Assessing the Legacy of The Queen's Gambit Netflix Miniseries with Probability Models

Applied Probability Models in Marketing
STAT-476
April 21, 2021

Grant Cho

Executive Summary

The Queen's Gambit is a Netflix miniseries that chronicles a reclusive, female chess prodigy's ascension to the top of the chess world. Since the show's release in late October of 2020, it has already become the most successful and popular chess-centered show in history, receiving nominations for the Golden Globe Awards and, within four weeks, becoming Netflix's most-watched scripted miniseries, with over 62 million households tuning in to watch the show.[1] Chess set inquiries spiked 250% on eBay while the number of new players on chess.com has increased fivefold.[2] That being said, the surge of interest in chess may not be as significant as what the media is reporting.

According to negative-binomial distribution models, experienced players – i.e. players who played at least one game prior to Queen's Gambit's release – are, on average, playing only slightly more after the show's release than before, and their propensity to play has only mildly increased. In contrast, new players who joined after the Queen's Gambit's release exhibited major signs of extreme heterogeneity, with most of the players churning following the creation of their accounts.

In other words, it's very possible that the Queen's Gambit attracted the interests of segments who are not the most inclined to play games like chess, resulting in the vast majority of new players playing a few games before ultimately losing significant interest. This is corroborated with the fact that accounts created after the show came out had a much less of an appetite to play than new players from before the show.

[1] https://deadline.com/2020/11/queens-gambit-62m-viewers-netflix-1234620378/

1 Introduction

1.1 Brief Motivation

I watched the Queen's Gambit upon its release on Netflix, and I absolutely loved the show. It ended my yearlong hiatus from the game of chess and also got me wondering if others felt the same.

I knew from the very start that, if successful, The Queen's Gambit was going to cause an influx in the number of players playing chess. However, I feel like that alone doesn't go far enough. Instead, the focus should be about understanding whether such a spike is there to last.

The main questions I thought of were: what are the long-term consequences of the show's release? Will I finally start seeing more people playing more chess in the long-term? How impassioning was The Queen's Gambit on its audience?

The game of chess is easy enough to pick up but tough – and even discouraging – to master, let alone learn independently. So, going into this research, I hypothesized that the surge of interest would be there but that it wouldn't last. This paper serves to be the answer to that hypothesis.

1.2 Data Overview and Methodologies

It should first be noted that, for a significant portion of this study, I used Canada as the country of interest as it was one of the regions that Netflix released the show to and because chess.com's API would not allow scraping of its American players. Virtually all of the models make extensive use of the negative-binomial distribution.

[2] https://www.cnn.com/2020/12/06/us/queens-gambit-chess-popularity-trnd/index.html

I obtained all my data from chess.com's API, randomly sampling 5000 new and experienced players. I defined new players as those who joined chess.com after October 23rd, 2020, which was the show's release date, and experienced players as those who had accounts prior to that date. The *t* for new players was set as the difference between the date of the scrape, February 22nd, 2021, and the show's release date, which winded up at 122 days. For the sake of consistency, I also classified any player who joined within 122 days before The Queen's Gambit's release date as a "new player" before the show's release.

For each player, I scraped the number of games played prior to The Queen's Gambit's release, the number of games after the show's premier, their join date, their ELO (skill) rating, and their username. I used the date the users joined to derive the *t* values before and after the show was released.

I then split the data into several count datasets as the foundations of the negative-binomial distribution models: two focused on experienced players who played before and after the show's release, while the other two focused on new players.

All of the models' objectives were to observe the change in the inclination to play before and after the show was released for new and experienced players alike.

A glimpse at a small, equally and randomly sampled selection of the data can provide some confidence that the scrape was done correctly. We would expect the majority of new players to not be stellar at chess and to, thus, have lower ELO ratings than experienced players, which is what the graph below confirms.

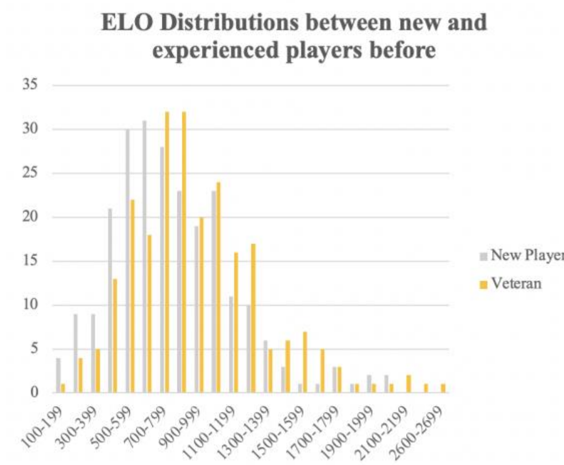

Figure 1: *ELO distributions among scraped players*

There are some outliers in the "New Player" data, which is expected since experienced players may want to create new accounts for the sake of crushing new players. Regardless, given the sparse proportion of outliers and the size of the overall data, there are not enough of these players to distort the overall interpretation of the models' results.

2 Model Results

2.1 Observations on New Players

	Before the show's release	**After the show's release**
r	0.14	0.10
α	0.09	0.13
Average games played	1.54	0.77
Median E(λ)	0.046	0.020
Spike at 0	0	0
Days until reach at 75%	1548	114,781

Table 1: *New Player NBD results*

After fitting two NBD models on a "New Player" population of 2828 players split into two groups – one before and one after the show's release – we can see that new players were already extremely heterogeneous in terms of how many games they would play before the show came out: people's propensities to play chess vastly differed across the population, meaning that there are a few who liked to play a ton of games and a lot who barely played at all.

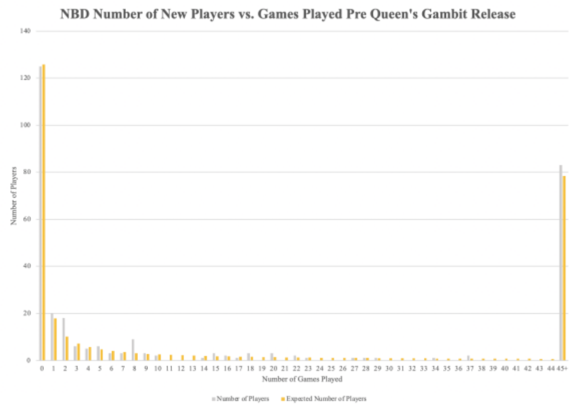

Figure 2: *NBD model on new players for observed vs. expected games played prior to Queen's Gambit's release*

The model above shows that the vast majority of observed chess games were played by a small percentage of people, implying that there were a handful of people with extremely high innate propensities to play the game. On the flip side, most new players would either play a few games before losing interest or create a new account and never play a real game.

Why would people create new accounts and never play? It's possible that there are players who create new accounts with the sole intention of either playing puzzles or to try to learn how to play through opening studies or video tutorials. A majority of people who create accounts at least play a game before calling it quits, but there's still a decent chunk who aren't as inclined to play at all, regardless of possessing an account or not.

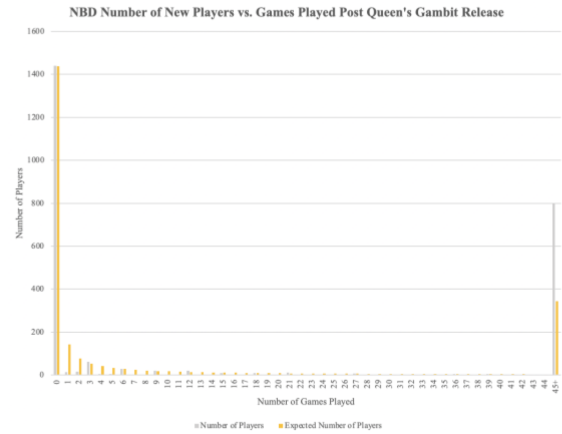

Figure 3: *NBD model on new players for observed vs. expected games played following Queen's Gambit's release*

The NBD model on new players post Queen's Gambit doesn't have a stellar fit to the data, severely underestimating the right-censored counts of games played. Regardless, the data the model was fitted on was not censored and thus the NBD retained the maximum amount of information available (i.e., the inferences derived from this model are not inconsequential).

Heterogeneity among new players increased following The Queen's Gambit's release. There were far more people who joined and played a few to no games than those who joined and got hooked.

This behavior could be a direct result of the show's successful ability to tap into populations that other outlets weren't able to historically reach. Moreover, these populations could be parts of segments that inherently have people who are less inclined to play board games like chess.

To clarify, suppose we take a few people who create accounts on chess.com since they've always enjoyed board games and compare them to people who enjoy watching shows on Netflix and created accounts because they happened to watch The Queens' Gambit. Intuitively, the former segment is more likely to get hooked to the game than the latter since the playstyle is more familiar and thus learning the game is not as serious a barrier to entry. The same can be said for people who previously understood the rules of chess and decided to play because the show sparked their interest.

Players who jump right into the game without knowing anything about chess except through the show, which barely explained the rules, may have a much harder time understanding what's going on and may feel less inclined to play after a game or two as a result.

The NBD models provide further evidence that upholds this sentiment. Indeed, the average games played per new player dropped by 50% after the show's release and the median E(λ) fell from 0.046 to 0.02, further validating our intuition that burnout has become more commonplace.

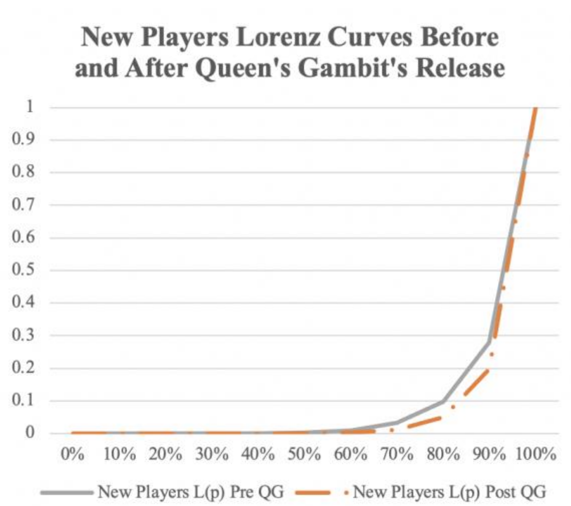

Figure 4: *Lorenz Curves on new players' percentiles on games played*

Furthermore, plotting the Lorenz Curves for new players before and after the show's release reveals how a smaller fraction of new players are playing the majority of games after the show came out. In fact, prior to the show's release, 10% of new players played 72% of total games and 80% of total games after.

Fitting a spike parameter to detect any indication of hardcore-never-players did not significantly improve the model's performance. The negative-binomial distribution fitted to the model already fits at zero well enough to the point where a spike is not necessary, implying that there's an absence of new players who never play. This is reasonable as new players wouldn't just create accounts to end up not using it at all.

Lastly, it's important to note that while the number of days it would take to get 75% of accounts to play a game increased to an extreme amount, it doesn't necessarily mean that overall penetration slowed down. The increase in number of days is more due to the increase in heterogeneity and a greater intake of people who have less propensity to play chess.

2.2 Observations on Experienced Players

	Before the show's release	**After the show's release**
r	0.19	0.20
α	0.13	0.11
Average games played	1.48	1.85
Median E(λ)	0.13	0.043
Spike at 0	0	0.19
Days until reach at 75%	161	106

Table 2: *Experienced Player NBD results*

Fitting NBD models on experienced players paints a completely different story compared to new players. In fact, the analysis reveals that there's only a small change in proclivity after the show's release.

To start, the model after the show's release indicates a very real presence of hardcore-never-players. This discrepancy could happen for various reasons. For one, players who create accounts have had at least a modicum of reasoning for joining in the first place and wouldn't create an account with at least wanting to try to experience the game. Once they've played a few games, they may feel burned out and decide to churn, which the second model subsequently picks up as players who may not play again no matter what.

This is precisely why we didn't observe hardcore-never-players for new players, because we can't use the same new players for before the show's release to after the show's release as then they wouldn't technically be new players anymore, and new players all exhibited at least some desire to play since each of their join dates.

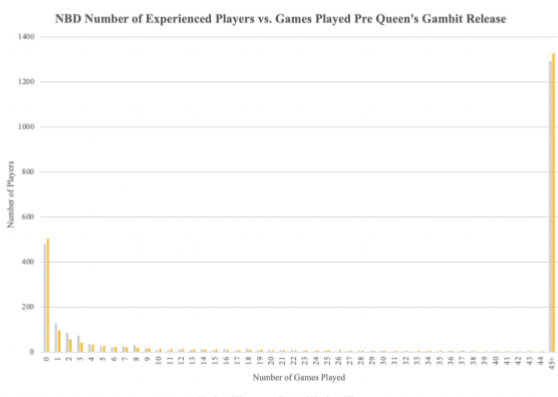

Figure 5: *NBD model on experienced players for observed vs. expected games played prior to Queen's Gambit's release*

There was only a slight uptick in homogeneity among experienced players following the show's release, with the average number of games played by each player increasing only slightly. This type of behavior could be associated with a renewed interest in the game that led to chess players wanting to play more; however, the change is not significant enough to be convincing that such is the case.

Interestingly, the majority expected propensity to play chess decreased after show was released, which could possibly be a byproduct of the hardcore-never-players who may have played a lot and eventually churned after experiencing burnout.

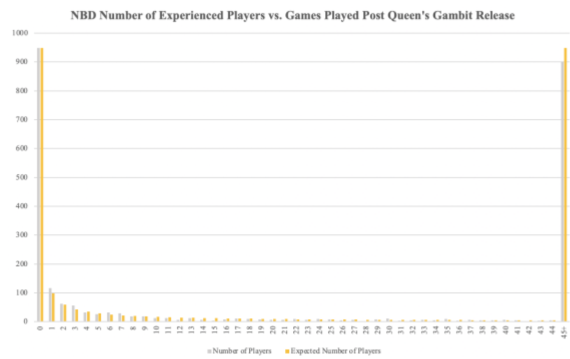

Figure 6: *NBD model on experienced players for observed vs. expected games played following Queen's Gambit's release*

We can see that after the show's release, the number of players who continued to not play actually increased. There are a few explanations for this phenomenon.

As mentioned before, it's very possible that there are a few players who already experienced burnout from chess before the show was released. In other words, some players played hundreds to thousands of games before either retiring from the game completely. Other potential explanations for the possible rise in zero games played include people forgetting their user info – creating new accounts in the process – or switching to a different chess platform, such as lichess.org, which has also seen a surge in popularity in recent months.

Regardless, these all tie in with a single major reason for the uptick in zero games played, which is "time". The median account lifetime before the Queen's Gambit was 690 days, which is close to six times the number of days following the show's release up to when the data was collected in mid-February. 690 days give more time for players to play at least one game, while 122 days provide more legroom for people to not have had the opportunity to log on and play. This uptick did not occur with new players since both models already had the same time of 122 days baked in.

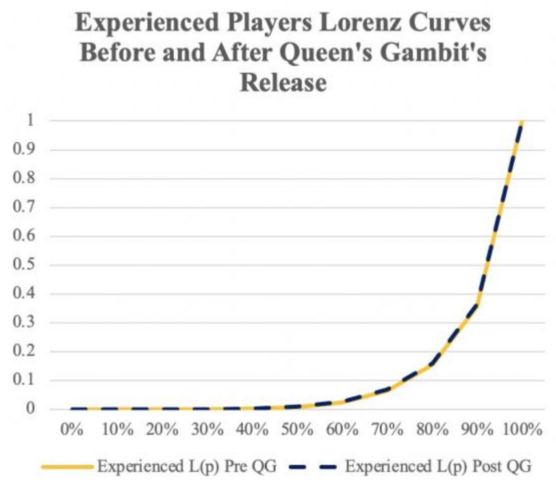

Figure 7: *Lorenz Curves on experienced players' percentiles on games played*

Moving forward with the intuition, plotting two Lorenz Curves on the model shows that prior to the show's release, 10% of experienced players alone played 63% of total games and 64% of total games afterwards. Not much has changed in terms of how many games existing players play, implying that people who played chess before did not feel more motivated to play as a result of the show's release.

Because of this greater increase in homogeneity seen across the experienced player population, the number of days it would take to have 75% of these players play at least one game decreased by over 50 days, which is relatively unremarkable considering the change in *r* values.

3 The Mainstream's Role in Chess

According to Weibull distribution models, people who joined chess.com prior to the Queen's Gambit's release exhibited extremely homogeneous behaviors and were duration independent, while those who joined after showed signs of both positive and negative duration dependence. This sudden shift in duration dependence amongst all observable segments is evidence of a possible network effect that can most likely be attributable to the show's mainstream popularity.

From a purely observational standpoint, the iOS and Android chess.com downloads follow distinct distributions. Indeed, upon further testing, it was found that the groups do have differing adoption rates: the former was extremely quick to peak in downloading the chess.com app soon after the show's release while the latter peaked almost 8 weeks later.

Downloads from iOS users were heavily influenced by the Queen's Gambit while downloads from Android users were less so, implying a two to three step process bucket in the chess ecosystem.

The first bucket has Queen's Gambit viewers downloading the app straight after watching while the second bucket has viewers requiring further mainstream convincing after hearing about the show before ultimately downloading the chess.com app.

This next part of the paper intends to reconcile this differing behavior between iOS and Android users by emphasizing the importance of mainstream network effects as well as better understand how this all ties back to the future of the promotion of chess.

3.1 Covariate Selection

Throughout the course of the study, I narrowed my covariate selection down to three primary, causal variables for the sake of parsimony: viewership count on The Queen's Gambit, the presence of holidays, and viewership of Grandmaster Hikaru Nakamura's Twitch stream.

The type of causal relationship between these variables and chess.com app downloads is important to note as they are the reasons someone may want to decide to get the app and not the means to doing so. This is why I did not include covariates such as Google Trends, as I believed such variables could be helpful in determining the number of downloads but not as helpful in the grand scheme of understanding *why* someone would download the chess.com app.

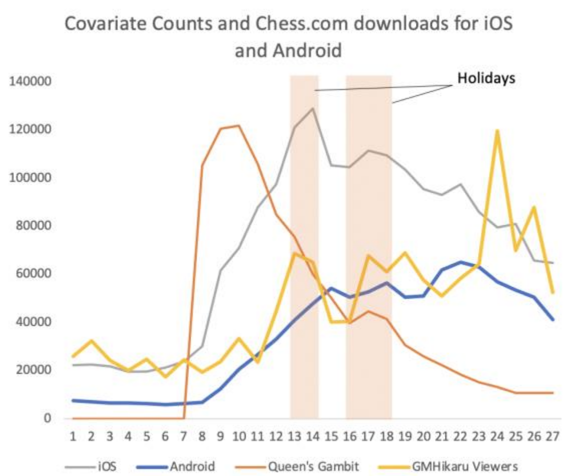

Figure 8: *Scaled covariate counts vs. Chess.com downloads among iOS and Android users*

The Queen's Gambit

For context, The Queen's Gambit chronicles a reclusive, American, female chess prodigy's ascension to the top of the chess world in the 1960's, torching Russian players along the way. Netflix's ability to reach broader audiences that the game on its own would never have been able to achieve opened the doors for chess to become mainstream. This should come as no surprise, as since the show's release in late October of 2020, it has already become the most successful and popular chess-centered show in history, receiving nominations for the Golden Globe Awards and, within four weeks, becoming Netflix's most viewed scripted miniseries, with over 62 million households tuning in to watch the show.[3] Chess set inquiries spiked 250% on eBay while the number of new players on chess.com has increased fivefold.[4]

Moreover, a spike in viewership in this show corresponded strongly with an uptick in downloads for the chess.com app for both iOS and Android users, so including this variable within at least one of the phone user segment's models was almost an inevitability.

Twitch

Hikaru Nakamura is an American chess grandmaster who was, at one point, ranked second in the world. His chess skills, detailed insights into the game, and entertaining content has made him by far the most popular individual chess streamer on Twitch, with thousands tuning into each of his livestreams. His channel represents the best chance at capturing people's growing interest in chess before, during, and after the Queen's Gambit.

It's worth noting that including the Queen's Gambit and Twitch covariates into the model simultaneously would not be very practical as that runs into the issue with multicollinearity, where the correlation between the two independent variables was $r = -0.71$. This can be interpreted as a decrease in viewership for Queen's Gambit comes with an increase in viewership for Hikaru's Twitch stream. Because of this inference, I will use the Twitch as a type of "lagged" variable,

[3] https://deadline.com/2020/11/queens-gambit-62m-viewers-netflix-1234620378/

[4] https://www.cnn.com/2020/12/06/us/queens-gambit-chess-popularity-trnd/index.html

wherein people who watch The Queen's Gambit may go onto Twitch to watch more chess before ultimately deciding to download the app.

Throughout this paper, I will often be referring to his channel's covariate as "Twitch". I was able to obtain this data by scraping Hikaru's twitchtracker.com page.[5]

Holidays

Holidays are often associated with an influx in the number of players who play games. Both the Queen's Gambit and Hikaru's stream may capture some of the effects of holidays on their own; however, there may exist segments who never watch Queen's Gambit or Twitch but just want to play with the spare time they have. Because of the potential uncaptured spike in these particular users, using "Holidays" as a covariate makes sense for additional fitting. Particular holidays commonly practiced in the United States and captured by the covariate include Thanksgiving, Christmas, and New Year's Eve.

3.2 Model Building

From an observational standpoint, we can see that the left-truncated nature of the dataset gives way for interpretation that, *prior* to the Queen's Gambit's release, the people downloading the chess.com app are homogeneous and duration independent.

This is because chess.com went online back in May 2007 and released their app in 2011, so people who were deeply interested in chess and/or loyal to the website would've downloaded the app closer to the release date and the imitators followed up within the next 10 years. Over time, homogeneity would've surely continued to grow, which is likely the reason why we can observe a flat line in incremental chess.com app downloads for both Android and iOS users.

After the Queen's Gambit's release, however, there is definitely heterogeneity within downloaders as the show was able to reach segments that are heterogeneous in nature, reshaping the innovator and imitator coefficients of new downloaders from a Bass Model viewpoint.

As such, we can intuitively conclude that an exponential distribution model is not ideal, no matter how many segments are added, because we need to observe a shake-up in duration dependence driven by a network effect. For further convincing, we can see how the absence of a c value results in a terrible model below, even with the Queen's Gambit covariate included.

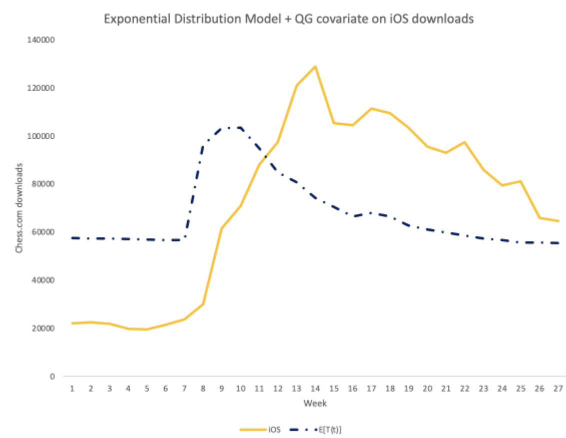

Figure 9: *Fitting an exponential distribution with covariates model on chess.com downloads results in a spectacularly bad fit*

The necessary inclusion of duration dependence narrowed the decision down to whether to use the Weibull or Burr XII model. The defining difference between the two aforementioned models is that the Weibull only requires a λ and c parameter, whereas the Burr XII requires an r, α, and c parameter.

[5] https://twitchtracker.com/gmhikaru/statistics

Because the dataset only includes 27 observations, using the distribution with the fewest parameters would not only prevent overfitting but also result in a more parsimonious final model.

In using the Weibull distribution, we can make the assumption that there is homogeneity within segments, which isn't an unguided intuition considering that the app has been out for over a decade and the λ values have likely been established by this point. In other words, the inherent propensity should not change at this point: you either want to download the chess.com app or not. The people who have already downloaded 10 years ago are not going to download again since they already obviously have the app, so that level of heterogeneity is gone. The people downloading within the 27-week period will likely have similar propensities to get the app, yet whether they choose to download depends on the effects of the covariates and the time that passes.

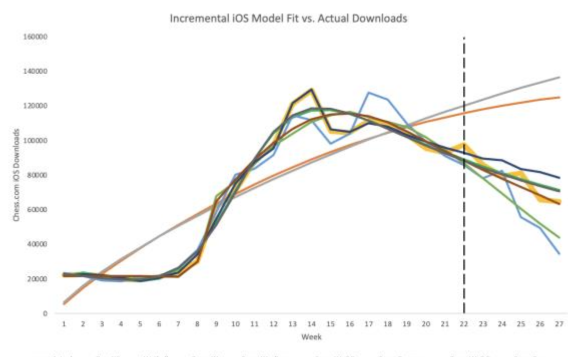

Figure 10: *All model fits on incremental weekly iOS downloads*

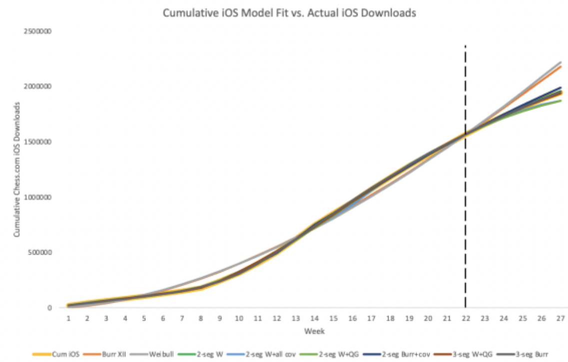

Figure 11: *All model fits on cumulative weekly iOS downloads*

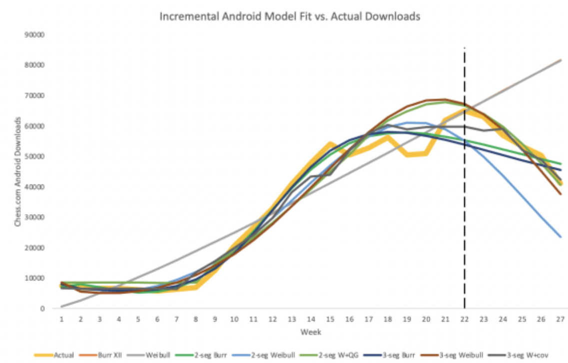

Figure 12: *All model fits on incremental weekly Android downloads*

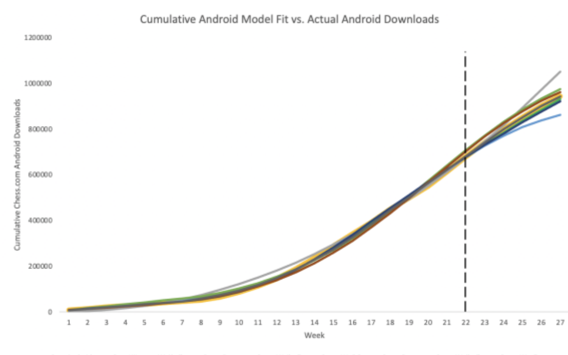

Figure 13: *All model fits on cumulative weekly Android downloads*

We can see in the figures above that the intuition is quite well-grounded. Visually, we can see that Burr XII models *with* covariates tended to overfit and did not perform as well on the sparse testing data as the Weibull distributions.

Also noteworthy is that having just one segment oversimplifies the differences in duration dependence experienced in potentially alternative segments.

For instance, some people may have watched and disliked the Queen's Gambit, resulting in a possible decreasing hazard rate. Those who liked the Queen's Gambit could have a positive duration dependence and may want to download the chess.com as more time passes. There may also exist hardcore-never-tryers who may never want to download the app. These are all scenarios that cannot be captured by just one segment. In fact, I found that the ideal number of segments based on what was mentioned above is three, which is what I used and will be presenting the findings for in a few sections below.

3.3 Parameters

Pre-Queen's Gambit c value

The Queen's Gambit was released on Netflix on October 23rd, which is represented as week 8 in the dataset. The assumption is that people most likely wouldn't have finished watching until a week or two after the show's release as that's when a spike was present in Google Trend mentions for the show. This may be attributed to the fact that there's leftover interest after watching the show (i.e., some people may want to ensure they fully understand the plot or people want to watch highlights of the show, etc.).

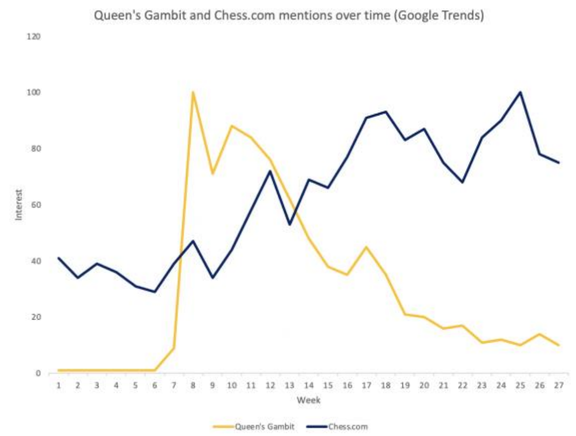

Figure 14: *Google Trend mentions of the Queen's Gambit and Chess.com*

Moreover, as mentioned previously, downloaders exhibited homogeneous behaviors prior to the show's release. I therefore held the prior *c* value constant at 1 until week 10. In doing so, I would capture the lag inherent in people finishing watching the Queen's Gambit as well as the homogeneity in downloaders prior to the show's release. This spared the need for an additional parameter, resulting in a model that's both more parsimonious and easily interpretable.

Android and iOS Covariates

Figure 1 depicts the trends between the covariates selected and the two types of users, to which there are several main takeaways. The first is that Android users were much slower to adopt than iOS users; the second is that the Queen's Gambit peaked much earlier than Hikaru's Twitch viewership; the third is that holidays coincided with local maxima for both Android and iOS downloads.

Combining these observations along with the fact that there exists multicollinearity between the Queen's Gambit and Hikaru's Twitch stream, we can deduce that 1) we don't put these covariates into the same model 2) it's more reasonable to include the covariate that peaks later into the Android model 3) only include additional covariates when absolutely necessary, and 4) the holiday covariate will likely be much more relevant to include within the Android model as the effect may be less strong than the Queen's Gambit covariate for iOS users.

3.4 Training and Testing

I selected a holdout period of five weeks as that would represent a training period of 22 weeks, which roughly equates to 80% of the provided data. There were, of course, some initial obstacles present in using the 5-week holdout period.

For example, the Android model proved much more difficult to fit on the training data as the holdout period of five weeks coincided with a peak in chess.com downloads. Because of this, adding additional covariates would often result in an overfitting of the model as shown below

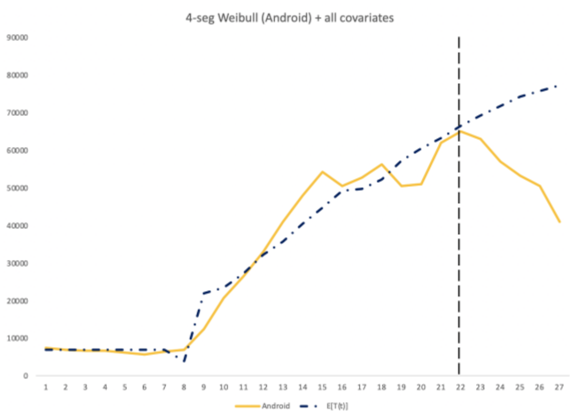

Figure 15: *Example of an overfitted model as a result of having too many covariates and the holdout period coinciding with a peak in downloads*

Correlation Matrix	Queen's Gambit	Twitch	Holidays
Queen's Gambit	1	-0.705637485	0.028658655
Twitch	-0.705637485	1	0.118347239
Holidays	0.028658655	0.118347239	1

Table 3: *Correlation matrix of all the used covariates*

This issue was resolved by using uncorrelated variables mentioned back in section 1.1, which, in this case, would be Hikaru's Twitch stream viewership and Holidays.

3.5 Model Selection

Summary of iOS Model Fit

Model	LL	#params	BIC	Testing MAPE	Testing MdAPE	Overall MAPE	Overall MdAPE
Burr XII	-12647377.58	3	25294805.59	64.02	51.21	40.44	30.14
Weibull	-12666279.69	2	25332593.01	75.69	60.80	43.09	28.34
2 seg Burr XII	-12524430.09	7	25048977.86	6.22	4.04	5.78	5.17
2 seg Weibull	-12603529.52	5	25207143.11	37.23	28.67	32.72	22.52
2 seg Burr XII+cov	-12524686.9	9	25049525.10	8.61	5.71	5.44	5.31
2 seg Weibull+cov	-12562651.03	7	25125402.93	33.82	8.41	13.98	8.41
3 seg Weibull+cov	-12524167.6	11	25048520.13	4.24	3.50	5.17	3.76

Table 4: *Summary of iOS models*

Summary of Android Model Fit

Model	LL	#params	BIC	Testing MAPE	Testing MdAPE	Overall MAPE	Overall MdAPE
Burr XII	-6740906.394	3	13481846.41	45.54	21.48	42.51	21.48
Weibull	-6740830.746	2	13481695.11	16.64	53.24	50.02	53.24
2 seg Burr XII	-6688812.884	7	13377743.45	32.15	34.17	55.57	64.58
2 seg Weibull	-6692953.624	5	13385991.3	31.54	29.84	20.29	18.47
2 seg Burr XII+cov	-6688219.98	9	13376591.26	30.59	33.13	55.33	65.17
2 seg Weibull+cov	-6701042.73	7	13402203.14	3.32	2.74	19.14	12.83
3 seg Weibull+cov	-6689767.987	11	13379771.33	3.85	3.14	10.15	7.13

Table 5: *Summary of Android models*

Comparing the iOS models, it's clear that the 3-segment Weibull model with the Queen's Gambit covariate performed the best: it had the lowest Bayesian Information Criterion; the lowest mean absolute percentage testing error (MAPE); and the lowest median absolute percentage testing error (MdAPE). This means that the final model was able to increase the log-likelihood to the point where the tradeoff of adding parameters was validated.

The 2-segment Burr XII performed extremely well given that it included no covariates whatsoever. There are a few reasons why this model nor a 3-segment Burr XII model was used as the final model.

Firstly, the 3-segment model wasn't used because the third segment, π_3, was only of size of 0.2 downloads.

Secondly, regarding the 2-segment Burr model, the 3-segment Weibull with covariates model performed much better on the testing data than the Burr.

Finally, from a managerial standpoint, it would make more sense to be able to interpret more direct causes of people's desires to play chess, rather than purely infer from segments and not exactly answer the question of why people are starting to play.

In the case of the Android models, the 2-segment Weibull with covariate model's testing MAPE and MdAPE are both lower than its 3-segment counterpart. However, the BIC of the former is much lower, implying that the tradeoff between parameters for a higher log-likelihood was warranted. Additionally, the overall MAPE and MdAPE for the 3-segment Weibull with covariates is much lower than any of the other models.

Taking everything into consideration, the best model to use for iOS downloads would be a 3-segment Weibull distribution with the Queen's Gambit as the covariate. The most appropriate model to fit against Android downloads would also be a 3-segment Weibull but with Twitch and Holidays as the covariates.

4 Inferences & Discussion

4.1 Final iOS Model Results

Final iOS Model Results & Parameters	Value
λ_1	0.0012
λ_2	1E-06
λ_3	0.00021
$c_{prior\ QG}$	1
c_1	1E-05
c_2	4.24
c_3	3.077
π_1	0.90
π_2	0.05
π_3	0.053
β_{QG1}	0.032
β_{QG2}	0.0038
β_{QG3}	0.0038
Overall Log Likelihood	-12,524,168
Overall BIC	25,048,520
Testing MdAPE	3.50
Testing MAPE	4.24
Overall MdAPE	3.76
Overall MAPE	5.17

Table 6: *Final results of the 3-segment Weibull with Queen's Gambit covariate*

Fitting the 3-segment Weibull while only using the Queen's Gambit as a covariate yields an excellent fit onto the data.

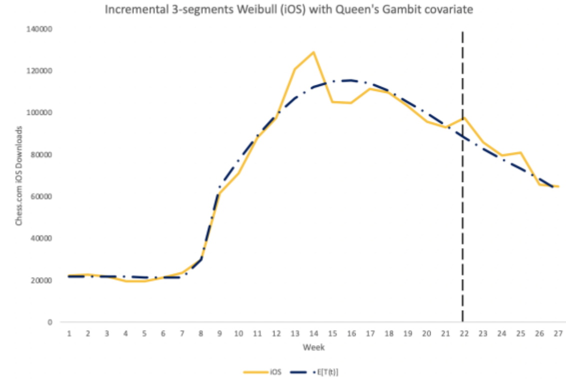

Figure 16: *Final model fit on incremental iOS downloads*

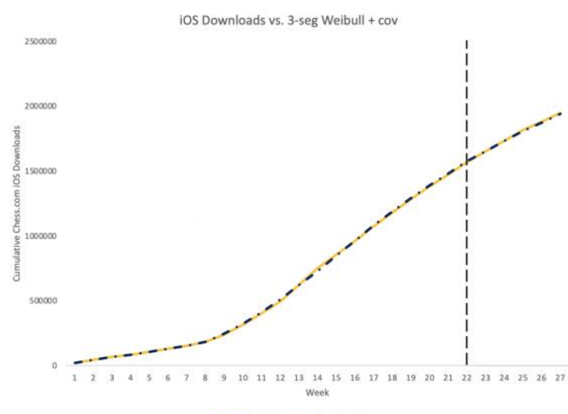

Figure 17: *Final model fit on cumulative iOS downloads*

The model, fitted on a training data of 22 weeks, was able to obtain a testing MdAPE testing MAPE, overall MdAPE, and overall

MAPE of 3.5%, 4.24%, 3.76%, and 5.17% respectively. Fitting the model onto the entire dataset yielded a model whose parameters were unchanged, indicating an extremely robust model.

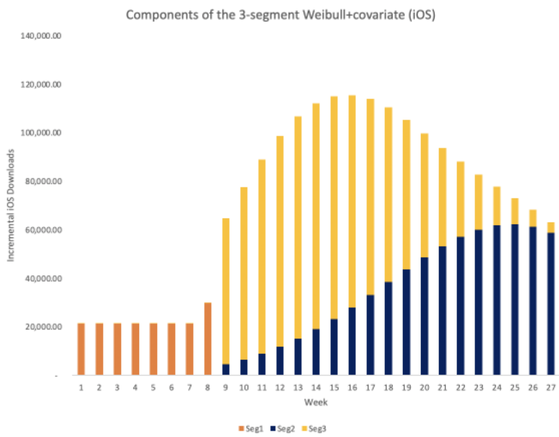

Figure 18: *Stacked bar chart on incremental iOS downloads*

We can immediately draw a few key takeaways from the derived parameters. To start, one commonality amongst all segments is the shared passion for the Queen's Gambit. The positive β values represent a "stretch" in time for however many people watch the Queen's Gambit. In other words, an increase in the number of people who watch the Queen's Gambit will result in people in a segment with a positive duration dependence – segments 2 and 3 – having a much stronger tendencies to download the app, which is strongly evident of a network effect. Those with a negative duration but with a positive β may feel a slight nudge to download the app when more people watch the show, but the probability they eventually download remains otherwise unchanged over time.

Next, iOS users' propensity to download is not particularly high. In fact, segment 1 has the highest innate propensity of any other segment to download the app at just $\lambda_1 = 0.0012$.

That being said, segment 1 also has the lowest duration dependence of the three segments, with a c_1 value of essentially 0. This makes the probability of someone from segment 1 downloading the app at time *t* given that they've survived up until that point also basically 0. This result should only be interpreted as the *c* value dropping to 0 as a result of a shakeup in the heterogeneity of the population.

Segment 1 users, in reality, are strictly duration independent. As revealed in Figure 11, we can see that segment 1 is the sole makeup of app downloads prior to the show's release, with a very flat incremental download trend.

This confirms the intuition that people who wanted to play on their phones would've downloaded the app a long time ago but given the left-truncated nature of the dataset, we're only left with homogeneity. Chances are, if you didn't download when the app was released, you're unlikely to download now. If you do decide to download, it's because of your innate propensity to play rather than you succumbing to your desires. This was the behavior of millions of iOS users prior to the Queen's Gambit's release.

Segment 1 completely disappears following the show's release. Again, this is because the very existence of segment 1 hinged upon the prior *c* value of 1 since the post-QG *c* dropped to 0.

iOS users in segment 2 displayed the highest levels of duration dependence but extremely low inherent propensities to download. This particular segment would not have otherwise downloaded the app had it not been for the growing urge to play and is most likely the quintessential example of a segment that downloaded because of directly watching the show.

To demonstrate, a person in segment 2 would probably never think to play chess until he/she hears about the Queen's Gambit on the news. They then either watch the show or see other posts online that tempt them to eventually try the game out.

People in segment 3 also have an extremely high, positive duration dependence but differ from segment 2 in that they are much more innately inclined to play the game of chess given their comparatively larger λ value. These are the type of people who are likely very impressionable and willing to try out new things once they come out (i.e., innovators).

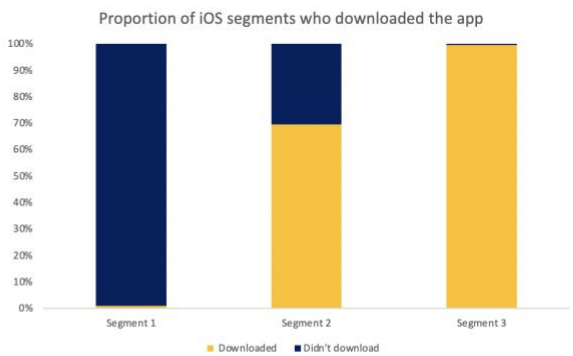

Figure 19: *Proportion of iOS users from each segment who did and did not download the chess.com app*

We can see how segment 3 download the app en masse right when the show comes out; however, as time goes by, segment 2 slowly becomes the dominant downloader of the chess.com app as the vast majority of segment 3 already downloaded the app. Segment 2 essentially becomes the new segment 1 as innovators begin to disappear and imitators start to take over, albeit with more heterogeneity than segment 1.

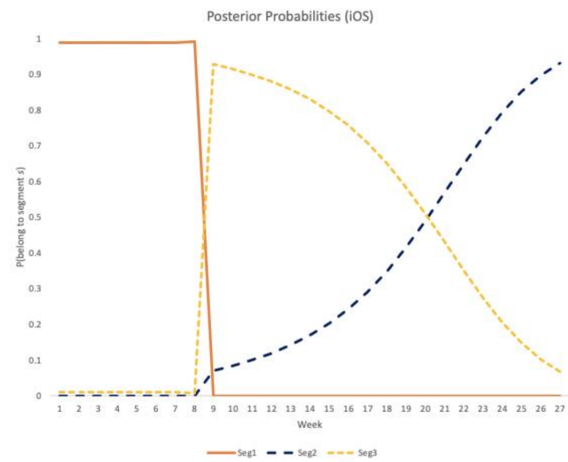

Figure 20: *Posterior probabilities of segment classes for iOS users*

Indeed, from the above graph, we can see how people who download later are most likely to be classified within segment 2 since people from segment 3 have already downloaded the app. Prior to the Queen's Gambit's release, people were basically guaranteed to be classified as part of the segment 1 population given the level of homogeneity at the time.

Finally, from a managerial standpoint, we can classify segment 1 as the disappointing segment. This classification may seem a bit harsh but considering the fact that sites such as chess24.com and chess.com invested so much resource into popularizing chess only to get the same kinds of people downloading is nothing but disappointing. People in segment 2 could be classified as new imitators because people in it likely succumbed to the network effect; and segment 3 as innovators for watching the show early, downloading the app quickly, and imposing the network effect onto segment 2.

4.2 Final Android Model Results

Final Android Model Results & Parameters	Value
λ_1	1E-08
λ_2	3.64E-05
λ_3	0.0063
$c_{prior\ QG}$	1
c_1	0.005
c_2	3.25
c_3	1E-05
π_1	0.89
π_2	0.055
π_3	0.055
$\beta_{Hikaru\ 1}$	0.0041
$\beta_{Hikaru\ 2}$	0.011
$\beta_{Hikaru\ 3}$	-0.012
$\beta_{Holidays\ 1}$	-0.1
$\beta_{Holidays\ 2}$	0.073
$\beta_{Holidays\ 3}$	-0.1
Overall Log Likelihood	-6,689,768
Overall BIC	13,379,771
Testing MdAPE	3.14
Testing MAPE	3.85
Overall MdAPE	7.13
Overall MAPE	10.15

Table 7: *Final results of the 3-segment Weibull with Twitch and Holidays covariates*

Just like in section 2.1, fitting a 3-segment Weibull model with Twitch and Holidays covariates onto the Android data also results in an excellent fit.

Recall that back in section 1.4, uncorrelated variables had to be used in order to prevent overfitting of the data, which is why I decided to opt out of using the Queen's Gambit as a covariate again since it didn't have the kind of lag seen in the Android data. Had more correlated variables been fitted to the data, the model would've overshot its forecast in an effort to fit better onto the training data, sacrificing overall fit for a better testing error.

As a result, the model fit extremely well onto the testing data, with testing MAPE and MdAPE of 3.85% and 3.14% respectively. The model fit only slightly less well on the overall data, with the overall MAPE and MdAPE of 10.15% and 7.13% respectively.

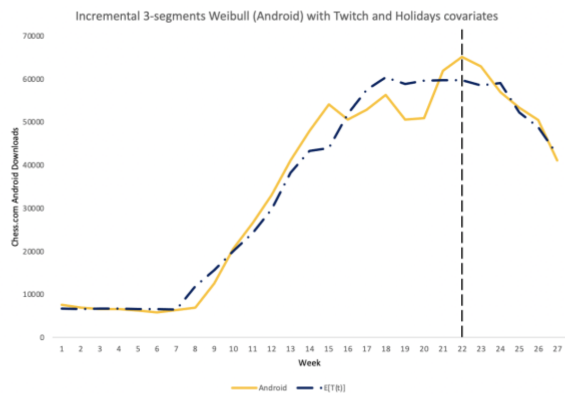

Figure 21: *Final model fit on incremental Android downloads*

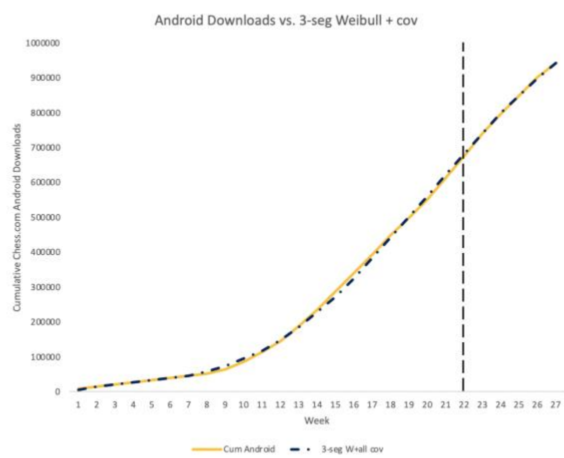

Figure 22: *Final model fit on cumulative Android downloads*

Fitting the model onto the entire dataset yields only a slightly lower overall MdAPE and testing MdAPE, which is expected. More importantly, however, is that the parameters barely changed, indicating a robust model.

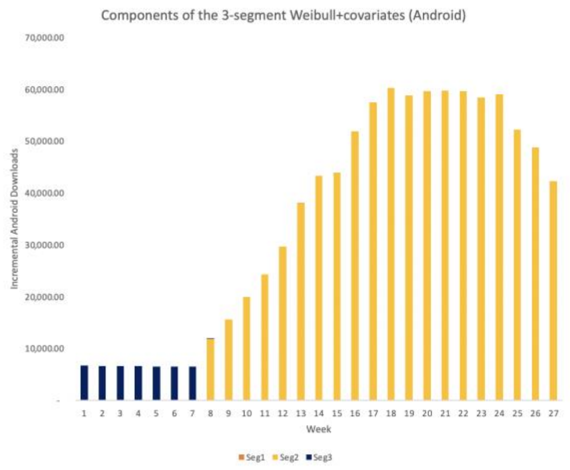

Figure 23: *Stacked bar chart on incremental Android downloads*

This model is slightly different from the iOS model in that there exists a presence of hardcore-never-tryers, as seen in the bar chart above. The extremely low $\boldsymbol{\lambda_1}$ of 1E-08 and *c* value of 0.005 strongly suggests that people in segment 1 simply have no interest at all in trying the game of chess, let alone the desire to download the chess.com app.

On the other hand, Android users in segment 3 are very similar to the iOS users of segment 1 in that they have a relatively high, innate propensity to play chess and exhibit strong signs of homogeneity as evident by the prior *c* value supporting the existence of segment 3 downloaders. In other words, people in segment 3 are truly duration independent and are homogeneous in nature. These are the kinds of people who are exclusively imitators.

People in segment 2 dominated the number of chess.com app downloads for Androids post-Queen's Gambit, and it's no wonder: these are the people who are the most impressionable and curious of the three observed segments.

While people in segment 2 had almost no initial interest in playing chess prior to the Queen's Gambit's release given their extremely low $\boldsymbol{\lambda_2}$, curiosity driven by the network effect motivated them to research more into the game before ultimately deciding to download the chess.com app.

To elaborate, people in segment 2 are more easily susceptible to the network effect given their positive $\beta_{Hikaru\,2}$ values. What might happen is that these people may have watched the Queen's Gambit or seen content show and then decided to look for alternative sources of chess, such as streams, before deciding whether or not to download the app. If they see a lot of viewers on Hikaru's channel, then they may feel more inclined/pressured to get in on the trend since lots of viewers is typically associated with credibility (i.e., *if other people are watching and know what's going on, then why don't I know?).*

Additionally, segment 2 people feel more inclined to download the app whenever there is a holiday, likely because they have would have more time during these periods to try new activities, such as chess.

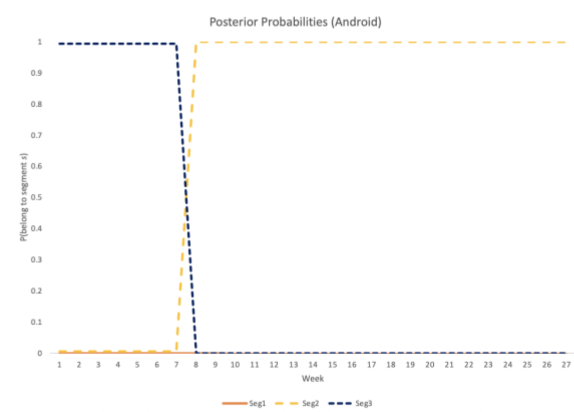

Figure 24: *Posterior probabilities of segment classes for Android users*

As expected, the posterior probabilities show that merely downloading the app would place a person out of segment 1 since that segment represents the hardcore-never-tryers. Again, prior to the Queen's Gambit's release, people downloading the chess.com app were essentially homogeneous in nature given how long the app has been around, so any downloads back then would place that person in segment 1.

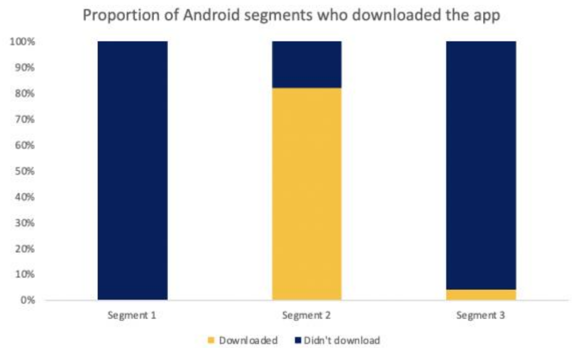

Figure 25: *Proportion of Android users from each segment who did and did not download the chess.com app*

Finally, we can see that, not only did people from segment 2 cumulatively download the app the most, but they were also more likely to download the app than any other segments.

In fact, upon viewing this data, we can observe just how powerful the network effect is on these types of people despite their initial, inherently low propensity to play chess. There are people who don't even know they want to play chess in both iOS and Android users, and it took something mainstream to make them realize.

5 Conclusion and Closing Remarks

From the NBD models, we're able to derive several inferences. The first is that new players who joined after the show's release experienced burnout and churned from the game of chess faster than historical estimates. The second is that experienced players, excluding the hardcore-never-players, are still playing just as much chess as they were before the show was released, if not a little bit more.

There are a few constants we can observe from the chess.com app download data. The first is that there will always exist people who will never try chess, no matter what you do to promote the product. The second is that, over time, homogeneity will be an inevitability as innovators fade and imitators' heterogeneity wane. The third and most important, however, is that there does indeed exist people who don't know they want to play chess until they're nudged to try it by the mainstream.
In the case of iOS users, the Queen's Gambit was mainstream enough to convince the main innovators to try the game out, compounding the network effect.

As for Android users, the Queen's Gambit may not have been enough for most, so they sought alternative mainstream sources, like Twitch, for further convincing. When given the time, like in holidays, they'll feel more inclined to give the app a try.

From a managerial standpoint, it's reasonable to assume that there will always exist the kinds of people who are unaware of their current propensities. Now is their best chance to keep captivating new audiences into the game of chess before homogeneity falls back into place.

If such homogeneity does return, trying old tactics won't do much to bring about new viewers. Instead, new and innovative developments, like the Queen's Gambit or something colossal like the 1972 World Chess Championship will certainly drive interest. The idea doesn't necessarily have to be mainstream per se; the platform does. If the new idea lacks a mainstream platform, then reaching the people who are unaware of their desire to play becomes much more of an impossibility.