Customer LTV Model

Google Sheet that houses the model:
https://docs.google.com/spreadsheets/d/1RXrM9dYTYcey-TaFGNqnI6kYPPvR9XHSIwACICjIaVM/edit#gid=720467450

Sigma workbook with modeling inputs:
https://app.sigmacomputing.com/qeepsake/workbook/Unit-Economics-Modeling-Inputs-279Au2Bvejyca0yZMZTUbs?:nodeId=5Wf0uDLuOkFf290yxrYNyY

💵 Archive of Customer LTV model updates

What is the purpose of this model?

We want to be continuously tracking how “valuable” our customers are to us. “Value” is a fuzzy notion that can include many factors (not all of them revenue), but generally the way that’s most useful for us to answer this question as a business is to try and project:

> ***How much revenue can we expect over the lifetime of a customer that begins on a paid Qeepsake plan today?***

This question is useful to ask because it lets us think and strategize about things like:

- How much should we pay to get a paying customer in the door with our advertising?
- How much revenue can the business count on from its existing customer base in the future?

What is the process for answering that question?

Data team keeps a model that attempts to answer this question for each year of Qeepsake's existence (ie, "How much lifetime revenue can we expect from customers that began on paid plans in 2016/2017/2018... etc.)

This model is updated for the current year every quarter, as more information from the current year comes in and we can more accurately project the current value of customers entering Qeepsake.

Data team may update the way the model works each time it does a quarterly update! We will try to keep this notion in line with the current methodology.

How does the model work?

How much revenue can we expect over the lifetime of a customer that begins on a paid Qeepsake plan today?

This is a difficult question to answer. It is a prediction about the future of customers who are starting their Qeepsake journey today, and we cannot teleport into the future to see what their Qeepsake journey "ended up being like". Instead, we need to use information about *past* customers and their journeys.

The advantages of this are:

- We can actually know what these customers ended up doing later on in their Qeepsake journey, and exactly how much revenue Qeepsake has gotten from them
- We have a growing history of users who have been on Qeepsake for many years, which gives us a lot of data to use for future projections
- The further back in the past we go, the more data we can collect and the longer "journey"s we can measure

There are also disadvantages:

- Even users who have been on Qeepsake for many years are not at the end of their Qeepsake journey. We still have plenty of users from 2016 who are actively journaling

- The further back we go, the more disconnected we are from the experience of the user entering Qeepsake *today*. Someone's journey will almost surely be different if they started on the app vs. over texting, and *the app didn't exist until 2018*. Same goes for Milestones (2021) and Video (2022).

Ultimately, we need to balance these tradeoffs and use the data we have available to make a **best educated guess** at how much revenue we can expect over the lifetime of a customer who onboards today. To start we break down this number into smaller parts, and estimate each of those parts individually:

`Total Lifetime Revenue = (Annual Subscription Revenue + Annual Book Revenue + Annual Gift Revenue) X (Expected tenure of a Qeepsake user in years)`

More specifically, we want to measure not just LTR, but LTV: The difference is Qeepsake's profit margins on its revenue. To calculate LTV, we use the formula above, but additionally we multiply each revenue type by Qeepsake's gross margin on that type of revenue. For the purposes of simple explanation, I'm omitting this additional step, but just remember that LTV is not exactly the same thing as LTR.

So to get Lifetime Revenue, we need to estimate 4 variables:

- Expected tenure of a Qeepsake user (in years)
- Annual Subscription Revenue per user
- Annual Book revenue per user
- Annual Gift revenue per user

Expected Tenure

This one is in one way simple, and in another way, complicated.

Expected Tenure is simply: 1/(1-Renewal Rate)

Renewal rate is extremely tricky to measure, but luckily, we already put a lot of effort into measuring it, so we can lean on that.

🔄 Renewals Methodology

To summarize/simplify, for the renewal rate we use in the LTV model, we ask the question:

`What % of users who started on a paid plan in year X also started on a paid plan in year X+1?`

(Where year X+1 is year-to-date as of the most recent quarter).

Notice — This formula is fundamentally using users who started on paid plans at least a year ago to try and predict the renewal rate of users starting on paid plans today. In the future, we may want to use a predictive model based on user behavior, that takes into account the differences in behavior (coming from a different product experience) of users today. But we're not there yet, so just keep in mind that this will necessarily be a *lagging* indicator of *future* renewal rates.

**Update as of Q4 2022: Moving forward, we will be adjusting how we project renewal rates by adding in an additional step. Previously, we were using the fully mature renewal rate for the most recent year's worth of subs, which was biased against the most up-to-date estimate, because the most recent data has had less time to mature. Now, we will look at *30-day*
renewal rates for the most recent year, and multiply those by the lifetime historical average of 30-day to fully mature renewal rates.

Annual Subscription Revenue Per User

Since we only charge users once annually, this metric is relatively simple. We simply take average purchase price (usually from the most recent quarter).

The one complicating factor is discounts. Since users *do not* renew at discounted rates, we also have to account for future renewals being at a slightly higher average purchase price. That's why you'll see two different "average purchase price"s, one for new purchasers and one for future renewals. The price for future renewals is simply *average purchase price without discounts*. Note that this does assume discounted purchasers renew at equal rates to full price purchasers, but so far the data backs up that assumption.

Average Book Revenue Per User

This one is trickier to project. We know exactly how much book revenue we get, but *from what base of users is it coming?*

For a yearly LTV projection, we have to basically use one number to mean “How many users were on a paid subscription long enough to order a book during this year”?

We define that as:

`All users who started or renewed on a paid subscription at some point during last year, and onboarded at least 6 months before the beginning of this year`

From there, the calculation is simple:

`Total annual book revenue / Eligible users = Book revenue “per user” for that year`

**Update as of Q4 2022: When we’re doing a model update partway through the year, we’ll project annual book revenue by multiplying revenue YTD by the average % of total revenue we HAD seen up to this point of the year in past years.

Average Gift Revenue Per User

This is done similarly to books, except we have a slightly different notion of who is “eligible” to purchase gift memberships for their friends. It’s simply:

`How many users started on a paid plan this year?`

Then we divide total gift revenue for this year by that.