. The ideal party streamer chart shows increasingly steep lines as later cohorts reach profitability earlier (indicative of improving acquisition funnel or increasing basket sizes).



Cohort payback line chart—from Eucalyptus Pitch Deck

Dissecting cohort performance—tell me what the story is!

If you're running a SaaS business you've probably come across a multitude of guides telling you what 'good' metrics look like.

		Consumer SaaS	SMB / Mid-Market SaaS	Enterprise SaaS
Logo Retention	Good	40%	60%	70%
	Great	70%+	80%+	90%+
Net Revenue	Good	55%	90%	110%
	Great	80%+	110%	130%+

Example of retention benchmarks from Lenny's Newsletter
<https://www.lennysnewsletter.com/p/what-is-good-retention-issue-29>

While these benchmarks are helpful for making a quick assessment and generally accurate for pure subscription software, there are some nuances to look out for. Every cohort tells a story and stories are more interesting if we understand the context; which in the case of SaaS, is about the product, the market and the business model. **Not all cohorts are made equal and here are some things to look out for.**

The pricing model and the Month-0 problem

Pricing model changes. Be conscious of how pricing changes can impact the shape of revenue cohorts. One of our companies experimented with a three-month discounted price ramp to encourage new customers. As users transitioned from their discounted plans to full-paying plans we see overall revenue jump

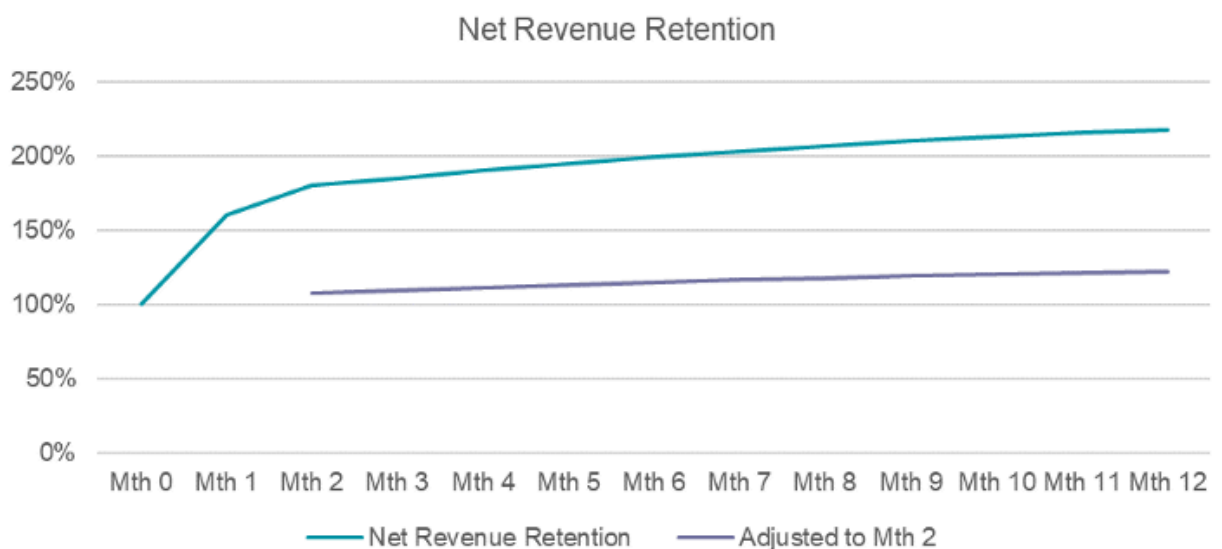
>200%, an exciting increase when seen in isolation. Zooming out we start to see these numbers trend back towards the pre Mar-20 cohorts. It's often worthwhile normalising the effect of this price ramp to track the drop off from month 3 otherwise the numbers may be artificially inflated.

Cohort	Initial MRR	Month 1	Month 2	Month 3	Month 4	Month 5	Month 6	Month 7	Month 8	Month 9	Month 10	Month 11	Month 12	Month 13	Month 14	Month 15	Month 16	Month 17	Month 18
Sep-19	9,497	100%	99%	115%	116%	114%	114%	114%	114%	104%	107%	107%	107%	101%	97%	97%	97%	92%	91%
Oct-19	6,241	100%	96%	96%	101%	101%	103%	97%	97%	97%	96%	96%	96%	96%	96%	96%	96%	96%	96%
Nov-19	4,320	100%	111%	109%	109%	109%	92%	92%	69%	69%	69%	69%	69%	71%	71%	71%	71%	71%	71%
Dec-19	2,094	100%	100%	93%	83%	83%	83%	83%	83%	83%	83%	83%	83%	89%	55%	55%	55%	55%	54%
Jan-20	3,126	100%	104%	100%	100%	108%	112%	107%	107%	107%	107%	107%	107%	95%	95%	105%	105%	78%	58%
Feb-20	5,334	100%	104%	105%	103%	93%	93%	91%	89%	80%	79%	82%	74%	69%	63%	63%	62%	60%	61%
Mar-20	3,706	100%	102%	146%	152%	156%	156%	144%	144%	143%	143%	131%	123%	122%	107%	89%	81%	83%	84%
Apr-20	7,005	100%	94%	177%	180%	169%	162%	163%	163%	167%	153%	150%	144%	137%	113%	116%	130%	115%	110%
May-20	4,502	100%	104%	217%	210%	207%	207%	204%	204%	203%	199%	199%	199%	186%	173%	198%	192%	173%	
Jun-20	6,151	100%	100%	207%	201%	199%	192%	173%	162%	160%	158%	158%	162%	141%	146%	148%	153%		
Jul-20	5,101	100%	100%	173%	198%	190%	186%	185%	185%	190%	190%	188%	168%	189%	184%	181%			
Aug-20	6,674	100%	98%	184%	188%	186%	186%	182%	171%	171%	114%	115%	122%	138%	148%				
Sep-20	4,080	100%	93%	157%	175%	179%	179%	186%	187%	173%	156%	167%	158%	173%					
Oct-20	5,517	100%	102%	171%	163%	159%	153%	142%	152%	154%	150%	146%	145%						
Nov-20	3,507	100%	100%	212%	240%	241%	242%	239%	235%	226%	240%	224%							
Dec-20	3,006	100%	118%	354%	357%	345%	349%	349%	339%	322%	307%								
Jan-21	3,653	100%	318%	326%	302%	268%	270%	261%	240%	239%									
Feb-21	4,344	100%	143%	128%	130%	131%	112%	135%	118%										
Mar-21	2,932	100%	108%	133%	438%	454%	459%	432%											
Apr-21	2,373	100%	142%	316%	331%	294%	284%												
May-21	4,187	100%	273%	350%	283%	286%													
Jun-21	3,529	100%	95%	95%	92%														
Jul-21	3,890	100%	139%	137%															
Aug-21	3,239	100%	117%																
Sep-21	2,864	100%																	
Average		100%	110%	131%	139%	137%	135%	131%	124%	121%	116%	115%	111%	111%	113%	111%	109%	106%	103%
Median		100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	98%	98%	95%	93%	89%	85%	83%	83%

Net Revenue Retention data table - EVP portfolio company

Time to value and the Month-0 problem. Outside of pricing impacts, it's often useful to start cohorts from a future data point (Month-2 or beyond) where revenue and usage is more stable. Particularly in models where usage tends to ramp up as customers experiment and adapt. For example, Hnry is a

business finance and tax app for the self-employed that charges based on a percentage of income through the app. Starting from Month-0 will show an increasing attach rate as customers start to migrate their invoicing and payments collection and ramp-up to steadier levels of usage. Looking at cohorts without adjusting to the steady-state would conflate the effect of attach rate with long-term retention.




Adjusting the start to Mth 2 reveals a steadier state of retention

Subscription retention alone can be misleading...

One of our best performing investments has below average logo retention and <100% net retention on subscription revenue. One of the reasons we love SaaS is the ability to wedge into the market and extend value through other offerings like payments,

marketplaces and services (David Yuan has a great article about this

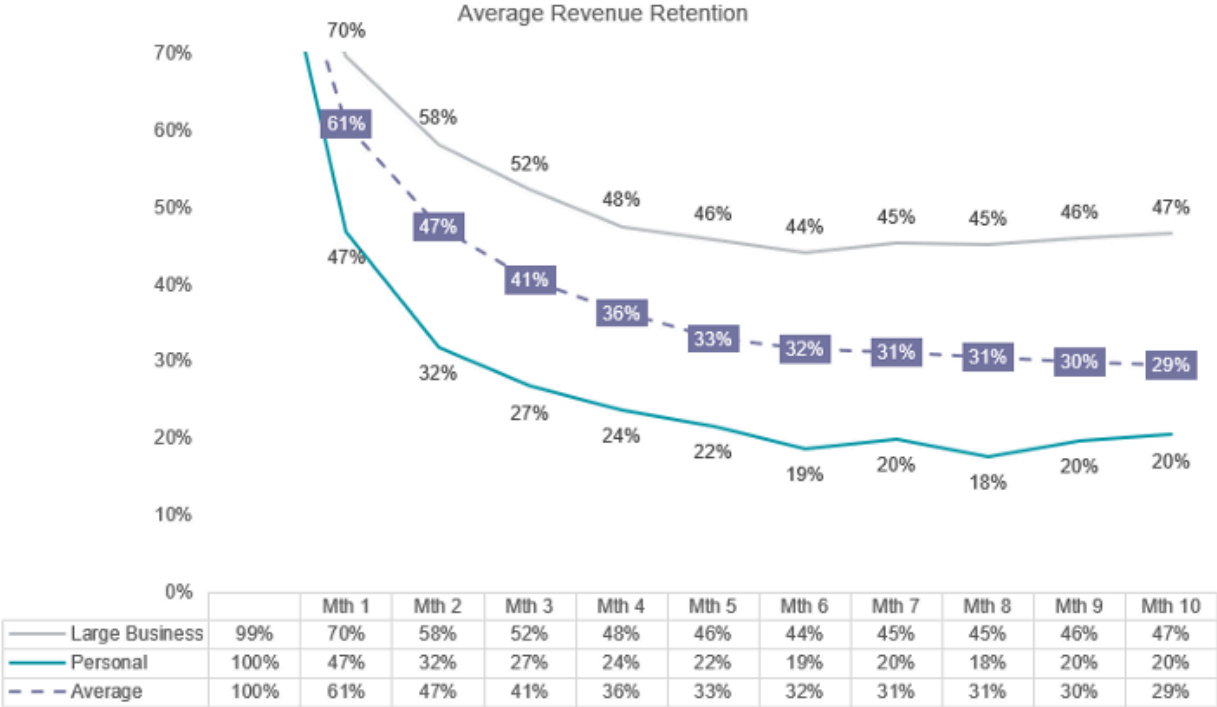
<https://www.linkedin.com/pulse/frontier-extend-through-value-chain-david-yuan/>). Here the users that stuck around were using the software to facilitate hundreds of payments a month and paying a small fee on every invoice they issued through the software. A large percentage of users churn but the value is sitting in the sticky customers who are running their entire businesses through the system.



Logo and Revenue Retention (split by Software only and Software+Payments)

Product-led and consumer facing

Some products are built for SMBs but end up attracting a stream of direct customers that are hobbyists, one-person shops or students. This is typically more pronounced in product-led models where the barrier to sign-up is very low. The price sensitivity is higher and time to churn can be much quicker as the product is used for one-off projects instead of ongoing business purposes.



Average Net Revenue Retention split by user type shows a different story when isolating personal/large accounts.

As such, logo retention can be very deceiving as low value customers get churned off quickly. Net revenue retention can also look abnormally low unless you split out the target users from the general users. We’ve invested in companies with these kinds of cohorts and look to split out the high value segments and make sure that acquisition funnel is efficient enough to acquire those high value customers.


There might be a steep decline [with retention lines] at first but the fact it eventually flattens out shows there are a group of customers for which there is great fit with the product.

—Mark Velik

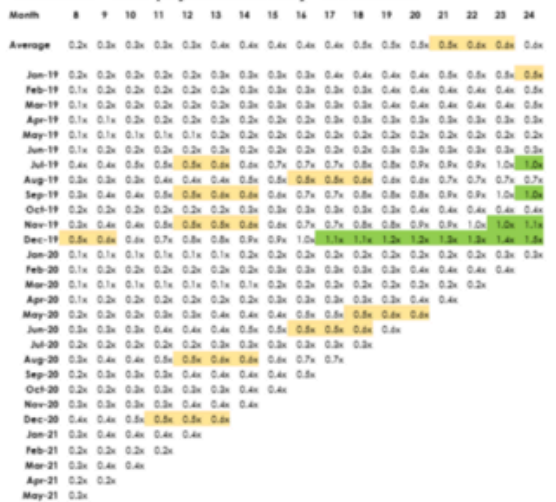
Variations across markets

One important reason to slice your cohorts by segments is that aggregate cohorts can mask problems. This exact phenomenon was surfaced during a follow-on raising by one of our companies. CAC payback (time it takes to recover the cost of acquiring a customer) across the company on aggregate was sitting at around ~16 months but looking cohorts split geographically showed that Australia took >24 months to recover.

New Zealand - CAC payback at 14 months



Australia - CAC payback currently >24 months



CAC payback split between New Zealand and Australia

This perspective helps inform how acquisition has performed. The more mature New Zealand business enjoyed the benefits of less competition and lower acquisition costs driven by wider brand recognition and better organic referrals. We've also seen this effect a couple of times across the portfolio as companies start to seed new markets and CAC improves over time.

I think payback period is an underused metric. Lifetime value is so hard to model or think about, but payback period is kind of factual. I think the point of the cohort analysis is to show customer benefit from the product but also to dictate how hard a company might be willing to go on acquisition.

— Justin Lipman

Pay to play—keep in mind cost to serve

Not all cohorts are made equal and neither is revenue. Investors like to focus heavily on acquisition channels and like to think of revenue through a paid vs organic lens. Unfortunately, we also tend to get overexcited when talking about expansion revenue and throw the frameworks out the window, forgetting to acknowledge that upselling can require its own independent sales team and significant resources to generate each incremental dollar. A dollar of organic revenue expansion (through clever

feature gating, product-led upgrades and naturally growing usage) is more valuable than a dollar of acquired upsell.

“I think it’s often missed that there might be large selling efforts to upsell or extend a customer contract, not really captured. So I think a distinction needs to be at least acknowledged by the analyser as to whether a customer is self-upgrading or being pushed. Obviously different value to the retention rates here.” —**Justin Lipman**

COVID impacted cohorts

It might seem obvious, but it’s really important to keep an eye out for COVID impact on cohorts. Fortunately, software has stayed quite resilient through the pandemic and there are plenty of products that benefitted from the accelerated adoption of technology (workplace SaaS, communications and e-commerce to name a few).

There are certainly software companies that felt the flow-on impact as their customers faced the crippling effect of lockdowns (travel, food and events). The cohort belows showed us that the cohorts that signed up prior to COVID (all are one year contracts) churned through their renewals as they came up during periods of high uncertainty and mass restrictions. The wave of new logos also decreases as we headed into and throughout 2020 / 2021.

Net Revenue Retention		Month 1	Month 2	Month 3	Month 4	Month 5	Month 6	Month 7	Month 8	Month 9	Month 10	Month 11	Month 12	Month 13	Month 14	Month 15	Month 16	Month 17
	Logos																	
Q1 2018	17	100%	107%	184%	171%	180%	210%	239%	243%	243%	252%	255%	261%	191%	258%	214%	288%	282%
Q2 2018	12	100%	106%	122%	142%	156%	160%	160%	168%	168%	168%	181%	181%	134%	220%	240%	248%	240%
Q3 2018	31	100%	131%	150%	180%	195%	203%	212%	220%	223%	272%	278%	280%	180%	155%	145%	117%	153%
Q4 2018	31	100%	103%	104%	114%	120%	130%	133%	170%	175%	187%	196%	217%	123%	121%	124%	119%	117%
Q1 2019	COVID 24	100%	102%	132%	135%	138%	139%	140%	140%	148%	148%	162%	163%	88%	82%	52%	48%	45%
Q2 2019	COVID 23	100%	101%	102%	105%	122%	123%	123%	125%	125%	126%	126%	104%	41%	40%	39%	38%	14%
Q3 2019	COVID 39	100%	102%	106%	115%	115%	115%	115%	116%	115%	114%	113%	112%	73%	73%	65%	58%	57%
Q4 2019	COVID 34	100%	143%	140%	148%	152%	155%	149%	147%	144%	142%	139%	125%	66%	84%	84%	83%	59%
Q1 2020	COVID 18	100%	103%	99%	98%	88%	88%	96%	96%	96%	95%	94%	93%	55%	55%	53%	53%	53%
Q2 2020	COVID 8	100%	112%	114%	112%	110%	169%	158%	153%	148%	148%	147%	140%	100%	124%	116%	108%	108%
Q3 2020	6	100%	101%	101%	100%	100%	101%	101%	100%	101%	102%	103%	103%	99%	126%	126%	126%	127%
Q4 2020	4	100%	97%	96%	160%	159%	161%	159%	159%	163%	165%	166%	102%	67%	0%	0%	0%	0%
Q1 2021	8	100%	121%	143%	185%	225%	245%	401%	405%	449%	510%	556%	560%	573%	542%	471%		
Q2 2021	21	100%	115%	84%	90%	88%	88%	80%	79%	78%	81%	96%	52%					
Q3 2021	24	100%	101%	102%	106%	105%	106%	115%	119%	111%								
Q4 2021	40	100%	104%	104%	152%	179%	176%											
Q1 2022	37	100%	104%	82%														
Average		100%	108%	113%	132%	139%	148%	159%	163%	166%	176%	187%	180%	137%	145%	133%	107%	105%
Weighted Average		100%	108%	108%	130%	139%	142%	140%	144%	145%	157%	162%	158%	107%	110%	101%	81%	76%
Median		100%	104%	104%	125%	130%	147%	140%	147%	148%	147%	155%	133%	99%	121%	116%	95%	83%

COVID shown in Net Revenue Retention

Seasonality

Seasonality is when a time series dataset exhibits regular and predictable changes that recur over a time period. Any predictable pattern that recurs or oscillates over a one-year period is regarded as seasonal. Cohort data tables are surprisingly great for revealing seasonality trends. The patterns can reveal how a product is being used and can alleviate fears that a product is broken. The table below shows a small-business product where the revenue model is a percentage of payments collected. The clear red diagonal that tracks up from May-18 (Mth 9) to Dec-18 (Mth 2) helps reveal a story about usage behaviour, revenue dips through the December holiday period and bounces back as people get back to work.

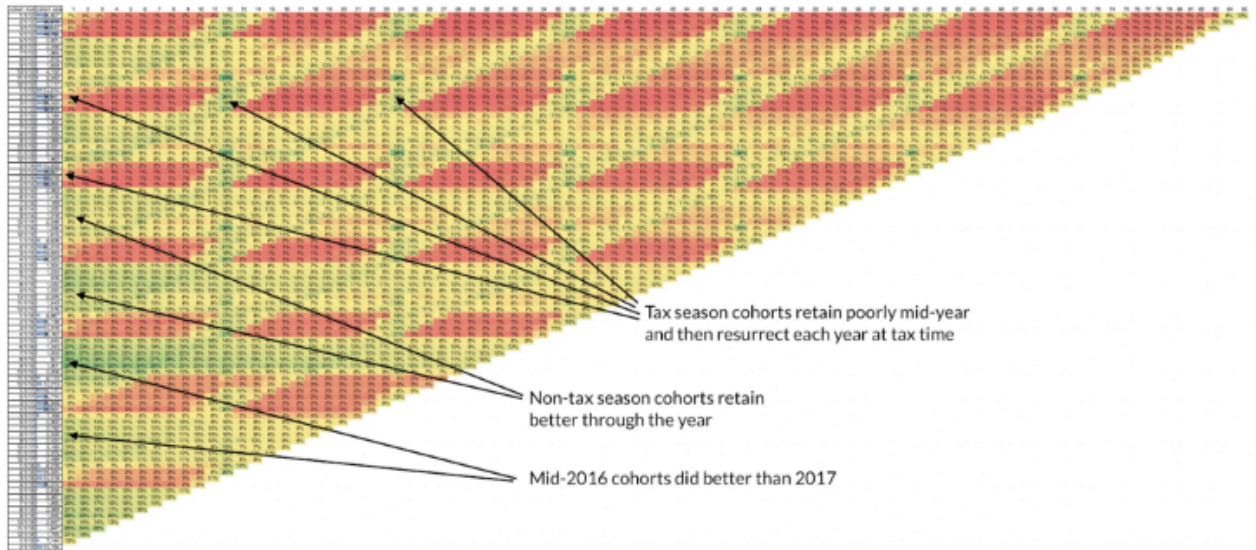
	Mth 1	Mth 2	Mth 3	Mth 4	Mth 5	Mth 6	Mth 7	Mth 8	Mth 9	Mth 10	Mth 11	Mth 12	Mth 13
May-18	100%	125%	137%	145%	86%	142%	167%	131%	66%	99%	99%	73%	83%
Jun-18	100%	180%	139%	124%	133%	126%	147%	64%	119%	120%	142%	148%	
Jul-18	100%	123%	104%	118%	102%	85%	53%	67%	44%	52%	43%		
Aug-18	100%	127%	115%	124%	94%	62%	82%	109%	109%	100%			
Sep-18	100%	122%	95%	92%	71%	78%	70%	89%	84%				
Oct-18	100%	131%	124%	69%	102%	99%	148%	146%					
Nov-18	100%	100%	62%	66%	83%	87%	75%						
Dec-18	100%	71%	82%	101%	98%	81%							
Jan-19	100%	183%	168%	193%	169%								
Feb-19	100%	125%	131%	128%									
Mar-19	100%	148%	139%										
Apr-19	100%	119%											
May-19	100%												
Average	100%	130%	118%	116%	104%	95%	106%	101%	84%	93%	94%	111%	83%
Median	100%	125%	124%	121%	98%	86%	82%	99%	84%	99%	99%	111%	83%

Seasonality has a significant impact. Dec-18 cohort retention in mth 2 looks very poor, but it is due to seasonal low in January. Similarly, the Jan-19 cohort is artificially strong.

Seasonality is exhibited across the diagonal red line from Dec-18 as customers churn into Christmas

Below we see seasonality incredibly clearly in Intuit’s 7 years of cohort numbers. Intuit is a global software company that specialises in financial software. The product revenues tracked below primarily include TurboTax (income tax returns product) and Quickbooks (SMB accounting and financial management software).

The seasonal effect for Intuit happens to occur around April which is not surprisingly tax return season. This pattern tells a story of usage behaviour in cohorts, who sign-up at the start of tax season to use TurboTax and switch it off until next year.



The ultimate example of Seasonality—table from <https://tribecap.co/a-quantitative-approach-to-product-market-fit/>

To wrap this up...

Hopefully you’ve taken away a more nuanced understanding of cohorts and can appreciate the stories that can be told. At EVP we monitor these metrics for every company we work with. Reach out to us if you want help understanding how your business is going and how you can improve performance.