David Liedell (djl296), David Castillo (dc689)
Albert Heidmann (ath55), Glenn Mathews (gjm98)

Yik Yak and Information Cascades: A Study of Viral Yaks
May 5, 2015

A. Abstract

Yik Yak allows its users to anonymously speak their minds on a public forum. The platform implements a scoring system, giving each participant an incentive to post worthwhile content. The score of a post is indicative of its quality due to the fact that other users vote on it. However, the content of a post may not be the only factor a user considers when they determine the quality of the post. In this study, we observe whether the score of a post can influence a user's vote. Using the principles of information cascades, we set out to find whether heavily influencing the score positively or negatively on various posts will cause other users to vote in a similar manner. We found that even a small artificial inflation of votes in either direction can drastically change how users will vote on a post.

B. Introduction

Yik Yak is a anonymous mobile social media application that allows users to make text posts of up to 200 ASCII characters without an account. This post is then added to a feed of roughly 200 recent posts, sorted chronologically or by decreasing post score. Users have the ability to 'upvote' or 'downvote' the post, with an 'upvote' equating to +1 and a 'downvote' equating to -1 for the post's score. The score of the posts is displayed next to the post and a user's overall score, or 'yakarma', is the total sum of their posts' scores. The score on a post can range from -5 to infinity. When the score reaches -5, it is immediately deleted from the feed so that no one else can vote on it. The posts that are available on the feed are based on your location and you can see all posts made within 5 miles. Users gain a sense of satisfaction from having the highest user score among their friends, so there is an incentive to both create posts, and to vote on other users' posts.

Yik Yak is most commonly used by university students and there are hundreds of posts created for Yik Yak on the Cornell campus each day. We wished to determine if it is possible to influence the votes of anonymous users through the use of information cascades. Specifically, we wanted to test if we can affect how users will vote on a post by having a large group of users 'upvote' or 'downvote' specific posts to pre-determined scores, then tracking the score of those posts over time after we artificially manipulated them.

Imagine the following scenario. Person A is faced with a decision. There is a building in front of them, and he can either enter the building or leave without entering. Person A observes 20 people make this decision before him, and he knows what decision each person made. Person A has a private information signal which influences his decision. He heard that

the building was entertaining, so he has a high signal to enter the building. However, all 20 people before him decided to not enter the building; a low signal. Person A therefore decides to not enter the building after all because the majority signal was the low signal. This is an information cascade, because Person A went against his private signal after witnessing 20 people make the opposite decision. This is the element of Network Theory we are trying to observe in our project. More specifically, we are observing whether or not users will vote differently simply as a result of a post having an inflated score. Will a user upvote a post he would normally downvote just because it has a high score? Will a user downvote a post he would have normally upvoted just because it has a low score? Is the magnitude of the information cascade affected by quality of the post? These are the questions we are trying to answer with this study.

Our model reflects score manipulation on a given Yik Yak post. We input high/low signals into the model for the purpose of observing the effect on the final score of the post. Our design involved feeding high and low signals to new posts that we created. We created posts spanning a variety of topics both popular and unpopular on Yik Yak to observe whether or not information cascades were amplified or dampened by the quality of a post. We tracked and recorded the score of each post over time in order to observe the rate at which the score increased. Initially, we tracked each post without affecting the score, then reposted it a week later and gave it an input of low signals. We then recreated the same content a week later and fed it an input of high signals. By spacing the trials a week apart each time, we ensured that a similar audience would be voting on the posts, making the trials as similar as we could. Therefore, each post was experimented with a control, high, and low signal inputs in similar settings.

In order to collect multiple samples, we gathered a group of roughly 20 individuals who had access to Yik Yak, and used their voting power to test the creation of information cascades. These users were instructed to vote on each posts as soon as it appeared on their feed. We made a few posts ranging various topics and immediately downvoted them so they were at a score of -2. We then tracked the post score as unbiased users voted on them. We simultaneously created a number of posts and immediately upvoted them to between 10 and 20 upvotes, and again kept track of the post scores. The third type of posts we observed were by taking a number of recent posts and observing them for a few minutes to collect initial data points, then quickly altering their scores using a large number of upvotes and downvotes. We then saw if this alteration changed the rate of the cascade, and saw if a popular post lost its popularity and was voted off, or if an unpopular post increased in popularity.

C. Analysis

Our experiment was conducted in the following manner. We came up with several yaks that had a wide range in topic and “quality.” Each of these yaks were posted 3 different times, each time at the same time of day on the same day of the week in an attempt to utilize a similar user base. The first time each yak was posted, they received no artificial votes. This functioned as our control group. The second time each yak was posted, we artificially inflated them with anywhere between 10 - 20 upvotes immediately upon posting. The third time each yak was posted, we deflated them with 2 downvotes each immediately upon posting. For each trial, we recorded the score of the yak at 0 minutes, 1 minute, 5 minutes, 10 minutes, 15 minutes, 20 minutes, 25 minutes, 30 minutes, 45 minutes, 60 minutes, and 120 minutes. After

collecting the data, we analyzed it to see if information cascades were created by comparing the trials with positive and negative artificial vote inflation with the control group.

.

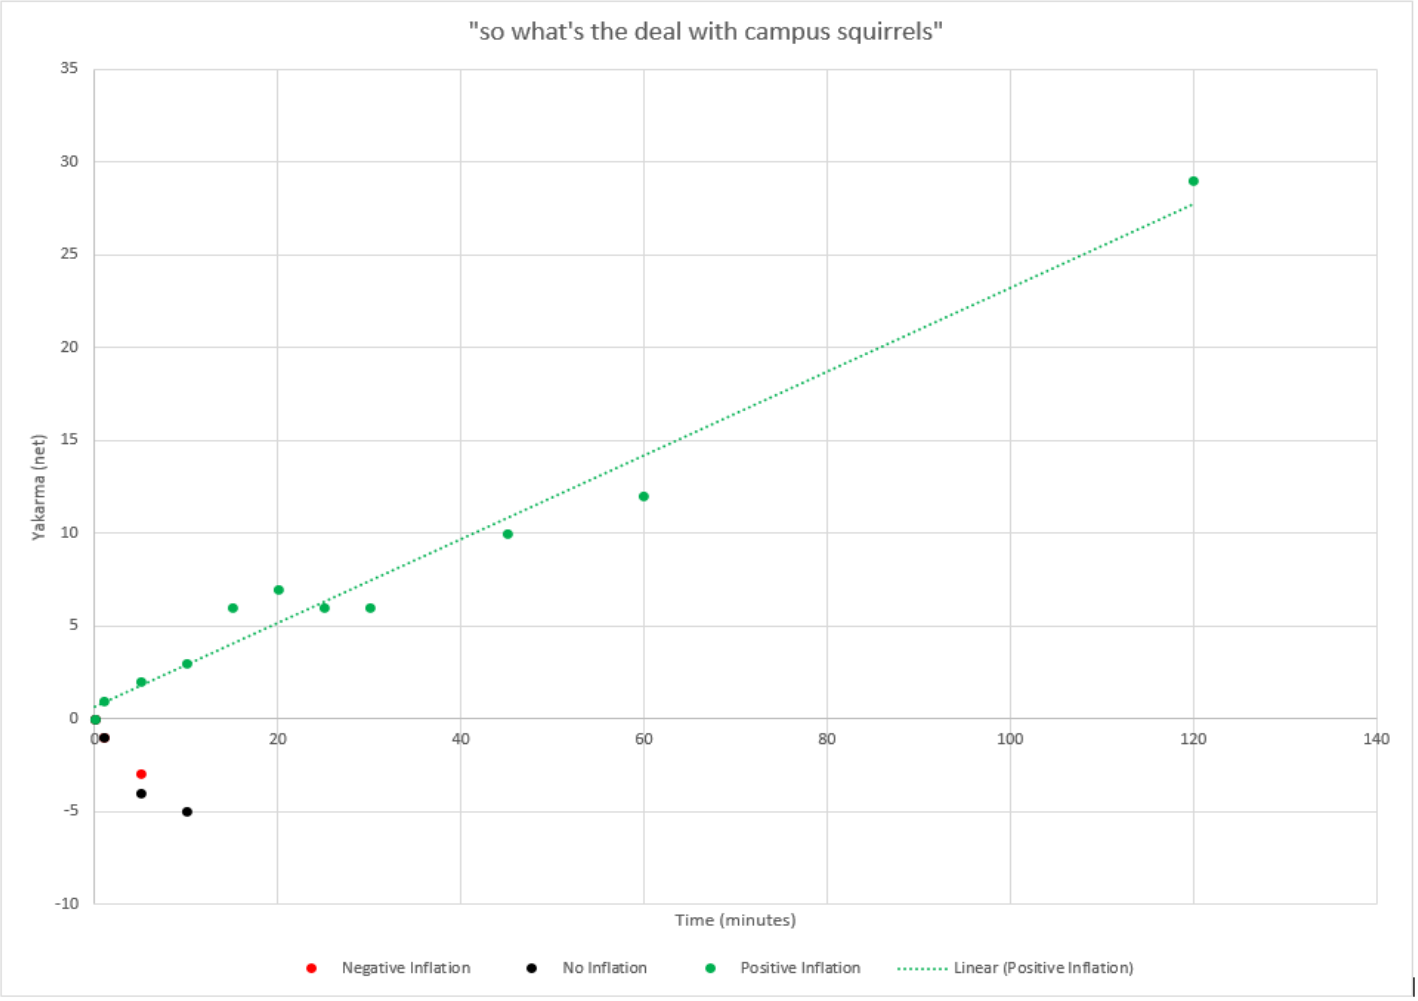

The graph above is an example of a low quality post that we created and tested in our study. Specifically, the post is “so what’s the deal with campus squirrels.” The x-axis represents the time in minutes, with 0 indicating the time when the post was created. The y-axis represents the net ‘yakarma’, or score, of the post. The artificial votes given to a post are factored out of this net score. For example, if this post was given 10 upvotes initially, these 10 upvotes are subtracted from the overall score for each data point for that particular trial. This was done for the purpose of only observing the votes of users that we did not instruct to vote a specific way. The black data points indicate the control group. As seen in the graph, this post was quickly downvoted to a score of -5 when we did not influence it’s initial score. This indicates most users would typically downvote this post. The red data points indicate the ‘negative inflation’ trial. In this trial we artificially inflated the post with two downvotes immediately after posting. The post almost immediately reached a score of negative -5 (the net yakarma is -3 because the two artificial downvotes are removed from net score). The green data points indicate the “positive inflation’ trial. In this trial we artificially

inflated the post with upvotes, specifically 10 for this post (to reiterate, the first green data point has a score of 0 because the artificial upvotes were removed from the net score). As seen by the graph, users voted much differently for this trial. The post had an increase in score over time, and finished with a final net score of 29. This is a drastic change from the other two trials, which achieved a score of -5 within 10 minutes. Assuming we had a similar user base for each trial, this indicates that users voted differently solely because the score of the post was higher. Ordinarily users had a low signal for this post, but they voted with a high signal when the post was artificially inflated to appear popular. This is an information cascade. All low quality posts were created showed the same results; they were ordinarily immediately downvote, but for the positive vote inflation trial they showed an increase in net score over time.

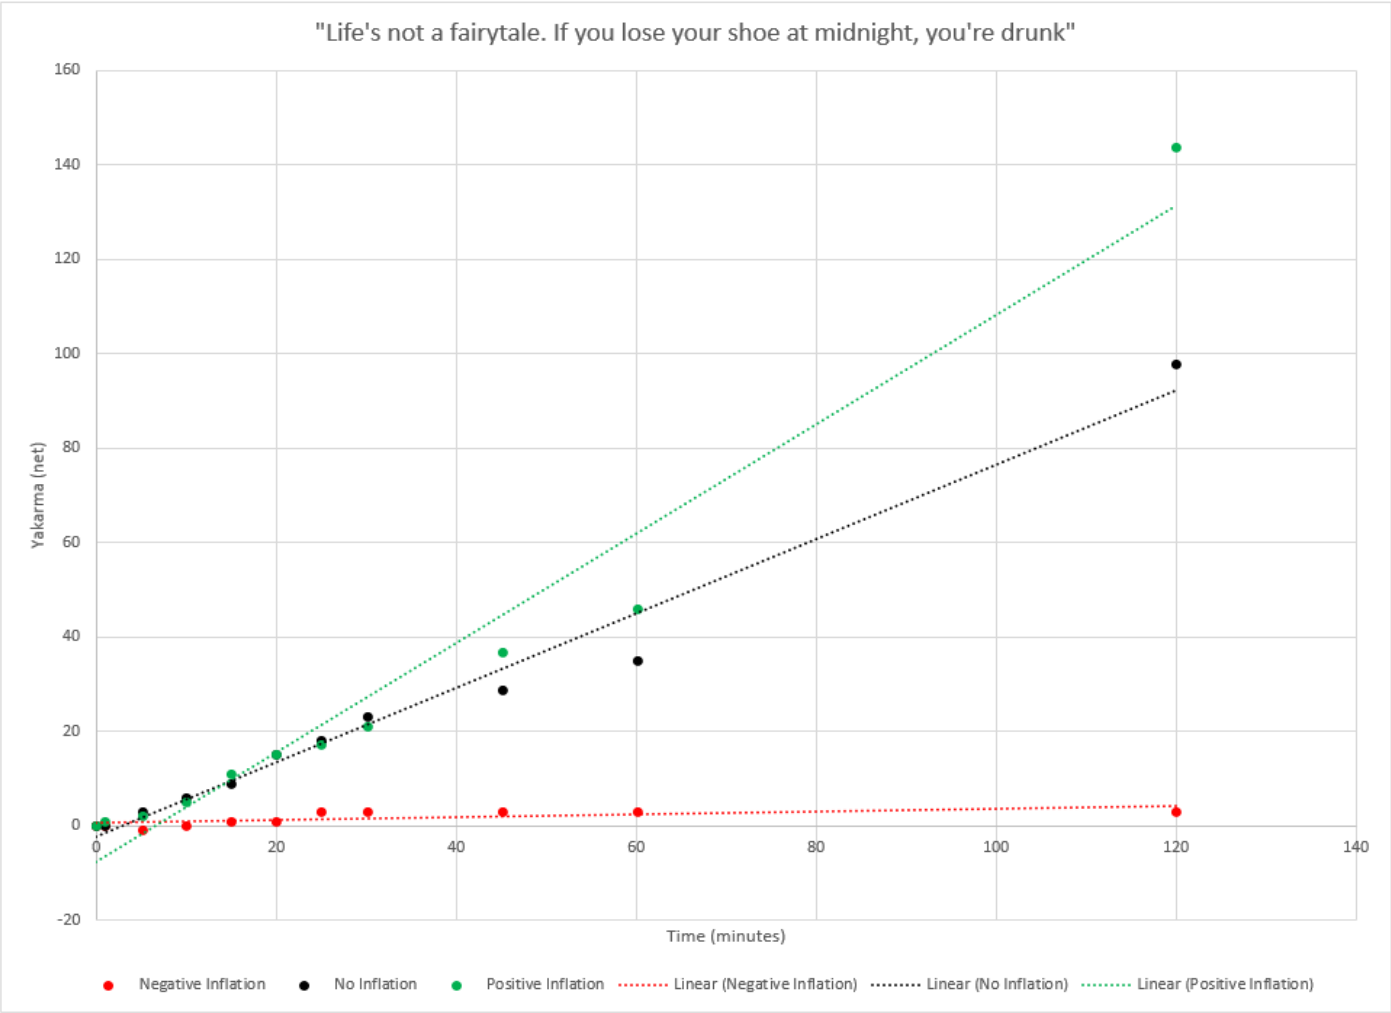

The graph above is an example of a popular post that we created and tested in our study. Specifically, the post is “Life's not a fairytale. If you lose your shoe at midnight, you're drunk.” As with the previous graph, the x-axis represents the time in minutes, with 0 indicating the time when the post was created, and the y-axis represents the net ‘yakarma’ of the post. The black data points indicate the control group. As seen in the graph, this post saw a large

increase in score over time, indicating that it is indeed a popular post. The red data points indicate the ‘negative inflation’ trial. Similar to the last graph, this post received two downvotes initially in this trial. As seen in the graph, these two downvotes had a very significant effect on the score of the post over time. Users voted much differently in this trial compared to the control trial, indicating the users voted against their own private signal signal simply because the post started with a score of -2. Therefore, can conclude that we observed another information cascade. The green data points indicate the ‘positive inflation’ trial. As with the last graph, this post received artificial upvotes immediately after it was posted. As seen in the graph, there was a faster rate of increase in this trial compared to the control group. This indicates that some users voted differently in this trial as a result of the higher initial score, suggesting there was an information cascade for this trial as well, but it should be noted that it is not nearly as significant as for the previous graph (the low quality post). This argues that the information cascade that results as a result of positive artificial vote inflation depends on the quality of the post, with information cascade having a greater effect when the post is of low quality and a lesser effect when the post is of high quality.

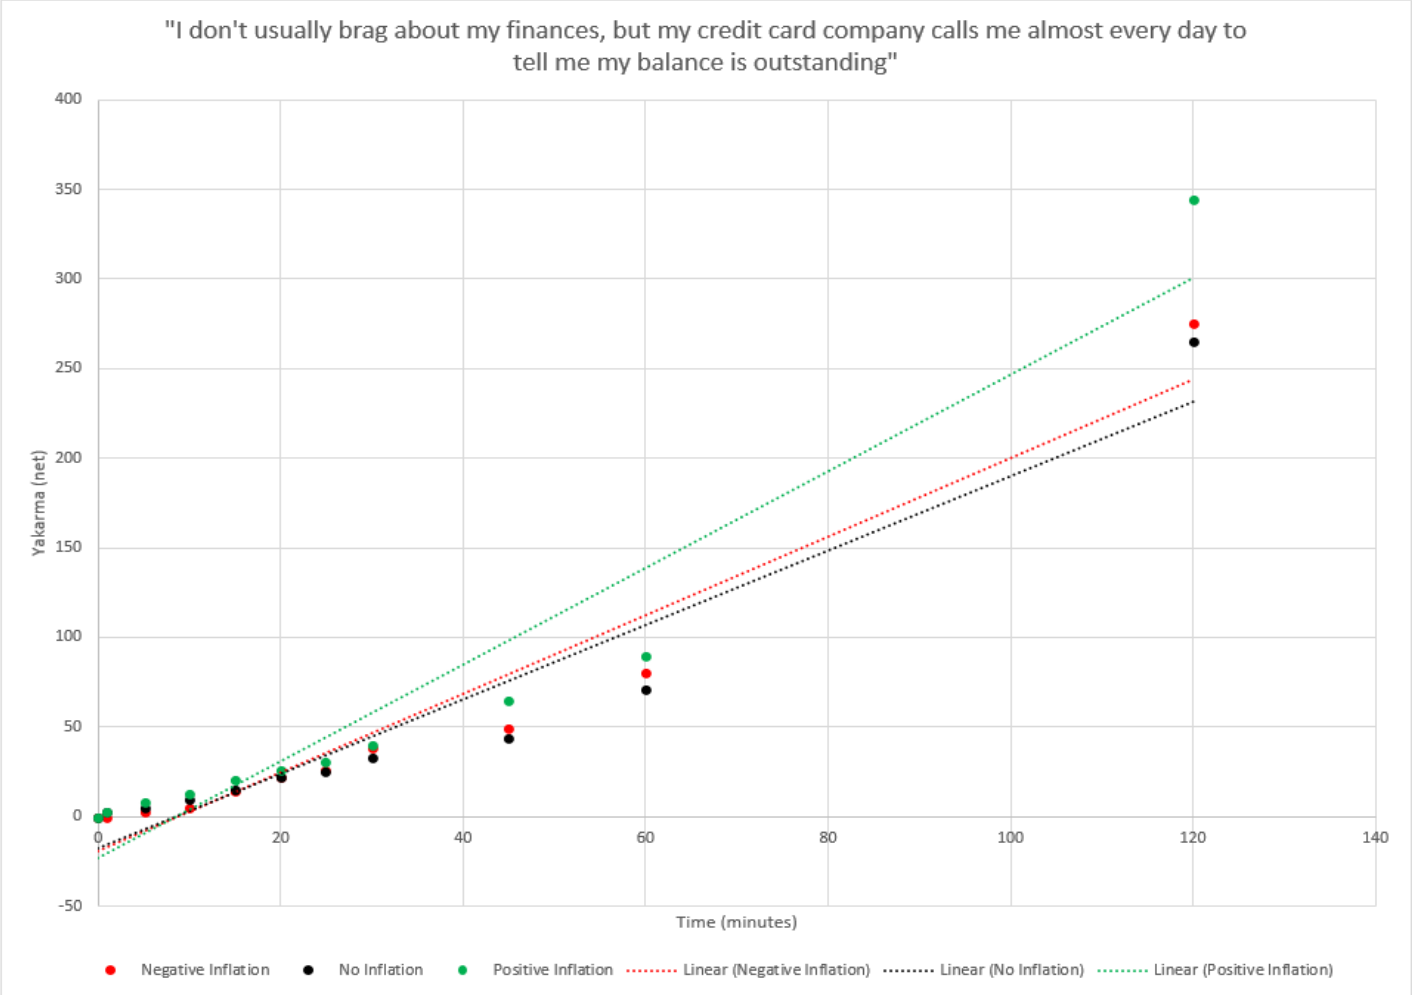

The graph above is an example of a very popular post we created and tested in our study. Specifically, the post is “I don't usually brag about my finances, but my credit card company calls me almost every day to tell me my balance is outstanding.” As with the previous graphs, the x-axis represents the time in minutes, with 0 indicating the time when the post was created, and the y-axis represents the net ‘yakarma’ of the post. The black data points indicate the control group. As seen in the graph, this post saw a very large increase in score over time, indicating that it is indeed a very popular post. The red data points indicate the ‘negative inflation’ trial. As with the last graphs, this post received two downvotes initially in this trial. As seen in the graph, the two downvotes had essentially no effect on the score of the post. This trial score very similarly to the control group, indicating that the small downvote inflation had no success in attempting to create an information cascade. The green data points indicate the ‘positive inflation’ trial. As with the last graph, this post received artificial upvotes immediately after it was posted. As seen in the graph, there was a faster rate of increase in this trial compared to the control group. This indicates that there was an information cascade, and strengthens the argument that the information cascade that results as a result of positive artificial vote inflation depends on the quality of the post, with information cascade having a greater effect when the post is of low quality and a lesser effect when the post is of high quality. All very popular post performed similarly to the example in this graph.

What we found is that subjectively ‘bad’ posts that were upvoted by a large number of users continued to be upvoted at a much higher rate than the unaffected control test. All across the board, posts that were immediately upvoted past a certain point quickly gained popularity and continued to increase in vote count. We did some experiments with smaller samples of posts and found this threshold for a cascade to be as little as 3, but with much more pronounced cascades visible when posts were upvoted to at least 10, and the most when upvoted to 20. Not only did this cause our posts to look very popular when sorted chronologically, but by upvoting them faster, posts made it to the top of the popular feed faster, causing even more users to see them and vote on them.

Similarly, we also found that while subjectively ‘good’ posts that were downvoted did not completely die out, they increased at a much slower rate compared to their control and upvoted counterparts. This also suggests an information cascade, because although we didn’t completely stop the popularity of a post, we affected it enough to change its rate of increase.

Our data points to a relationship between post quality and the significance of the effect of the information cacde. For low quality posts, negative vote inflation was essentially meaningless, because users were going to downvote the post anyway. However, positive vote inflation had a much larger effect for low quality posts. For high quality posts, positive vote inflation had little effect, and negative vote inflation had almost no effect. Positive vote inflation for high quality posts and negative vote inflation for low quality posts having little effect makes sense. The lower bound of -5 is most likely the reason negative vote inflation had no effect on high quality posts.

We also did comparison tests to determine what the threshold value was for the start of an information cascade in the down direction. After multiple samples and experiments, we found this threshold to be between -2 and -3. Posts that were initialized to -2 occasionally made a comeback and became very popular, whereas posts initialized to -3 were almost immediately voted off. As with the positive direction, this has a surprising amount to do with where exactly the post appears on the feed. When a post is first created, it appears at the top of the chronological feed, and if it’s immediately downvoted, users will see that and continue

to downvote it, as it must not be very popular. We found that if a post lasts at least 5 minutes, it tends to stick around and never die out. This is due mainly to most users not scrolling down very far on the feed to read it, and since it never became popular, it won't show up at the top of the popular feed either, meaning it will simply sit there until enough new posts have been created and the post eventually disappears from the feed due to age.

Surprisingly, very little research has been done with regards to information cascades via social media. Most of the studies we've found involve looking at the popularity of individual users and what effect it has on the outcome of a post[1,2]. However, since Yik Yak is anonymous, information cascades can have a vastly different effect, and we were unable to find any additional research to back up our claim.

D. Conclusion

From our analysis, we can conclude that information cascades can be produced in this setting by influencing the scores on posts. There were multiple findings that stood out to us, including how the quality of a post affects the severity of a cascade. Given an objectively good post, we were unable to affect it much in the negative direction, even when downvoted to -2. This suggests that there is some post value at which previous signals are ignored and personal opinion trumps what the cascade suggests. In the case of posts downvoted to -2, all it took was a few people in a row to upvote it back to 0, then the post followed its control sample at a slightly lower rate.

Furthermore, it was interesting to find that there was a sharp cutoff between -2 and -3 for negative cascades, but a much more indefinite line for positive cascades (between 3 and 10 upvotes). This is probably due to the fact that there is a limit at -5 when posts are deleted, and a post at -3 only requires a user and one other person to vote on it for it to disappear, whereas a post at -2 requires a user and *two* other people to vote. This doubling of the number of other users makes it much harder for a negative cascade to start at -2 than it is to start at -3.

Although we have shown that it is possible to create information cascades in this environment, our procedure had a few limitations. One limitation was the consistency of the user base. This was an issue because at times the user base could either be too consistent or not consistent enough. If the user base was too consistent, then the audience viewing our posts would notice that we were repeating them more than once. This negatively affected our posts because users of Yik Yak generally do not encourage posting content that was on the forum recently. Also, the quality of a post is much lower for a user when they have already seen it before. If the user base was too inconsistent, then it was hard to determine the quality of a post, since its quality varies from person to person based on their interests.

Another limitation we ran into was that of certain events causing more users to vote on Yik Yak posts. For example, we made a few posts immediately after the Organic Chemistry exam ended, and there was a large influx of users voting on posts. Since this was not the normal user base, the rate of change for the scores were vastly different from anything we had seen before. We concluded that we would be unable to duplicate the same user base, and therefore marked the posts as invalid and discarded the results.

When setting up this experiment, we did as much as we could to track the rate of change in votes, but unfortunately it's hard to get data points for every minute. Instead, we were forced to make cutbacks and recorded the post scores every 5 minutes up to 30 minutes, then after an hour, and finally when the post disappeared due to time. If we had been able to access the Yik Yak framework, we could have designed a program to record the

post score as a continuous function. However, it's also a good thing that we were unable to access the framework, because that would mean that others could as well, and thereby make it so that they can vote an unlimited number of times on each post. By restricting access of their application to mobile users in a small geographic area, Yik Yak has ensured that votes are authentic, and that each user only gets to vote once. This further backs up the idea of information cascades, since each user has an equal say in the final score of a post, so the final score of a post is the net number of users that voted for the post.

Another question that arose during this experiment is "what would happen if there were no limit on the number of downvotes a post can receive before it is deleted?". Typically when a post is downvoted to -5 and deleted, it is because the post is objectively bad or it was offensive to an individual or group. By removing the lower limit for posts, users could continue voting on the posts, and there could be multiple waves of information cascades, i.e. from an upward cascade to downward cascade, and visa versa, depending on the popular opinion of the group of users voting.

Observing information cascades in social media implies that a popular opinion may have gained it's status due to it's popularity and not only it's content. This makes it possible to take advantage of a social media platform that has a voting system by inflating the votes. Misleading and inaccurate information could potentially gain too much popularity if this sort of inflation is applied to it and exposed to the public. However, the key point to understand is that although it is possible to create information cascades in this manner, the content of the post still has an impact on the user's perception of its quality. Given a high quality post, an information cascade will affect the post's popularity less than the content itself.

E. Sources

1. Lerman, Kristina, and Rumi Ghosh. "Information Contagion: N Empirical Study of the Spread of News on Digg and Twitter Social Networks." Association for the Advancement of Artificial Intelligence. N.p., 2010. Web. 5 May 2015.

2. Romero, Daniel M., Wojciech Galuba, Sitaram Asur, and Bernardo A. Huberman. "Influence and Passivity in Social Media." (2011): n. pag. Web. 4 May 2015.