MKTG 776 PROJECT 2

April 2019

Modeling the cumulative adoption curve of ‘The Binding of Isaac: Rebirth’ after its launch in November 2014

Author: Steven Ripplinger

MKTG 776 Project 2

MODELING THE CUMULATIVE ADOPTION CURVE OF 'THE BINDING OF ISAAC: REBIRTH' AFTER ITS LAUNCH IN NOVEMBER 2014

Assignment: Using the given dataset containing the number of new adoptions per week over 113 weeks, develop a model to explain and predict the adoption curve.

EXECUTIVE SUMMARY

Model

The final model is shown below. The 'Pareto II + 3 Covariates' model provides a compelling story of the adoption curve for the in-sample period, as well as a powerful predictor of the hold-out period.

This model takes into account the expansion pack, summer sales, and the Xmas+NY effect. It has 7 parameters and results in an in-sample MAPE of 2.9% and an out-of-sample MAPE of 1.2%.

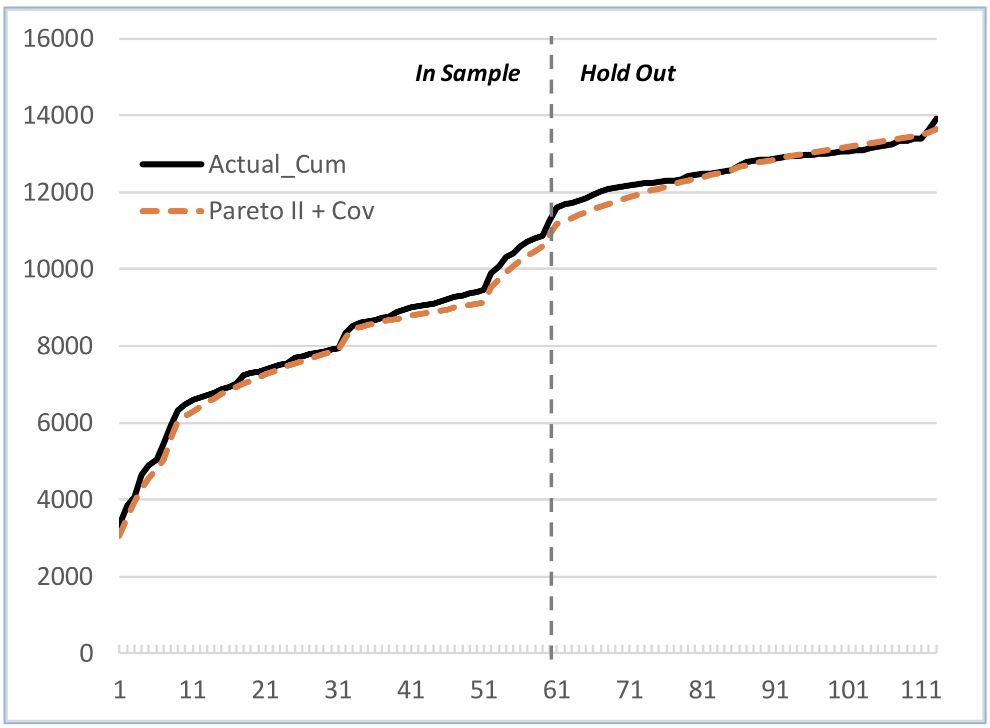

Parameters	
r	0.023
Alpha	3.290
Pre-Sale	0.025
B_Expansion	2.630
B_XmasNY	1.094
B_SummerSale	1.708
Expansion_Delta	0.249
LL_Sum	-74,133.52
BIC	148,278.55
Parameters	7
IS_Cum_MAPE	2.9%
OOS_Cum_MAPE	1.2%

Conclusions

There are several conclusions relating to the model. First, the model suggests that heterogeneity, and not duration-dependence is driving the declining hazard rate in the data, with implications including the ability to increase propensity to adopt among users. Second, the low r of 0.023 means that heterogeneity is high. There are some hard-core adopters, and varying degrees of those who may be less likely to try the game. Third, the pre-sale parameter is 2.5%, a useful metric that allows a manager to compare to other launches and provides significant insight for planning and forecasting future launches. Lastly, we can turn off the expansion parameter, allowing us to quantify new users that adopted because of the expansion pack launch.

SETTING UP THE PROBLEM:

Some Basics:

Assumption: N=100,000

Test period = week 1-61

Holdout period = week 62-113

Expansion pack = week 32

Accounting for Pre-Sales

The dataset starts with week 1 sales of 3,344, which includes many weeks of pre-sales prior to the week 1. There are two ways in which we can handle this: use a shifted model or add spike at week 1:

Shifted model: Through some research, we know that pre-sales started on September 5, 2014 (source: techraptor.net), approximately 9 weeks prior to the launch date (and start of dataset) of November 4, 2014. Therefore, we could start the model at t = -8 and still fit the model for weeks 1-61. The parameters of the model would allow us to backout the pre-sale sales by week. The benefit of this approach is that it would not burn any parameters, and it could provide some insight into the pre-sale adoption curve. However, the problem is that the nature of the pre-sale period adaptors is likely different than post-launch adapters, and back-casting the model without the weekly data to validate would be somewhat speculative. The managerial benefit of this process is questionable without validation.

Spike at week 1: The other approach is to simply aggregate all pre-sales and provide a probabilistic estimate of the amount of pre-sales. To do this, we add a spike at week 1. While this adds a parameter, this is a useful metric to have. In fact, it is extremely useful for a manager to know what % of the population (and/or how many people) purchased the product prior to launch. This metric could be compared to other launches and provides significant insight for planning and forecasting future launches.

The approach we will use across all models is to allow a spike at week 1 for pre-sales. It makes the most sense because while it burns a parameter, it is simple to understand, easy to implement, and it is a useful metric to know. Further, if a particular model doesn't need the spike, it will just be zero, and we can re-run it to save the parameter.

Expansion & Covariates

Afterbirth, the expansion for Rebirth, was launched on October 30, 2015 (week 32). There is a clear uptick in adoption after this expansion pack.

Promotions are run during holiday season, early summer, and other times.

We will discuss accounting for these effects and other covariates in later sections.

APPROACH #1: BASELINE MODELS

The Individual-Level Model

First, we should establish that at the individual level, we can observe time (since product launch) of adoption. This is what we are modeling.

We first consider the exponential distribution. The exponential distribution produces a near-linear CDF during the analysis period and fails to adequately incorporate the actual shape of the adoption curve. This is not a good distribution on its own because it has a stationary hazard function (static at λ), whereas the actual data has a declining hazard function. Exponential distribution essentially says that the probability of adoption, given that it has not occurred, does not change over time (duration independence). We will discuss the merit of heterogeneity versus duration dependence in the following section.

The Weibull distribution incorporates duration dependence. Fitting this model results in significantly better fit. This says that everyone has the same propensity which declines over time. The declining adoption curve is a function of people being less willing to buy as time gets further from the launch date.

Figure I: Exponential vs. Weibull Distribution

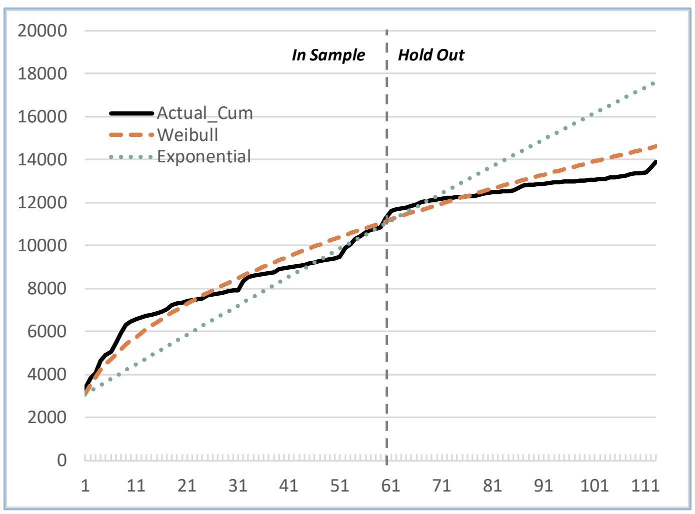

	Exponential	Weibull
Lamda	0.0015	0.0106
c		0.5425
Pre-Sale	0.0294	0.0203
LL_Sum	-76,840	-76,091
BIC	153,691	152,193
Parameters	2	3
IS_Cum_MAPE	13%	6%
OOS_Cum_MAPE	14%	4%

It is useful to look at the hazard function to see what the curve looks like. In this case, with c = 0.54, it has a downward sloping curve, negative duration dependence, which means that as time goes on, the probability of a given individual adopting, given he/she has not adopted yet, decreases. It decreases rather quickly for the first ten weeks, and then stabilizes. The story here could be that potential adopters are most likely to adopt close to launch, and then less likely over time, perhaps because they will disadvantaged by being too far behind other gamers that have already been playing for a while. We will discredit this story shortly.

Figure II: Weibull Distribution Hazard Function

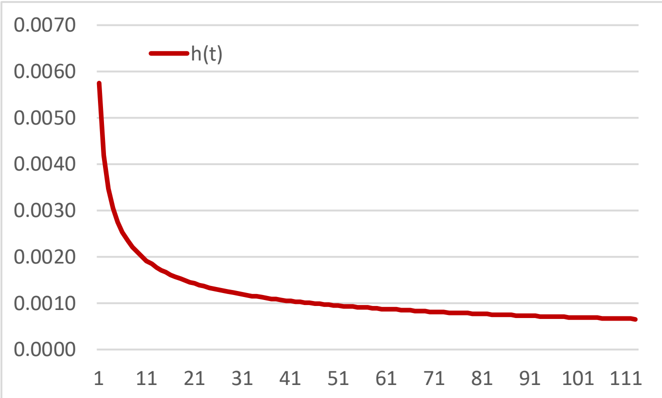

Adding Heterogeneity – Parametric Approach

It is worth looking first at the Pareto II (Exponential Gamma). Adding heterogeneity to the exponential via a gamma function results in the Pareto II model, which has strong fit as depicted by the green dotted line in Figure III. By incorporating heterogeneity into the exponential, there is a resulting declining hazard function consistent with the data. However, by using the Pareto II, we would be assuming that the decreasing hazard rate in the data is entirely due to heterogeneity and not due to changes in individual level propensities to adopt over time (duration dependence).

On one hand, duration independence may not make sense because one could argue that games such as this are popular for a time and then fade away. They have a finite shelf-life, and thus it stands to reason that as a game gets older, the probability of a given individual adopting changes.

On the other hand, it is not unreasonable to believe that heterogeneity, and not duration dependence, explains the declining hazard function in the data. The case here is that a game such as *Rebirth* is not a typical game from the old cd/xbox/disc type era where games become obsolete quickly and new complete versions of games are released every couple years. Rather, games such as *Rebirth* that are released on Steam are continuously improved with "over-the-air" updates. Indeed, we see that expansion packs were released in October 2015 and again in January 2017. Because of this, we put forth the case that *Rebirth* has lower duration dependence, and instead, it is heterogeneity among gamers that is largely causing the shape of the adoption curve.

Turning back to the Weibull, we can add heterogeneity via the gamma function, resulting in the Burr XII. The Burr XII model results in ultra-high r and alpha parameters, converging back to the Weibull. This could indicate that there may not be any heterogeneity. However, we know that Excel Solver has a difficult time separating the effects of heterogeneity and duration dependence, so we cannot reject heterogeneity outright.

Indeed, we could qualitatively make the case that there is heterogeneity among the population. Some people are hard-core gamers and will buy every game that will come out. Others are more selective. Some might only adopt if they get a discount. People are generally heterogeneous, even when it comes to cult games.

<u>For these reasons, while in reality there is likely some mixture of heterogeneity and duration dependence, we move forward with the assumption that heterogeneity is the overwhelming force driving the declining hazard rate in the data. Thus, the Pareto II model is a viable model to move forward with.</u>

Figure III: Exponential / Pareto II / Weibull / Burr XII

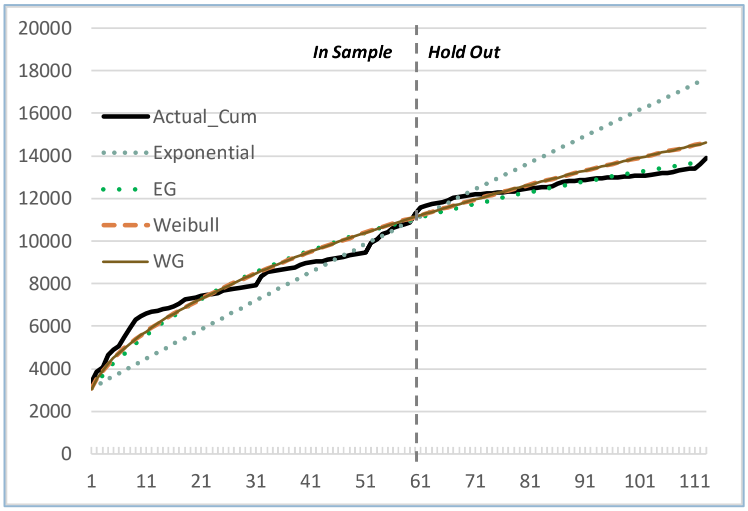

	Exponential	EG	Weibull	WG
Lamda	0.0015		0.0106	
r		0.0585		864.2138
alpha		16.6420		#########
c			0.5425	0.5425
Pre-Sale	0.0294	0.0294	0.0203	0.0203
LL_Sum	-76,840	-76,421	-76,091	-76,091
BIC	153,691	152,853	152,193	152,193
Parameters	2	3	3	3
IS_Cum_MAPE	13.1%	6.8%	5.7%	5.7%
OOS_Cum_MAPE	14.3%	1.9%	3.8%	3.8%

The Pareto II (EG) is the lead model at this point. But we can do better.

ADDING HETEROGENEITY – LATENT CLASS MODELS

The other way to incorporate heterogeneity is via a latent class model. After extensive analysis, it was determined that these models are inferior with poor out-of-sample fit. The data does not need multiple segments.

Figure IV: Latent Class Models

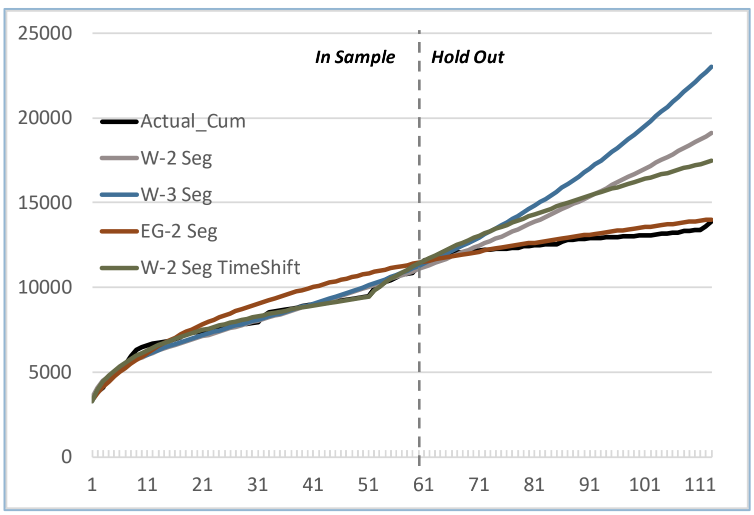

ADDING COVARIATES

Moving forward with the straight Pareto II model, covariates are added to account for other factors driving adoption.

Expansion

On October 30, 2015 (source: wikipedia), week 52, the expansion pack, *Afterbirth*, was released. This is an important covariate because the expansion pack significantly improved the game and added new features. This, along with publicity and reviews that went along with the release, increased the adopters to the game. This is clearly evident in the data.

This is not a one-week event, but rather a structural shift. We therefore take this into account with a time-varying co-variate.

$$X(t) = 1 – d[1-e^{-(t-t52)}]$$

Where d = the fraction of the baseline

This results in a structure shift that decays over time starting at week 52.

This coefficient is statistically significant based on an LRT between the full model and without this covariate.

Summer Sales

Steam has annual summer sales. Through google search history and news articles, we can determine the exact dates of the summer sales and include this as a covariate:

- 2015: June 12-21 (Weeks 32 & 33) – 50% discount on *Rebirth*
  - Source: https://wccftech.com/pree3-weekend-gaming-deals-price-cheats/
- 2016: June 23 – July 4 (Weeks 86-87) - 40% discount only for one day (Jun 24)
  - Source: https://www.vg247.com/2016/06/30/steam-summer-sale-deals/

When including this as a covariate, there is a nuance. In 2016, the format of the Steam summer sale changed where there are discounts throughout the two weeks, however each day has a daily deal with bigger discounts. These are more meaningful. Rebirth was featured as a 40% discount on June 24, occurring in week 86.

We will account for this change in sale practice and the difference in discounts as follows:

- Week 32 & 33: 1
- Week 86: 0.4 = 0.8 * (0.5) (reflects 40% discount vs. 50% and cut in half to reflect 1 day)
- Week 87: 0.2 (half of prior week to reflect steam sale but no daily discount this week)

This coefficient is statistically significant based on an LRT between the full model and without this covariate.

Holiday Season

With the hypothesis that adoption increases during the holiday season, we added a binary 1/0 covariate with the last 4 weeks of December coded as 1 each year.

This coefficient was not statistically significant based on an LRT between the full model and without this covariate. We discard this covariate for the final model.

Xmas+New Years

Continuing with the theory that a season effect may be a broader holiday increase and then another effect during Christmas week, another covariate was added as a 1/0 binary with Christmas and New Year weeks coded as 1.

This coefficient is statistically significant based on an LRT between the full model and without this covariate.

Covariate Summary and Testing Coefficient Significance

First, the covariates were tested individually to determine significance.

Figure V: Covariate Coefficient Significance

Coefficient Testing						
	LL	P	LRT	DF(dif)	p value	Significant?
Full	-74133.5	8				
-B_Exp	-75399.0	6	2530.929259	2	0.0000	*Significant*
-B_Hol	-74133.5	7	0.028275338	1	0.8665	*Not Significant*
-B-XmasNY	-74447.2	7	627.4522022	1	0.0000	*Significant*
-B-Summer	-74614.9	6	962.713082	2	0.0000	*Significant*

Holiday coefficient was thus discarded.

Figure VI: Pareto II versus Pareto II + Covariates Significance Test

We can see that adding the covariates to the EG model creates a better model

EG vs EG+Cov(excl holiday)						
	LL	P	LRT	DF(dif)	p value	Significant?
EG	-76420.51	3				
EG+Cov	-74133.52	7	4573.976359	4	0	*Significant*

FINAL MODEL AND INTERPRETATION

The final model is shown in Figure VII. The 'Pareto II + 3 Covariates' model provides a compelling story of the adoption curve for the in-sample period, as well as a powerful predictor of the hold-out period.

This model takes into account the expansion pack, summer sales, and the Xmas+NY effect. It has 7 parameters and results in an in-sample MAPE of 2.9% and an out-of-sample MAPE of 1.2%.

Figure VII: Pareto II + Covariates - Cumulative

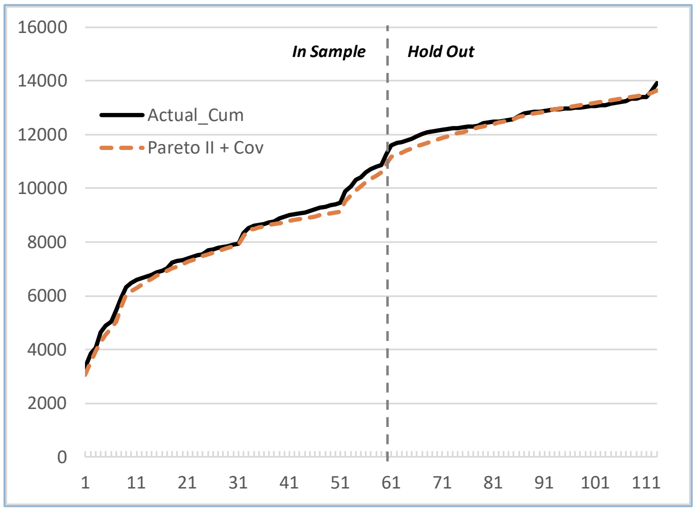

Parameters	
r	0.023
Alpha	3.290
Pre-Sale	0.025
B_Expansion	2.630
B_XmasNY	1.094
B_SummerSale	1.708
Expansion_Delta	0.249
LL_Sum	-74,133.52
BIC	148,278.55
Parameters	7
IS_Cum_MAPE	2.9%
OOS_Cum_MAPE	1.2%

Another way to analyze this model is by looking at the incremental adopters. We can see that this model is fairly good – it captures the general trend and many of the spikes.

Figure VIII: Pareto II + Covariates – Incremental Line Chart

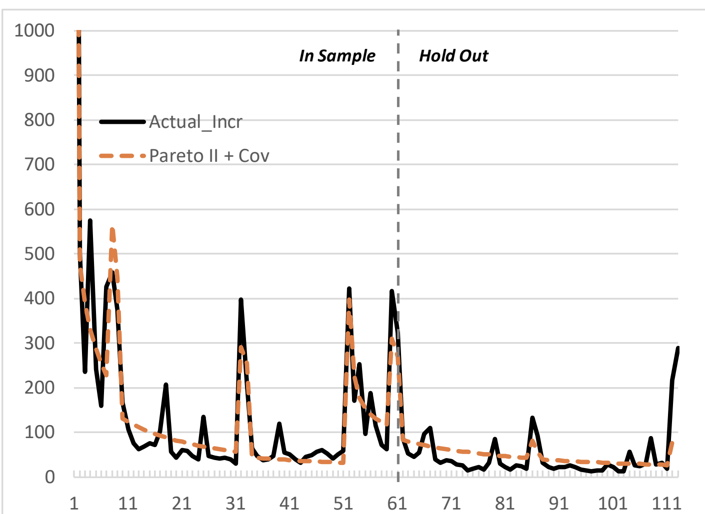

Admittedly, this model is not perfect. In particular, the model slightly underpredicts adopters just prior to the expansion release in October 2015, and overestimates adopters in the weeks after the release. This is likely just a timing effect (some adopters from the expansion just adopted early).

It is possible to backout covariates until the fit is perfect, however this is overfitting and is not going to be overly useful, especially as a predicter.

Potential improvements for further analysis

There are some spikes not accounted for the model that are likely the result of flash sales or reviews in key trade publications. With more data, these effects could be captured more accurately.

Of note, the there is an uptick in adopters in the last two weeks of 2016. The model only partially captures this with the Xmas+NY covariate. What is likely happening here is the forthcoming release of the *Afterbirth+* expansion pack in early 2017. Some sort of "hype" covariate could be added to capture sales in advance of expansion releases; however it is difficult to include this without overfitting.

CONCLUSION & KEY MANAGERIAL INSIGHTS

There are a few key insights from the model.

First, the EG model suggests that there is no time-dependence, and rather heterogeneity is the key explanation behind the declining hazard rate. This is positive as it means the company can promote the game and attempt to increase buying propensity among those with lower propensity, which could increase adoption.

Second, the low r of 0.023 means that heterogeneity is very high. There are some hard-core adopters, and varying degrees of those who may be less likely to try the game. Similar to the previous point, it means that the company could segment customers and try to increase propensities to adopt through marketing. High heterogeneity also means that the declining incremental adopters is primarily a shakeout of heterogeneity and not because the game is becoming less popular or not doing as well. If marketing of the game was static, it may appear that the return on marketing dollars is declining, but this is just heterogeneity.

Third, the pre-sale parameter is 2.5%. It is useful for a manager to know what % of the population (and/or how many people) purchased the product prior to launch. This metric could be compared to other launches and provides significant insight for planning and forecasting future launches.

Lastly, an interesting insight is to see the effect of the expansion pack on the adoption curve.

Figure VII: Effect of the Expansion Pack on Adoption Curve

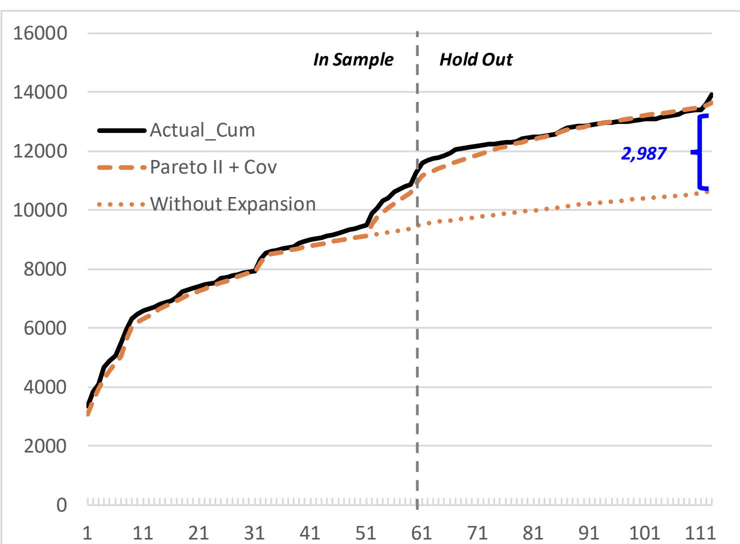

We can see that there were almost 3,000 additional adopters because of the expansion pack. This should be factored into the decision to launch new expansions (in addition to keeping current users playing).