


# Daniel McCarthy

Assistant Professor of Marketing at Emory University - Goizueta Business School
22 articles

There has been a lot of buzz surrounding Blue Apron's S-1 for an initial public offering. Many analysts have dived in it to glean insights into Blue Apron's valuation. However, virtually all the analysis has focused upon traditional financial metrics such as revenues and net profits (e.g., "revenues and losses are both going up!") or surface level analysis of their disclosed customer metrics (e.g., "average orders and average revenue per customer were down year on year!").

But what did the S-1 actually say about the true underlying propensity of subscribers to acquire and stay with Blue Apron? There has been far less written about this, because as Tren Griffin and PitchBook Data had noted, churn metrics were surprisingly absent from the filing (in stark contrast, for example, to the treasure trove of customer data disclosed in furniture e-commerce retailer Wayfair's S-1). By the end of this post, I will show you that we can back out an estimate of what their retention curve looks like from the few figures that Blue Apron has disclosed. While business intelligence firms such as 1010data, SecondMeasure, and Cardlytics have estimated what Blue Apron's and subscription meal delivery businesses' "survival curve" may look like, they are doing so off of panels of users that may not necessarily be reflective of Blue Apron's overall subscriber base. Working with audited data from Blue Apron itself about Blue Apron's entire subscriber base can be a useful complement, despite the limited nature of the disclosures they have put out to the public.

Cutting to the punch, this is what Blue Apron's data tells us about subscriber retention, and more below on how we get here:

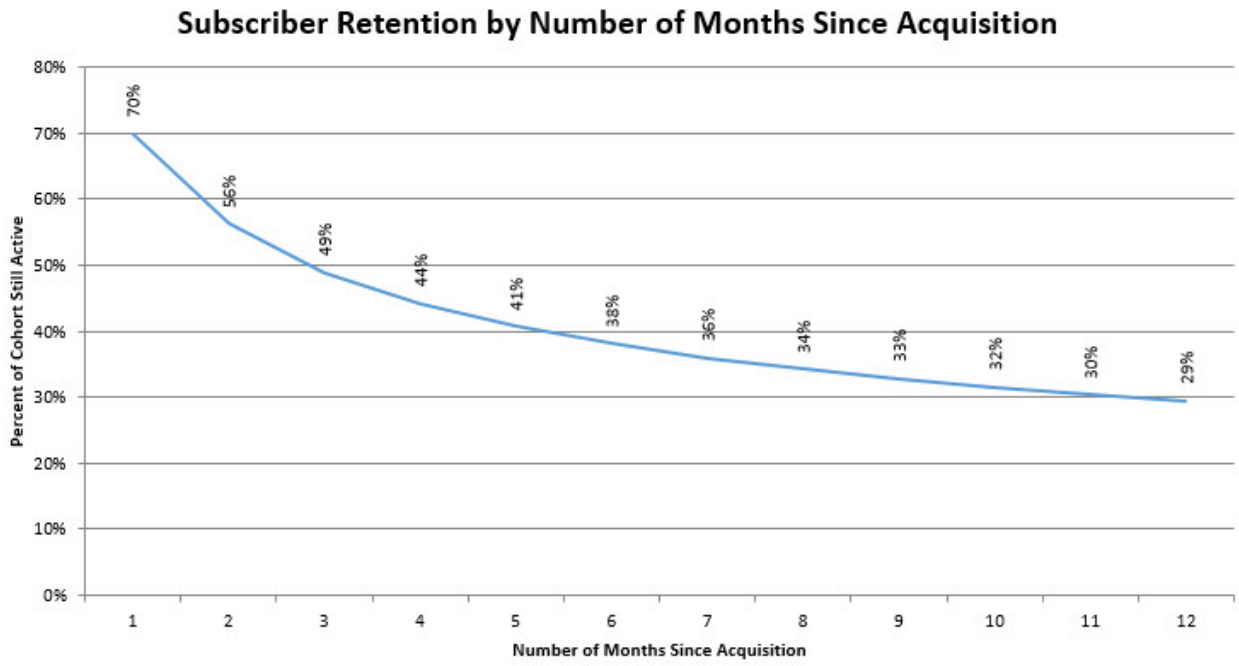


**Blue Apron Discloses Customer Churn... Sort Of**

While Blue Apron didn't *explicitly* disclose churn metrics, they did *implicitly*. By disclosing their "cost per customer" and their historical marketing expenses, we can back out how many customers they acquired between Q1 2014 - Q1 2017.

Doing out the math, they acquired 2.9MM customers in this period. Assuming they had a very small number of customers at the start of 2014 (a seemingly reasonable assumption), this would

suggest that they acquired 2.9MM customers and then lost another 1.9MM of them to finish Q1 2017 with 1MM subscribers. But is this good or bad?

**Using Probability Models To "Back Out" Blue Apron's Customer Retention**

Here's the fun part - using some survival model math from my paper on customer based corporate valuation paper for subscription businesses with Peter Fader (whr.tn/CorpValPaper1 for download; here for Journal of Marketing version) , we can get the quarterly customer behavior that is most consistent with the customer data that Blue Apron has disclosed... even though Blue Apron never disclosed it! Check out our paper for more on how we do it. Wonkish comments are at the bottom of this note.

**The Results: Low Customer Retention**

Fitting the model for customer acquisition and retention at Blue Apron yielded some eye-opening findings about Blue Apron's customer retention: 62%+ of customers churn within 6 months (*NB: given the sparsity of Blue Apron's disclosures and the front-loaded nature of the acquisition curve, I would only be confident in the retention curve up to one year out)*. If retention of future customers is consistent with that of historically acquired customers, then this would suggest that Blue Apron's future success will be heavily reliant upon future customer acquisition – finding many new customers to acquire, and being able to acquire them cheaply. It will be hard to get off the "acquisitions treadmill" because the firm loses customers very quickly after acquiring them.

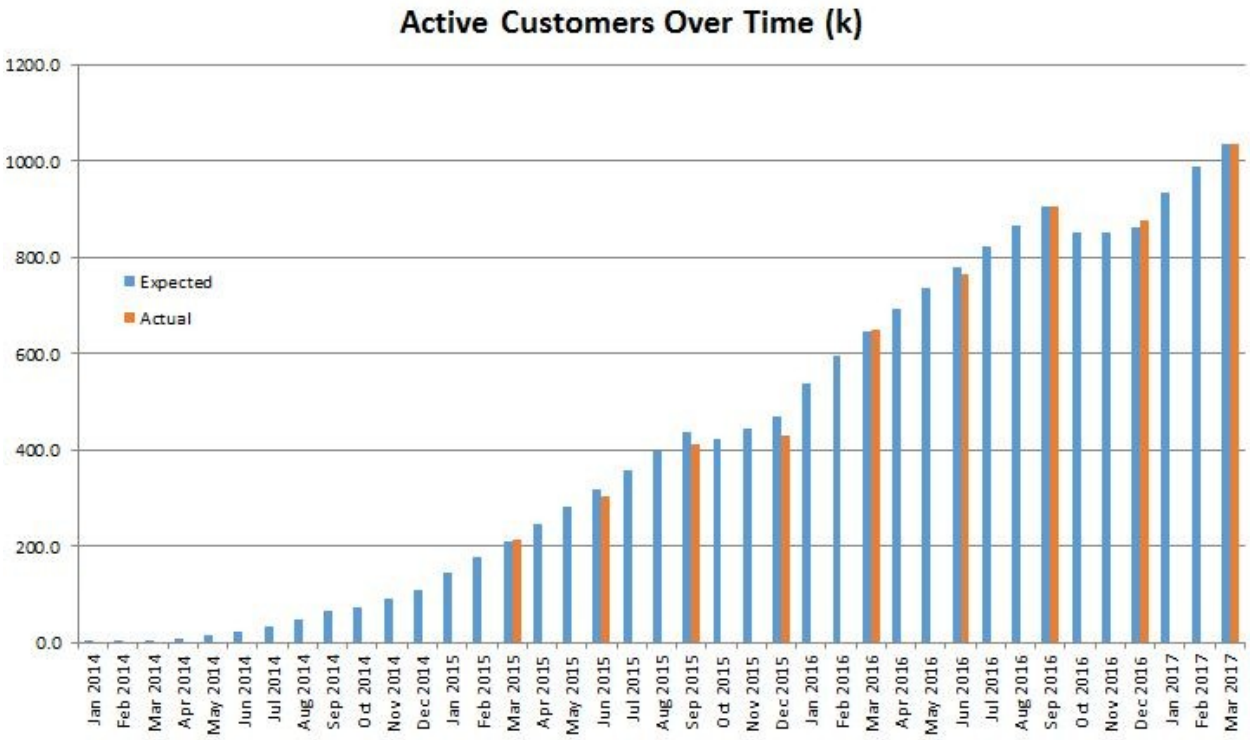


More fundamentally, it will be harder for Blue Apron to turn a profit in the future because it will need to continue spending large sums on customer acquisition expenses. Blue Apron has spent $94 to acquire new customers on average historically, yet makes around $25 per month in gross profit per active subscriber (using Q1 2017 numbers). Even with very optimistic assumptions about what costs are variable in nature, Blue Apron likely hits break-even when customers have been around for at least 4.5 months.

Moreover, as pointed out by Valery Rastorguev, our model would also suggest that the cost per customer acquisition has been rising lately. While $94 is the historical average amount that Blue Apron has spent between Q1 2014 and Q1 2017, more recent acquisition costs are likely well north of this (and if they aren't, then acquisitions are even more front-end loaded than my model would suggest, implying even worse retention figures). The plot below contains my estimates for CAC over time. I estimate CAC in Q1 2017 to be $169, which would imply a break-even point on newly acquired customers of ~8 months, implying that Blue Apron will lose money on 66% of its customers.

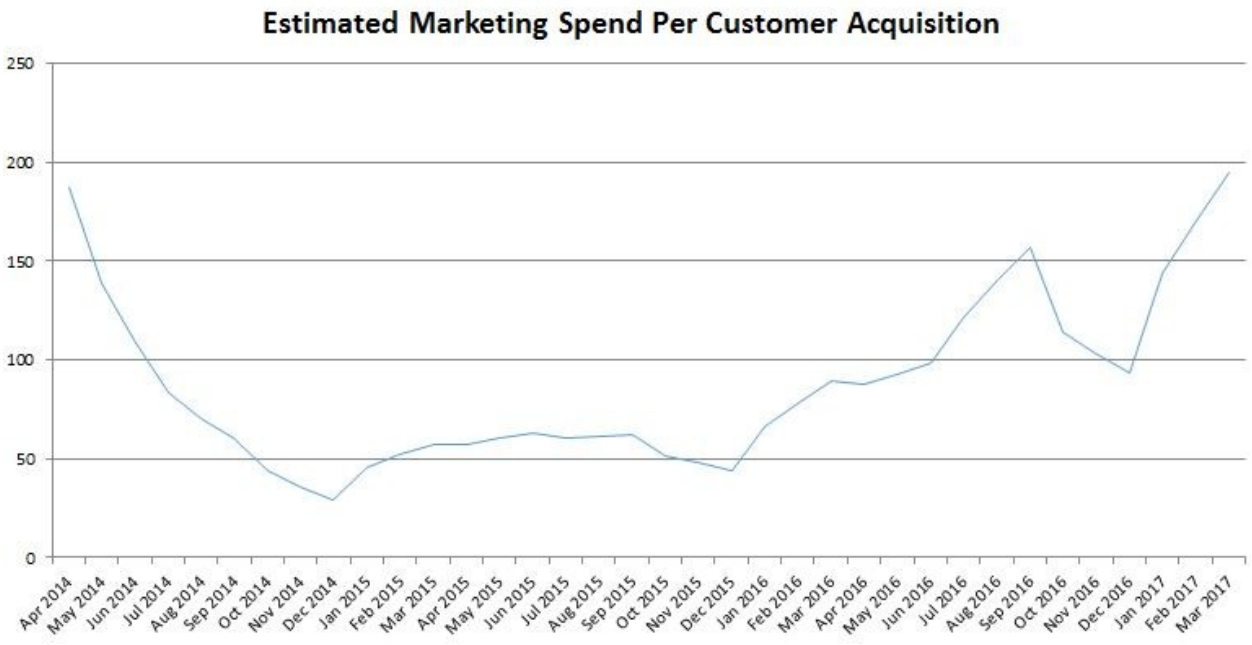


While these data points make it hard to be bullish on Blue Apron's valuation, they make me even more bullish on "customer-based corporate valuation." Investors should be demanding metrics such as the number of active customers and the number of customers acquired and/or lost over time. More data would certainly lessen the uncertainty associated with the estimates provided in this analysis, which begs the question - is Blue Apron offering a relatively meager serving of customer metrics because of the troubling implications those disclosures would have for customer retention?

*I am very grateful for many insightful conversations with Mark Zubenko, Eric Schwartz, and Valery Rastorguev. Errors and omissions are all mine.*

Wonkish Comments on model:

1. I use two hazard models -- for the acquisition and the retention of customers over time -- to estimate monthly customer behaviors most consistent with Blue Apron's quarterly disclosures.
2. The estimation procedure is virtually the same as what I laid out in the customer-based corporate valuation paper.
3. I used "smoothed" monthly estimates of Blue Apron's marketing expenses based on Blue Apron's quarterly disclosures to get estimated monthly CAC figures.

# A Detailed Look at Blue Apron's Challenging Unit Economics

- Published on June 27, 2017

Credit: Blue Apron


## Daniel McCarthy

Assistant Professor of Marketing at Emory University - Goizueta Business School
22 articles


Good companies can acquire many customers cheaply, retain existing customers for extended periods of time, and generate a lot of revenue while those customers are alive. Putting it simply,

the litmus test of any company's financial success is the ability to acquire many high lifetime value (LTV) customers. Being LTV-centric is at the heart of being customer centric.

Does Blue Apron, which recently priced its IPO at a very healthy ~$3 billion implied valuation (or almost 3.5 times trailing twelve month revenues), pass the test? In my last note on Blue Apron, which was recently cited in the Wall Street Journal, I showed that while Blue Apron disclosed nothing *explicitly* about its customer retention, and very little about how its customer acquisition cost (CAC) has been changing over time, it disclosed just enough to use an extension of the modeling approach that I advocated in a recent journal article to "back out" what these figures are most likely to be. The conclusion: Blue Apron doesn't retain customers for very long, and the cost to acquire customers has been on the rise lately. These are important ingredients to the overall customer-based corporate valuation recipe. At the same time, there is a lot more that we can learn from Blue Apron's S-1 disclosures.

I went back and built a much more complete model to leverage all the data that Blue Apron has disclosed. I explicitly model how customers are acquired, how long they remain customers before churning, how many orders they make while they are retained, and how much they spend on each of those orders. This more general model allows us to incorporate all the metrics that Blue Apron has disclosed, such as six-month cumulative revenue for annual customer cohorts. It allows us to refine answers to previous questions, such as what Blue Apron's retention curve looks like, and answer new ones, such as how the post-acquisition profitability of customers has been changing over time, and whether younger customers generate more revenues as they age or not (e.g., that the customers who stick for a long time around reorder a substantial amount).

The results continue to suggest challenges ahead – retention is even weaker than I had originally estimated it to be, new acquisition cohorts are generating less revenues than old ones, and as customers age, they spend less and not more with the firm. In recent months, I estimate that Blue Apron is losing money on ~70% of the customers that it acquires. I dive into the model briefly next, before expanding on these conclusions.

**The Model**

My model for the acquisition and retention of users remains the same, using only the cost per acquired customer, historical marketing expense, and active customer data as inputs. However, I built additional models for how many orders customers make while they are alive, and how much they will spend on a particular order. I estimate parameters for each of these models so that what we expect the data to be is as consistent as possible with the disclosed data.

The resulting relatively simple composite model does an excellent job of fitting the observed data. As shown below, it provides a very reasonable fit to all the data – the number of active customers, total customer acquisitions, orders, revenues, and cumulative revenue per acquired customer metrics. I provide a series of charts summarizing this performance below.

Quarterly total number of active customers:

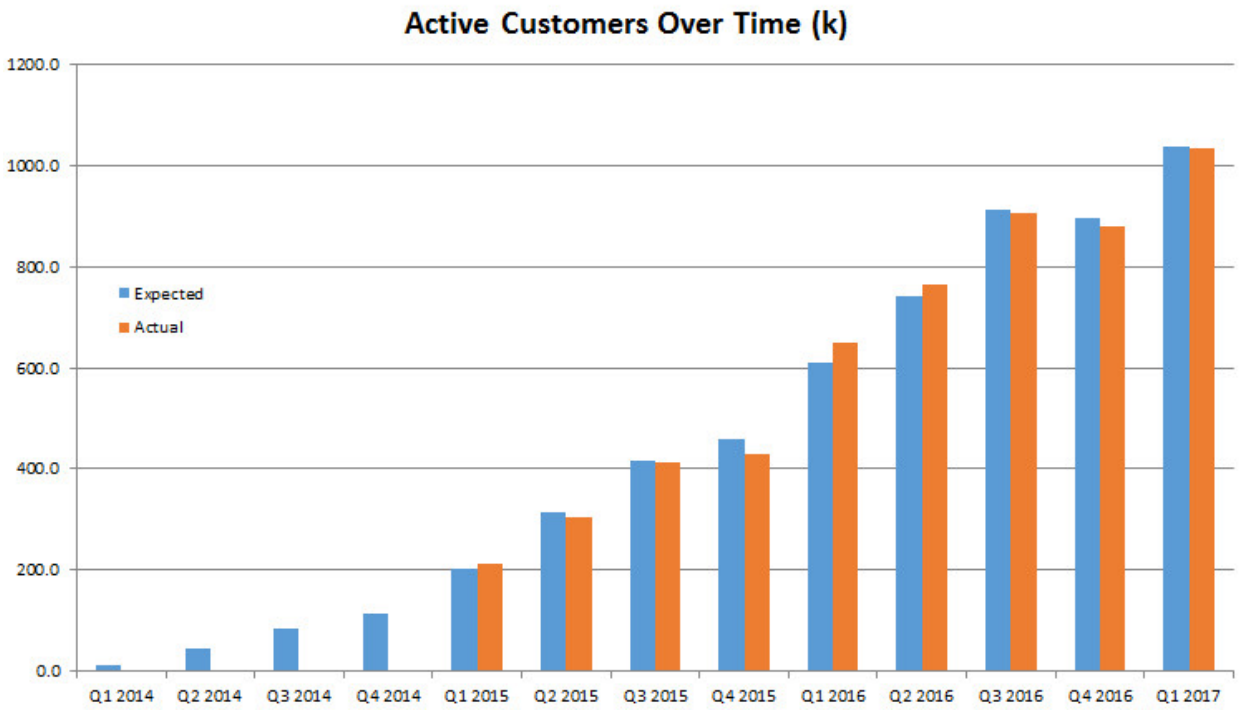


Cumulative customers acquired, Q1 2014 to Q1 2017:

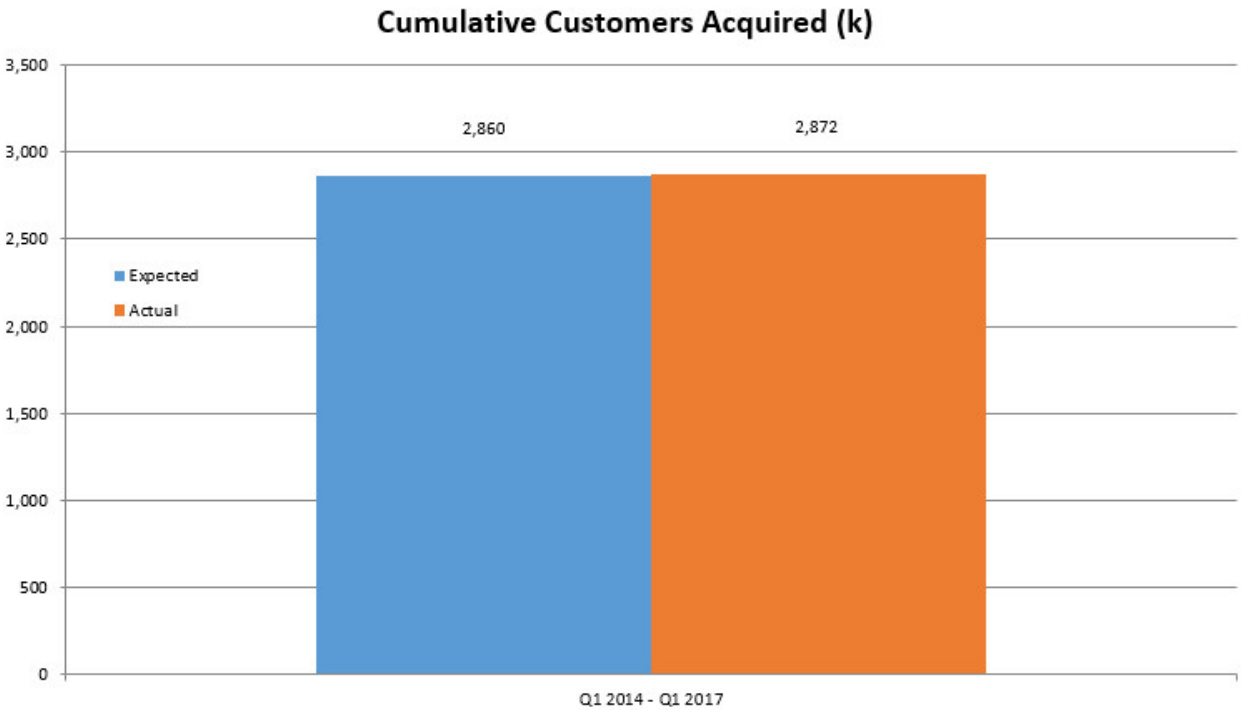


Quarterly total orders:


Quarterly total revenue:

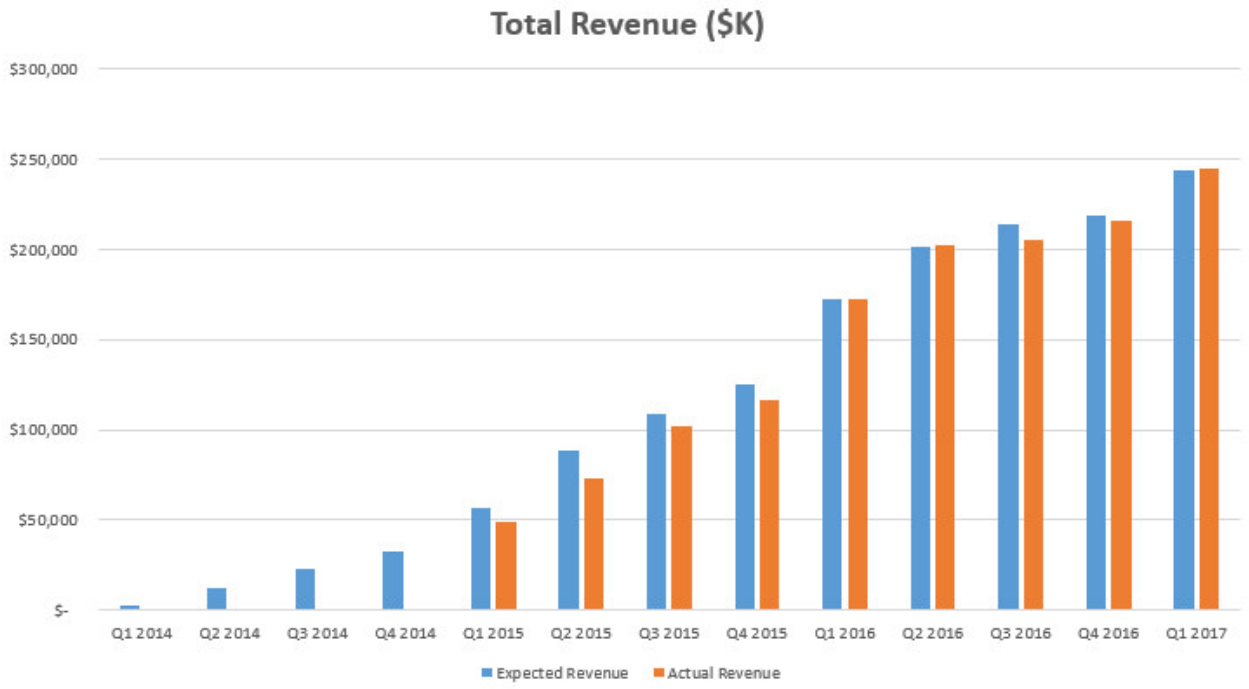


Cumulative net revenue per acquired customer for customers acquired between Q1 2014 and Q1 2017, 6 to 36 months out:

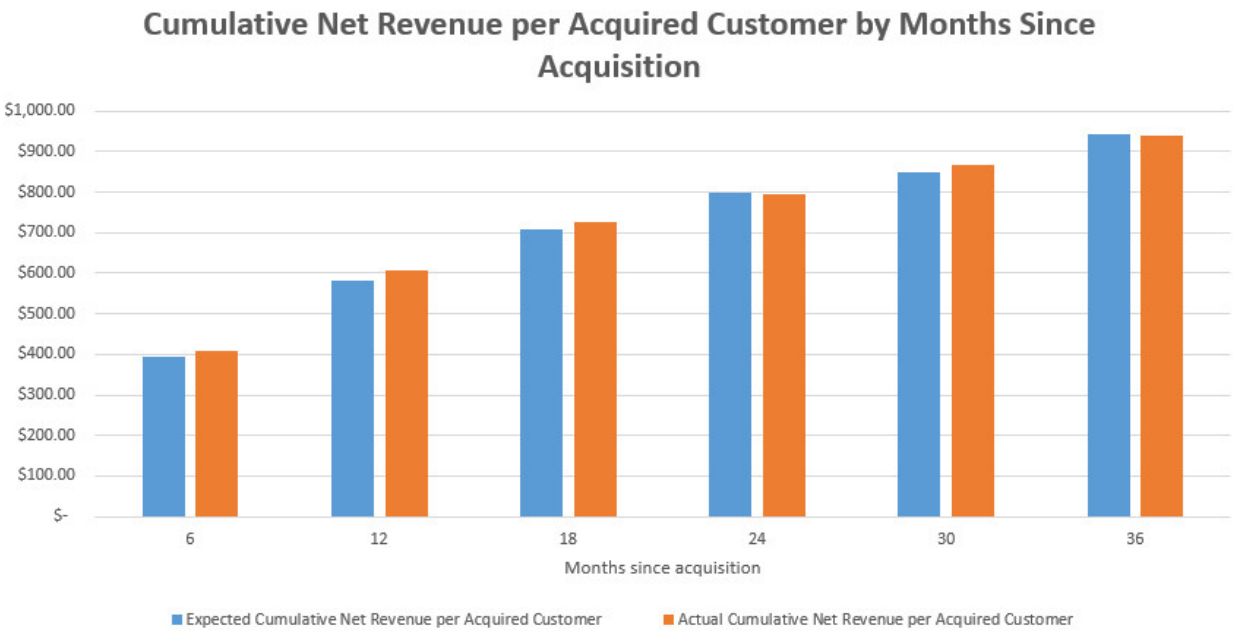


Cumulative net revenue per acquired customer over next six months for customers acquired in 2014, 2015, and 2016:

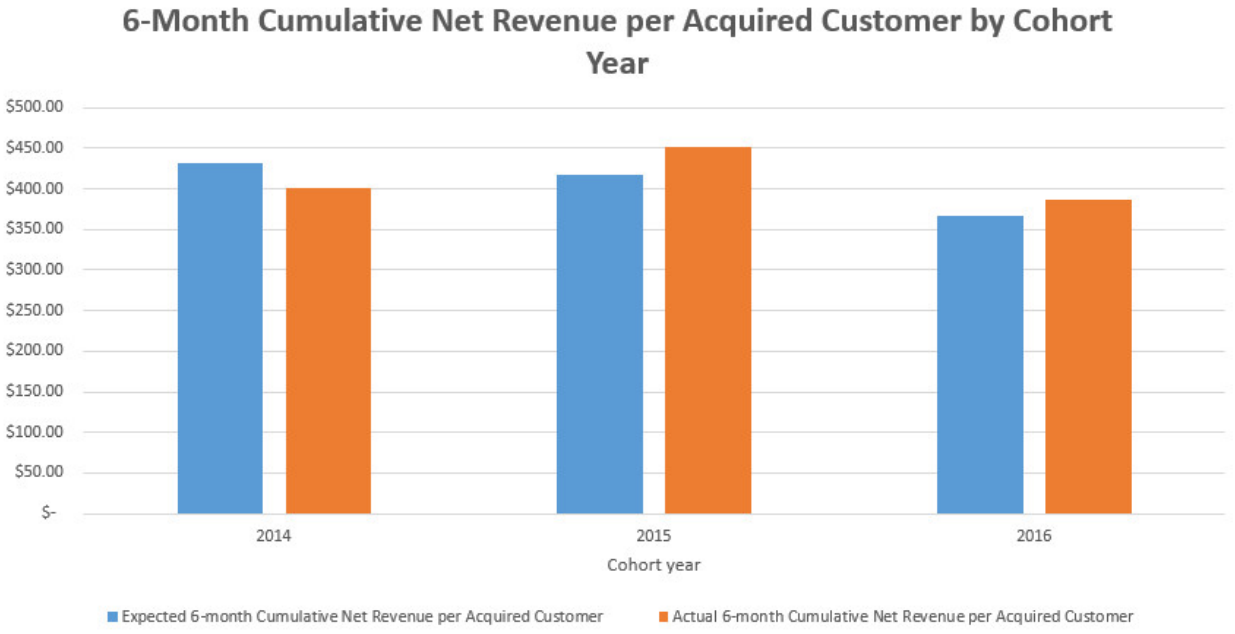


The fact that my relatively simple model is consistent with the data along so many key dimensions at the same time provides some comfort that we can trust the results of the model. Let's discuss those results next.

**The Results: Anti-stickiness – Low Retention and Declining Revenue per Customer, Over Time and Across Cohorts**

Here is a summary of what I found from the deeper dive:

1. *The retention curve is worse than I originally had estimated it to be*. While my substantive conclusion remains the same, I estimate that **72%** of customers will churn by the time they are six months old. Because Blue Apron cannot retain customers for extended periods of time means that CAC is effectively part of cost of goods sold. CAC should go down relatively sharply over time as a percentage of sales at healthy businesses, as sales are increasingly derived from loyal customers who have been around for a while. When customers churn out very quickly, that pool of loyal customer revenue remains small, making CAC effectively variable in nature.

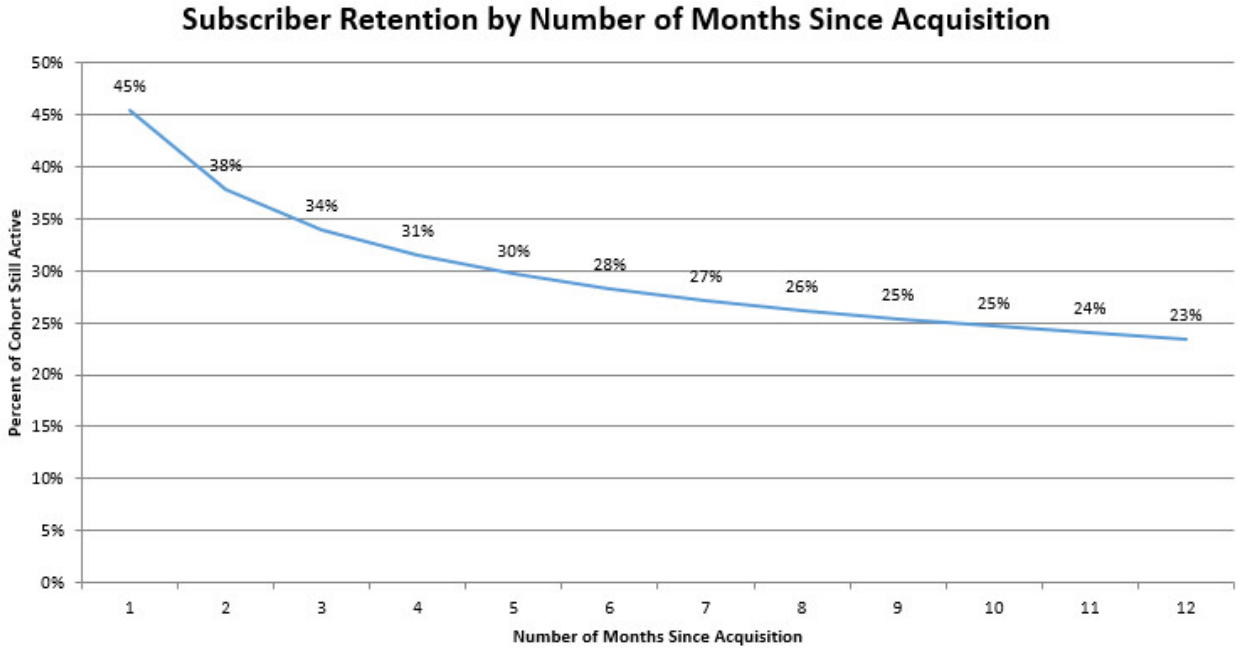


2. *The revenue that Blue Apron is generating from more recently-acquired customers is less than from customers acquired in the past*. Every new acquisition cohort generates, on average, about $7 less in revenues over the next 6 months than the cohort which preceded it, which adds up quickly over time. In other words, while the cost to acquire new customers is going up, the go-forward value of those newly acquired customers is going down. Both trends are driving LTV lower over time. I suspect that this is due at least in part to the vast sums of money that Blue Apron is spending upon subscriber acquisition expenses (SAE). It is very common to see LTV go way down when SAE goes way up.

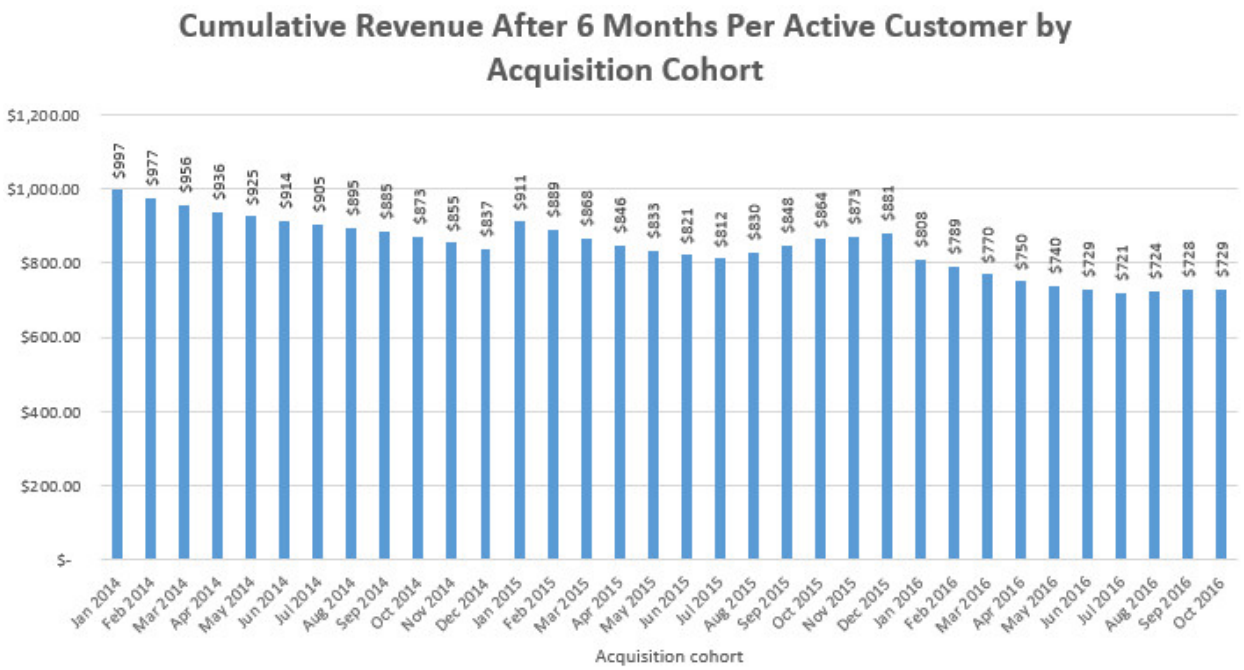


3. *While customers are alive, the amount of revenue that Blue Apron generates from them tends to go down, not up, over time.* This makes it unlikely that long-time loyal customers will "bail out" the firm because they are also high spenders, a common trend at mobile gaming companies, for example – in fact, we infer that the opposite has been taking place. As customers get older, they place fewer orders on average, which is only slightly offset by a marginal increase in spend per order over time. Customers are not "sticky." Moreover, at subscription-based businesses like Blue Apron, there is only so much that big spenders can spend, while there is no such upper bound at non-subscription businesses.

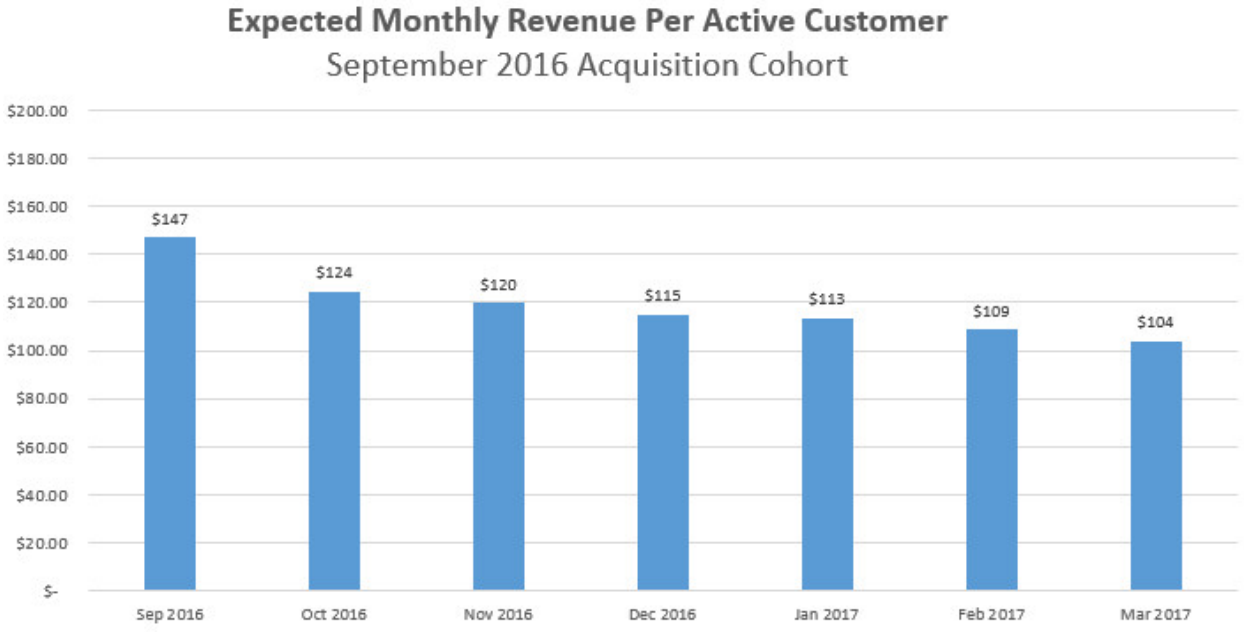

4. *70% of recent Blue Apron customers will not break even*. We estimate that CAC in Q1 2017 is $147. To break even at this CAC, new customers must generate at least $565 of net revenue (i.e., gross revenue minus returns and promotional discounts), assuming Blue Apron's variable contribution margin is equal to ~26%. The chart above shows that newer customers must remain subscribed for about 4.5 months to generate this much revenue. However, almost 70% of customers churn by this time and thus do not break even. Even though Blue Apron turns a profit on the remaining 30% of customers, the break-even point is moving farther away with every new cohort due to declining revenue and growing CAC for newer customers.

In summary, this customer-based analysis spells trouble for Blue Apron, with important measures of customer health in decline. Amazon's recent acquisition of Whole Foods is likely to make it even more difficult to keep those Blue Apron subscribers coming back. I recommend that Blue Apron redouble its efforts upon activities that will make customers "sticky" in the long run. Investors are clamoring for customer metrics so that they can go beneath surface-level financial metrics to better understand Blue Apron's underlying unit economics. I hope that this analysis takes investors a step closer to what they are looking for, and that Blue Apron will begin disclosing a few more.

*A big acknowledgement goes to Valery Rastorguev. All errors and omissions are mine.*