STANFORD UNIVERSITY

Symbolic Systems Masters Program Final Project

Temporal dynamics of adoption and diffusion patterns in online petitioning

Chi Ling, Chan

Primary Advisor and First Reader

DR. TODD DAVIES

Second Reader
PROF. LEE ROSS

August 29, 2015

Acknowledgement

The author wishes to offer her sincere gratitude to Dr. Todd Davies for his invaluable guidance and advice, Prof. Lee Ross and Prof. Howard Rheingold for timely feedback, and last but not least - Bryan Hooi for his contribution and time to this project.

1 Introduction

The combination of networked individualism with Web 2.0 digital tools has created a new public space (Castells, 2012) in which social agents organize social change. Petitioning sites like change.org and wethepeople.gov are enabling citizens to participate in social campaigns online, facilitating a new type of collective action in the interstitial space between the digital and the urban. The availability of temporal data on these activities has made it possible to observe dynamics of diffusion and adoption crucial for understanding how momentum is generated for successful social campaigns.

This study observes temporal dynamics of diffusion patterns in online petitions-signing, in order to understand what drives growth momentum of support and help petitions gain traction. Using aggregated data from the White House's WeThePeople petition platform, we observe temporal characteristics of online petitions and test the explanatory power of the traditional S-curve model of mobilization. Thereafter, we test hypotheses on the effects of thresholds and distill systematic diffusion patterns that are correlated with a petitions' eventual success. Finally, we generate broad contours of diffusion dynamics by running simulations based on a combination of diffusion models.

2 Related Work

In the literature on diffusion and contagion, the dominant model predicts that new adopters are influenced by previous adopters within a population, giving rise to an S-shaped curve (Coleman et al. 1957, Rogers 1962, Bass 1969, Valente 1995, Young 2009, Iyengar et al. 2010). An S-shaped curve exhibits an initial period of exponential growth which eventually levels off when the population runs out of potential adopters. Mathematically, it may be defined by the following sigmoid function:

$$S(t) = \frac{1}{1 + e^{-t}}$$

This results in an S-shaped curve, as in Figure 2:

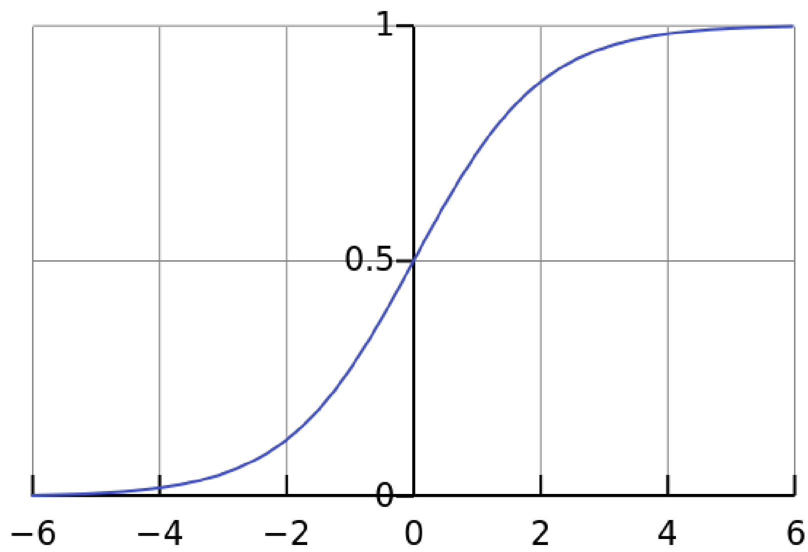

Figure 1: Standard logistic sigmoid function

Whether this diffusion pattern applies in the online setting has been a question of interest in recent years, especially given the increasing prominence of online collective mobilization. Other investigations into the the role of social networking sites in the growth of mobilizations have identified an S-shaped curve, with critical mass reached only after participants have responded to evidence of contribution of early participants (Gonzalez-Bailon et al., 2011) However, evidence remains mixed. In a study that tracks the growth curves of 20,000 petitions on the government petitioning site in the United Kingdom, Yasseri et al. (2013) notes that a petition's fate is virtually set after the first 24 hours of its introduction. The decisive impact of the first day on overall success of a campaign has been echoed in a separate study by Hale et al. (2013), which tracked 8000 petitions on UK petitioning site No. 10 Downing Street. The period of stasis before reaching a critical point appears largely absent, a finding consistent with most studies on online petitioning. These findings collectively call into question the explanatory power of the S-curve model. Given that research on online petitioning is a relatively new area of study, however, there remains a lack of empirical calibration and external validity - a point acknowledged by authors of most of these studies (Gonzalez-Bailon et al., 2011; Hale et al. 2013).

A second question has to do with the diffusion mechanism behind observed patterns. Empirical studies have largely interpreted the S-curve as evidence of social contagion (Rogers 1962; Bass 1969), but others suggest that the same curve could arise from broadcast distribution mechanisms such as mass media sharing (Van den Bulte and Lilien 2001). It remains ambiguous as to whether viral diffusion is in fact the diffusion mechanism driving growth momentum, as is typically interpreted in classical diffusion studies. Related research has shown that information cascades in online networks occur very rarely, and studies on online petitioning have found that the vast majority of signatures on online petitioning platforms are dominated by a tiny fraction of massively successful petitions (Yasseri et al., 2013; Jungherr et al., 2010; Hale et al., 2014). Events that 'go viral' are an exception rather than the rule, and a rare one at that. Untangling broadcast from viral diffusion mechanisms, however, has been met with difficulty largely because these studies have based themselves on aggregate diffusion data.

Recent studies have been able to overcome this limitation using cascade data and identify the diffusion patterns by observing their structures. For example, Goel et al. (2013) samples cascades from a data set of a billion diffusion events on Twitter and offers a more fine-grained analysis of how viral and broadcast diffusion interact. The study finds that large diffusion events exhibit extreme diversity of structural forms, and demonstrate a continuous distribution of viral and broadcast diffusion, such that the S-curve is but one out of many conceivable combinations. Extending Goel's analysis of viral and broadcast diffusion, we might ask related questions on how patterns develop over time: when does viral or broadcast diffusion set in, and how do they combine to generate the adoption patterns observed in aggregated online petitioning data? Hale et al. (2013) observes that signatures are typically gathered via the repeated occurrence of 'punctuated equilibrium', specific points that trigger large cascades in signatures within a short time with the result of having leptokurtic distributions (i.e. characterized by sharp peaks in signature counts). Several questions remain standing: When in the span of a petition do these points occur? Are there consistent diffusion patterns or interactions between viral/broadcast diffusions that systematically characterize successful petitions? Given that online petitions have temporal thresholds, understanding not just which diffusion models work best, but also when they do, can tell us something about the underlying momentum that drives the success and failure of a petition.

Finally, a related but less well-studied aspect of online petitioning and diffusion has to do with the effects of threshold requirements. Signature thresholds are a typical feature in most online petitions crucial for goal-setting for campaigners, and function as important social information on how proximate a petition is from its target. There are reasons to believe that thresholds matter and, more specifically, that proximity to goals is positively correlated with willingness to participate. For example, various studies in social psychology have demonstrated that people invest greater efforts as they approach a goal (Kivetz, Urminsky and Zheng, 2006; Cheema and Bagchi, 2011; Koo and Fishbach 2012). This phenomenon has come to be known as the goal-gradient hypothesis, first described by behaviorist Clark Hull (1934) in his observations of rats running faster as they approach a food reward in a maze. A recent study on microlending online by Cryder, Loewenstein and Seltman (2013) finds that people are more willing to pitch in as charitable campaigns approach their goals. In the context of online petitioning, there is suggestive evidence in Hale (2013) for threshold effects at the 500-signature mark (the minimum number of signature required for official response on No.10 Downing Street), but it remains unclear if this can be extrapolated across other petition platforms such as the one this study focuses on (i.e., We The People).

3 Questions and hypotheses

In light of existing literature, we identify three primary questions and, where applicable, relevant hypotheses that are testable with our data set:

- What patterns do we observe in online petitions across time, and how well does the S-curve model explain these patterns?

- What are the effects of different thresholds on diffusion?

  Hypothesis 1: Diffusion momentum slows after crossing a temporal threshold - the number of signatures increases at a slower rate after a petition deadline has been reached.

  Hypothesis 2: Diffusion momentum slows after crossing a signature threshold - number of signatures increases at a slower rate after having reached its target number of signatures.

- Are there systematic differences in diffusion patterns between successful and unsuccessful petitions? We hypothesize that less successful petitions receive an initial burst of signatures but taper off quickly, while more successful petitions build momentum over time through viral effects. In particular:

  Hypothesis 3: Less successful petitions have a more peaked distribution (higher kurtosis) and fewer signatures in their right tail (lower skewness) compared to more successful petitions.

  Hypothesis 4: More successful petitions take longer to reach their peak (i.e. the day when they receive the most signatures)

4 Data and Methods

4.1 We The People petitions data

As previously mentioned, studies on diffusion patterns and social contagion in social movements tend to be limited by the quality of data analyzed, particularly around time dynamics. We find that online networks and petitioning platforms offer a great opportunity to explore diffusion patterns in an empirical setting. In this observational study, we rely on an aggregated adoption dataset of 3682 publicly searchable petitions on the public API of WeThePeople.gov, the official online petitioning platform of the White House. These petitions were created within a time period that began from the inception of the platform on September 20, 2011, to March 30, 2015. Within this time period, the White House raised the threshold for official review from 25,000 to the present 100,000 threshold in January 2013. Our analysis accounts for this policy change by separating and sub-setting data points accordingly.

Two critical thresholds for We the People should be noted. First, to cross the first threshold and be publicly listed and searchable within the site, a petition must reach 150 signatures within 30 days. Second, to cross the second threshold for review by the White House and to be distributed to the appropriate policy officials within the Administration for an official response, a petition must reach 100,000 signatures within 30 days. This response will be posted and linked to the petition on WhiteHouse.gov, as well as emailed to all petition signers. Petitions that do not cross this threshold in the given timeframe will be closed and removed from the site after 60 days. All petitions in this dataset crossed the 150 signatures threshold, and are publicly listed - excluding petitions with fewer than 150 signatures was less a decision and more of a feature of the data available through the API: the API allowed the retrieval of only petitions that were publicly searchable. Apart from that, the data retrieved from the API includes a vast majority of petitions and is essentially unbiased.

4.2 Variables in signatures and petitions dataset

The WeThePeople API provides data on both petitions and the individual signatures that were submitted via the petition site. Data on petitions include timestamp of creation, the actual body of text that campaigners submit, present status (open/pending response/responded/closed), and the signature count. The data tied to individual signatures logs the time a signature was submitted, as well as the geographical details (state, zipcode, country, city) of the signatory.

Pulling data from the API, we put together a dataset comprising all signatures and petitions and organized them into the following variables: 1) petition ID; 2) signature ID; 3) Unix timestamp of signature; 4) zipcode of signatory. This was then merged with a dataset of petitions containing the following data: 1) petition ID; 2) petition title; 3) petition description; 4) signature count; 5) signature status; 6) Unix timestamp of creation (Unix timestamps were recoded to reflect number of days since a petition's creation). These variables were chosen for their relevance to the research questions outlined in Section 3.

The "success" variable was coded based on the following condition: a petition is considered successful if it reaches the 100,000 threshold (the threshold necessary for official review by the White House). Based on this condition, we find that a large majority (98.4%) of petitions fail. Of all visible petitions, 1.6% eventually successfully reached the 100,000 signature threshold. This success rate is consistent with the predicted pattern that only a small fraction of campaigns eventually succeed (Yasseri et al. 2013).

5 Empirical Findings

5.1 Overall adoption patterns

Overall, a total of 24.5 million signatures were collected by 3682 petitions. Of these 3682 petitions, 1.6% reached the 100,000 signatures threshold required by an official response from the White House. Figure 2 shows the overall distribution of signatures in the dataset by plotting the total number of signatures for each petition against the rank order of the petition by total number of signatures. The distribution is left-skewed towards successful petitions, with successful petitions (i.e., petitions with more than 100,000 signatures, which was 1.6% of all petitions) taking 31.8% of total signatures. Two discontinuities are observed at two thresholds - the 100,000 signatures threshold and the 25,000 signatures threshold, suggesting threshold effects that will be discussed in greater depth in Section 5.2.

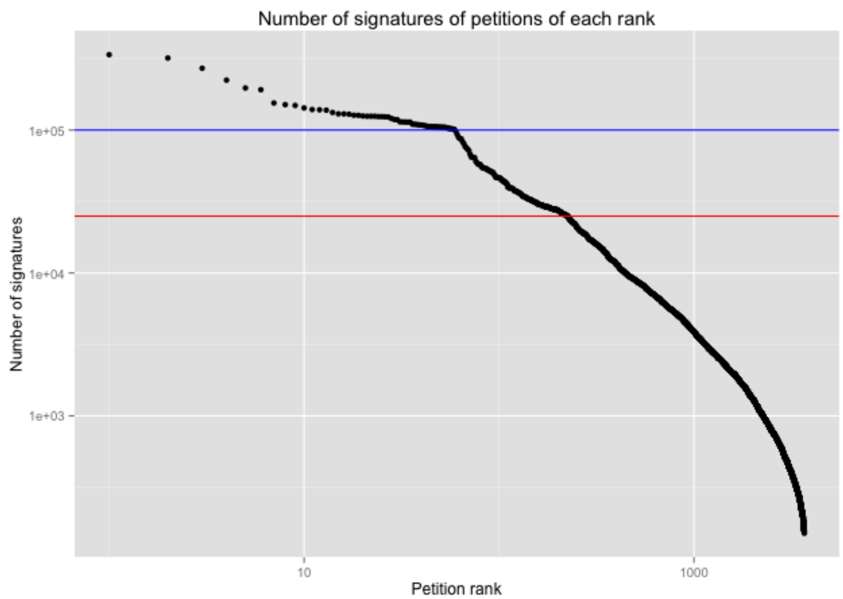

Figure 2: Total number of signatures per petitions plotted against the rank order of the petition based on total number of signatures.

Figure 3 shows the cumulative curves of signatures gathered for each petition over a 60-day period, on a normalized y-axis. It is clear from this figure that a large number of signatures are collected shortly after the launch of a petition, and growth in support levels off soon after. On reaching the 30th day deadline, petitions exhibit a sharp decrease in rate of growth in support.

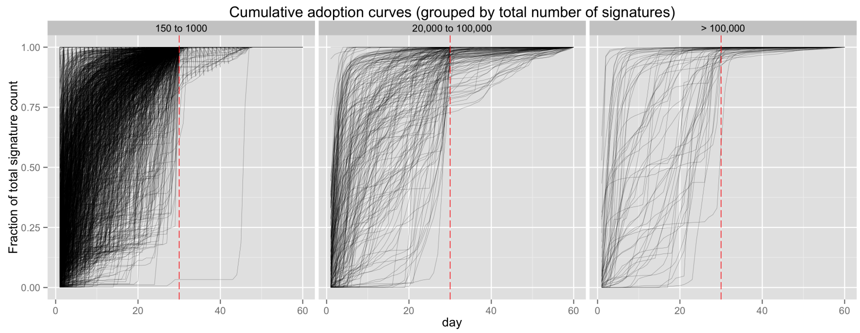

Figure 3: Cumulative density function of signatures gathered for all 3682 petitions across a 60-day period, separated into subgroups based on their total number of signatures received. The y-axis is the cumulative fraction of signatures received. In all subgroups, the petitions receive fewer signatures after crossing the 30-day point.

We return to key questions outlined in Section 3. Firstly, to what extent is the S-curve model a good characterization of growth momentum in online petitions in our data set? Based on Figure 3, the typical pattern of early onset of support does not confer support for the classical S-curve model that predicts a slow build-up period followed by accelerating growth and eventual decay. Rather than a slow accumulation growth of supporters building up to critical mass, petitions demonstrate rapid early growth and decelerates over time. This is consistent with findings made by Hale et al. (2014) demonstrating rapid dynamics in the growth of signatures among petitions.

Secondly, are there systematic differences in the cumulative density functions among highly successful petitions (those with more than 100,000 signatures) and less successful petitions? In Figure 3 we subset plots of cumulative density functions based on a petition's final number of signatures. However, the small number of successful petitions does not facilitate a fair comparison vis-a-vis the much larger pool of unsuccessful petitions does not allow us to generalize conclusive systematic differences based on a visual comparison of cumulative density functions. Nevertheless, we will proceed with a more formal analysis of the shape of adoption curves and a petition's eventual success in Section 5.

5.2 Temporal threshold effects

What are the effects of temporal thresholds on the growth momentum of support for petitions - we may ask, for example, whether growth momentum drops after crossing the threshold (Hypothesis 1)? The key temporal threshold here is the 30-day mark - all petitions are given 30 days to reach the signature threshold for an official response, and petitions that fail to reach the threshold will be removed from the site on day 60.

Figure 4 shows the cumulative adoption curves of all 3682 petitions, color-coded according to the final number of signatures accumulated. As noted previously, most petitions are sensitive to the 30-day deadline for official response and the level of support typically falls drastically after crossing the 30-day mark.

With reference to the third question in Section 3 - are there differences in temporal patterns between successful and unsuccessful petitions? Our plot shows that activity tends to persist for highly successful petitions with a high eventual signature count, exhibiting residual momentum indicated by the purple, blue and violet lines that dominate the graph after the day-30 mark (Figure 4.)

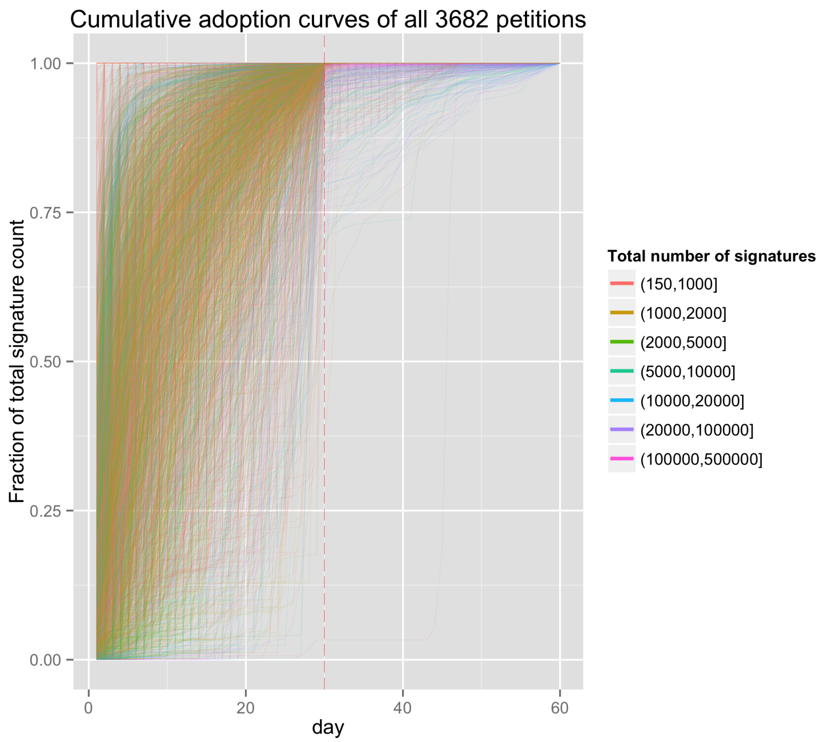

Figure 4: Total number of signatures per petitions plotted against the rank order of the petition based on total number of signatures.

Figure 5 shows adoption curves capturing signatures gathered daily. Overall, we see that the number of signatures for most petitions decline over the 30-day period, before experiencing a spike as it approaches the 30th day. This surge in support may be explained by a last- minute campaign for support before the petition loses any chance of receiving an official response from the White House. This, along with the sharp fall in number of signatures after the 30-day mark, indicates high sensitivity to the temporal threshold, confirming Hypothesis 1.

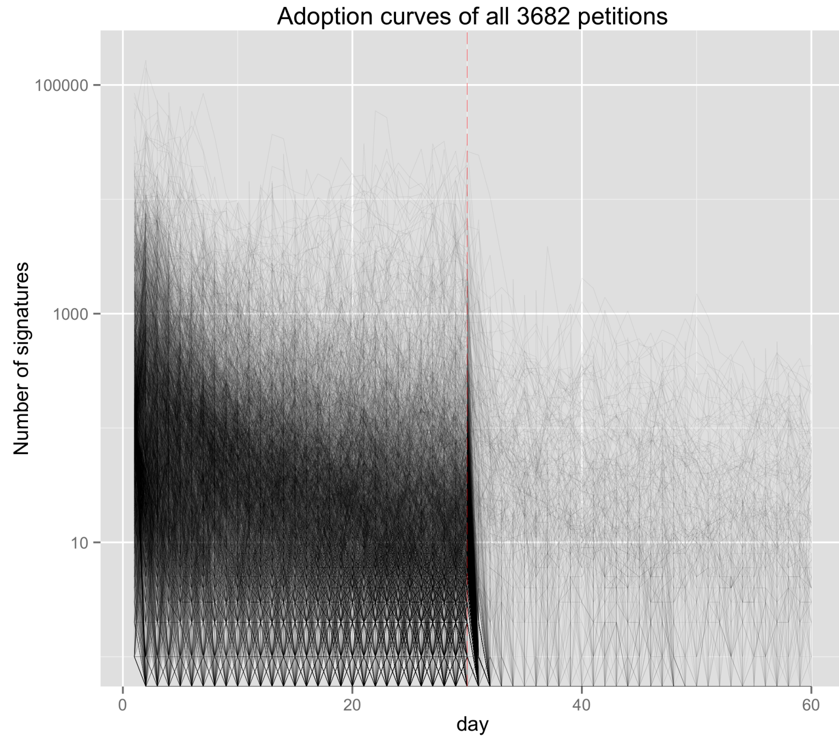

Figure 5: Adoption curves capturing daily accumulation of signatures in 3682 petitions across a 60-day period. A clear spike indicating a surge in support for a large number of petitions right before the 30-day deadline.

Which petitions are most sensitive to temporal thresholds? Figure 6 exhibits the same adoption curves as Figure 5, color coded based on the total signature count. We observe from the relative preponderance of blue, purple and violet curves after the 30-day mark that highly successful petitions are less sensitive to the temporal threshold. By contrast, less successful petitions exhibit little or no activity once the 30-day mark is reached. This makes sense intuitively - for less successful petitions, reaching the 30-day mark without accumulating $100,000$ signatures means they will not receive a reply from the White House; thus we expect less interest in such petitions from users. In contrast, for highly successful petitions, they may have accumulated $100,000$ signatures even before reaching the 30-day mark; thence, the 30-day mark becomes irrelevant since they have already reached the goal.

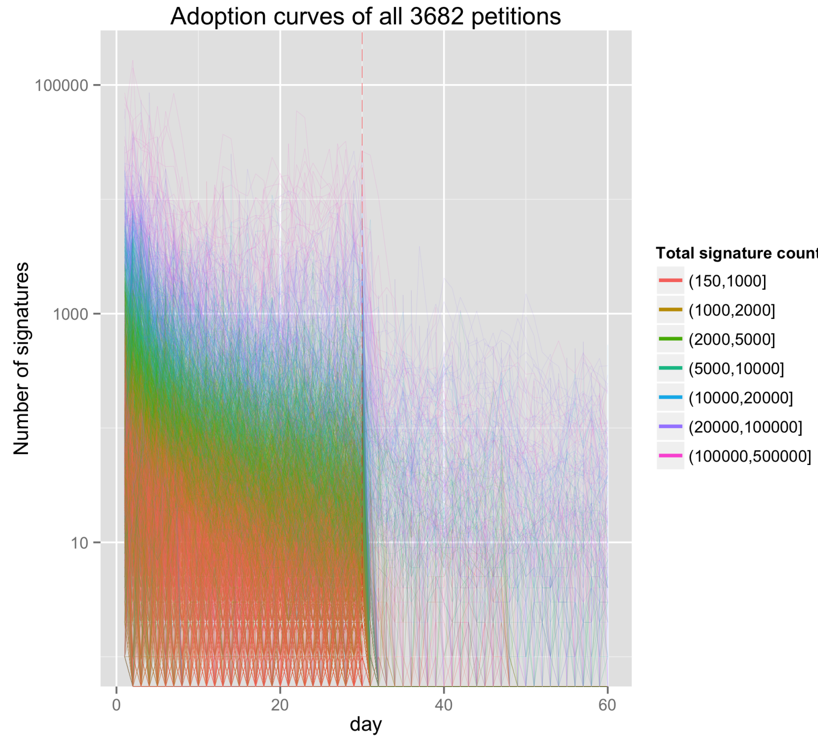

Figure 6: Adoption curves of all 3682 petitions; each line represents a single petition and shows its signature count per day over a 60-day period, colored by the petition's total signature count. There is in general a drastic drop in number of signatures around the 30-day mark, but highly successful petitions are less sensitive to this effect.

5.3 Signature threshold effects

Under many standard diffusion models, the virality of a spreading petition (indicative of how likely it is to spread from one individual to another) is governed by a fixed parameter such as the transmission probability or basic reproductive number, which does not change over time. On the "We the People" petition platform, however, due to the fact that a petition must reach $100,000$ signatures to receive an official response from the White House, we conjectured in Hypothesis 2 that petitions exhibit threshold effects, losing momentum for increasing the number of signatures once the $100,000$ target is met. This agrees with Hale (2013), who studied the No. 10 Downing Street petition site and provides suggestive evidence for threshold effects at the 500-signature mark. Understanding and quantifying signature threshold effects would have important implications in allowing petitioners to more carefully set targets for their petitions, as well as for how petition platforms display and make use of these target thresholds.

Indeed, the discontinuities in Figure 7 suggest the existence of threshold effects at the $25,000$ and $100,000$ signatures mark. Figure 7 shows the number of signatures gathered by each petition, plotted against the rank of the petition based on signature count, e.g. the leftmost point is the number of signatures received by the petition with the most signatures, and so on. The low density of points on the upper left of each plot demonstrates that only a few petitions make it beyond the 100,000 signature count threshold. In the right plot, above the 100, 000 signature threshold, the slope of the curve is noticeably shallower, indicating a drop in momentum for increasing the number of signatures after a target threshold is met. This point coincides with the 100, 000-signatures threshold required for an official response by the White House, consistent with a similar pattern previously observed in UK's petition platform, No. 10 Downing Street (Hale et. al, 2013).

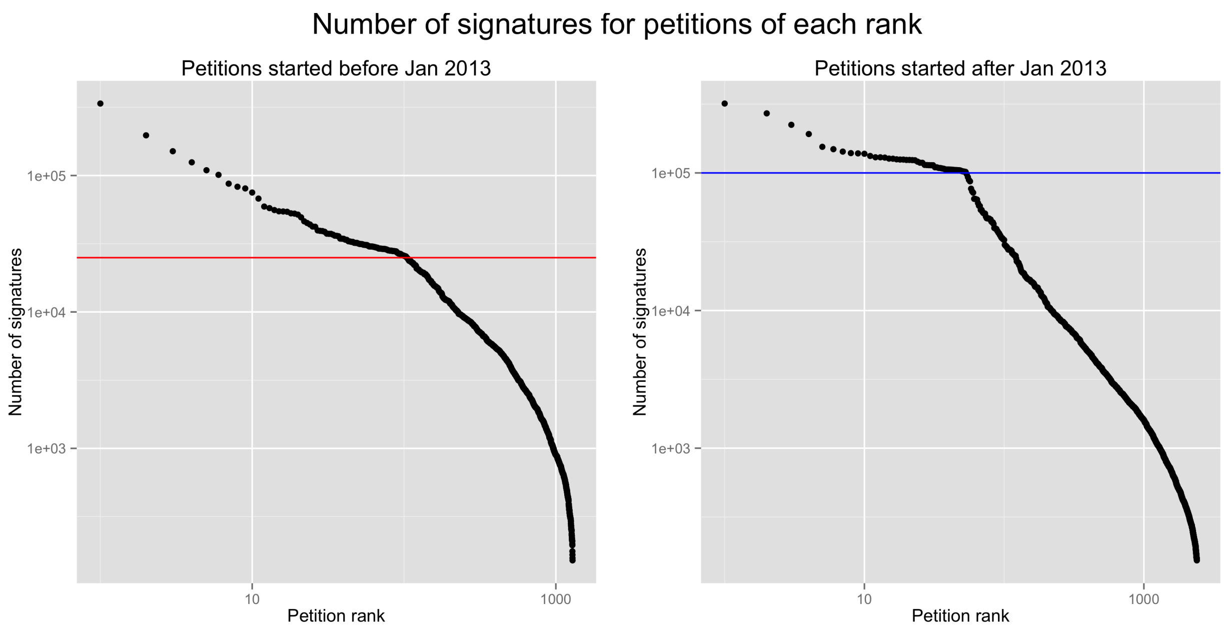

Figure 7: Threshold effects: for petitions started before Jan 15 2013 (left plot), when the threshold for a response was 25, 000, the graph has a discontinuity around the 25, 000 mark. After 15 Jan 2013 (right plot), the threshold became 100, 000. In each graph, the leftmost ('rank 1') petition is the one with the most signatures, the rank 2 petition has the second most, and so on. Horizontal lines indicate 25, 000 signatures (red) and 100, 000 signatures (blue). The discontinuity suggests a high concentration of petitions just above the thresholds, suggesting that petitions lose momentum after crossing the thresholds.

For petitions started before Jan 15 2013 (left plot), we instead observe a discontinuity at the 25, 000 signature mark. This point coincides with a previous 25, 000 signature mark threshold required for official response before the White House changed its threshold policy in Jan 2013.

Figure 8 also shows the drop in momentum that occurs after the threshold is crossed. In this plot, we consider only the subset of petitions started after Jan 15 2013. Each point represents the signature count of a petition on the days before and after crossing the threshold; most points lie below the line $y = x$, suggesting that almost all petitions receive fewer signatures on the day after they cross the threshold than the day before.

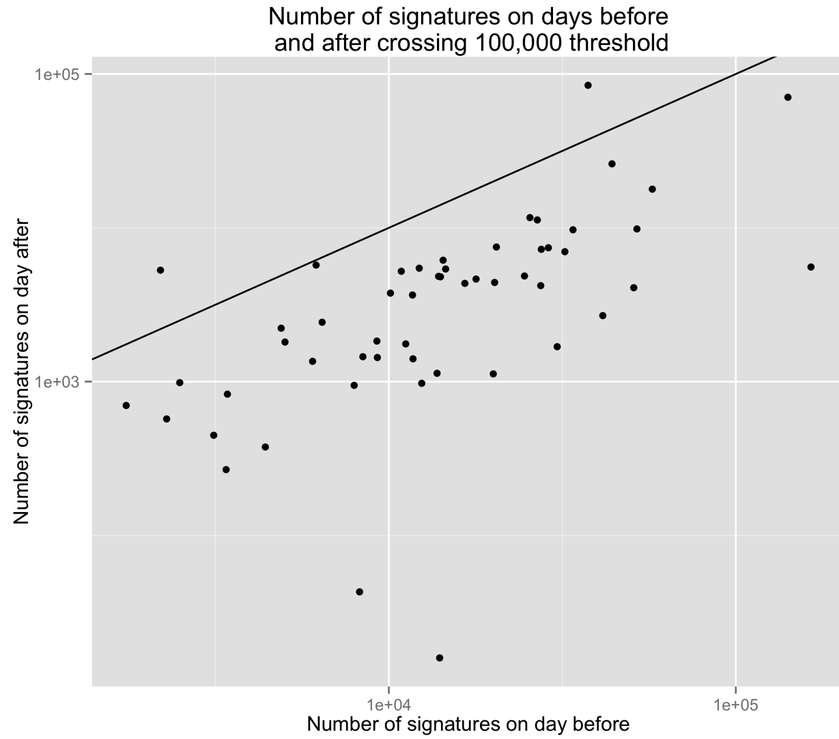

Figure 8: Petitions lose momentum after crossing the 100,000 signature threshold: each point represents the number of signatures a petition has on the day before and after it crosses the threshold. Almost all petitions receive fewer signatures on the day after they cross the threshold. The diagonal line is the line $y = x$.

However, a possible objection to this analysis might be that for most petitions, the number of signatures they receive per day gradually decreases over time; hence, they may receive fewer signatures on the day after the threshold purely due to this fact, rather than a genuine threshold effect at 100,000. Observe that if a threshold effect exists at the 100,000 mark, we should expect the drop in signatures per day as we cross the 100,000 mark to be distinctly greater than the corresponding drop that occurs as we cross the 80,000 or 120,000 signatures mark, for example. Indeed, as shown in Figure 9, the relative decrease in number of signatures at the 100,000 signatures mark tends to be greater than that which occurs at similar thresholds at the 60,000, 80,000 etc. levels, suggesting a genuine threshold effect at the 100,000 mark.

To formally test this point, we use the Mann-Whitney U-test (Mann et al., 1947), a nonparametric test for whether one population has larger values than another, which does not require any distributional assumptions, e.g. normal distributions. The test reports that the relative changes at the 100,000 threshold are significantly lower than at the other thresholds (p-value $\approx 2 \times 10^{-10}$). In all, this demonstrates that a signature threshold effect exists at the 100,000 signature level, confirming Hypothesis 2.

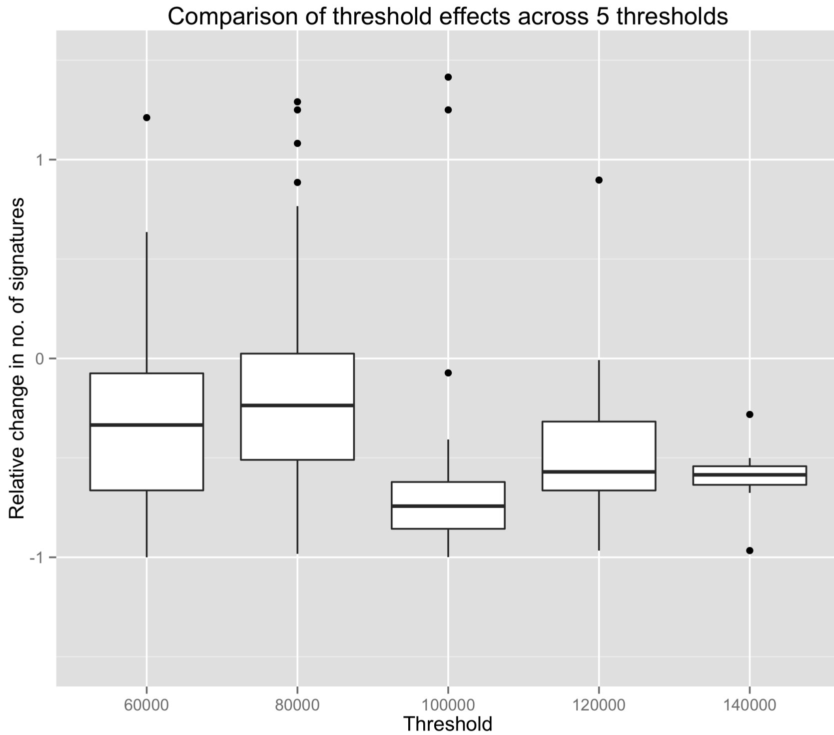

Figure 9: The threshold effect at 100,000 is stronger than at other nearby thresholds: the middle boxplot shows that in the day most petitions cross the 100,000 signatures mark, their signature count decreases by around 60% to 90% (as indicated by the box, which shows the 25th to 75th percentile). Furthermore, this drop is significantly stronger than that at neighboring thresholds (Mann-Whitney U-test: $p \approx 2 \times 10^{-10}$)

5.4 Temporality and signature thresholds

In the previous section, we observed threshold effects in petition adoption: the rate of receiving signatures substantially decreases after the signature threshold is reached. A follow-up question one might ask is: when in the life cycle of a petition do we observe the crossing of these thresholds?

Limiting our analysis on temporality and thresholds to petitions that were created after the White House raised the threshold to 100,000 signatures, we find that a vast majority of successful petitions made it past the threshold within the first thirty days of its maximum 60-day cycle. As shown in Figure 10, of the 57 petitions that successfully crossed the 100,000 signatures threshold, only 1 petition crossed the threshold after the 30-day midpoint.

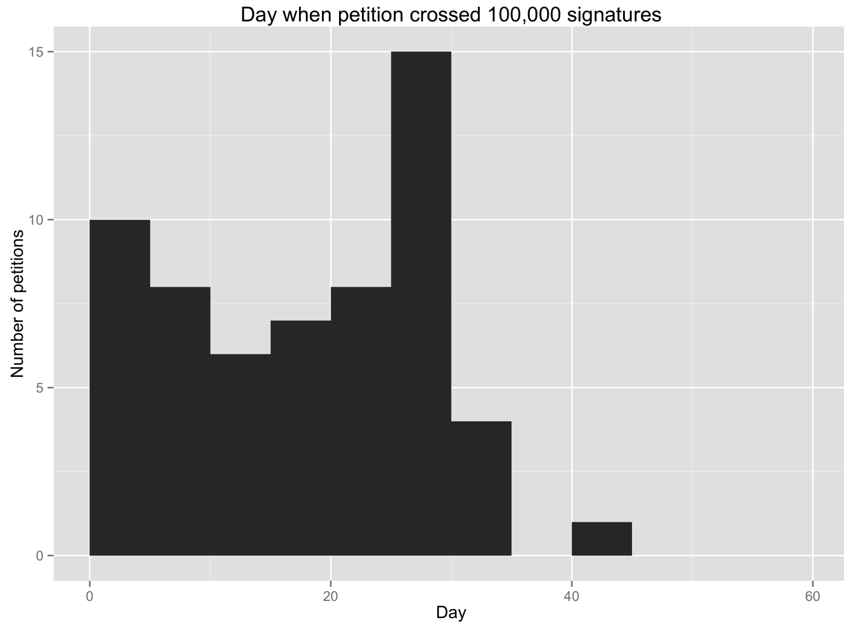

Figure 10: With the exception of one petition, all successful petitions created after White House revised its threshold crossed the 100,000 mark by the 30th-day midpoint.

5.5 Regression analysis of temporal distributions and petition success

Our next goal is to determine if there are systematic differences in the diffusion patterns between successful and unsuccessful petitions. For example, consider Figures 11 and 12: the top 5 are randomly chosen petitions with $\geq 100,000$ signatures, while the lower 5 are randomly chosen petitions with $< 100,000$ signatures. Our goal is to determine if petitions in the first group exhibit systematically different 'shapes' than the petitions in the second group: for example, do their number or location of peaks differ?

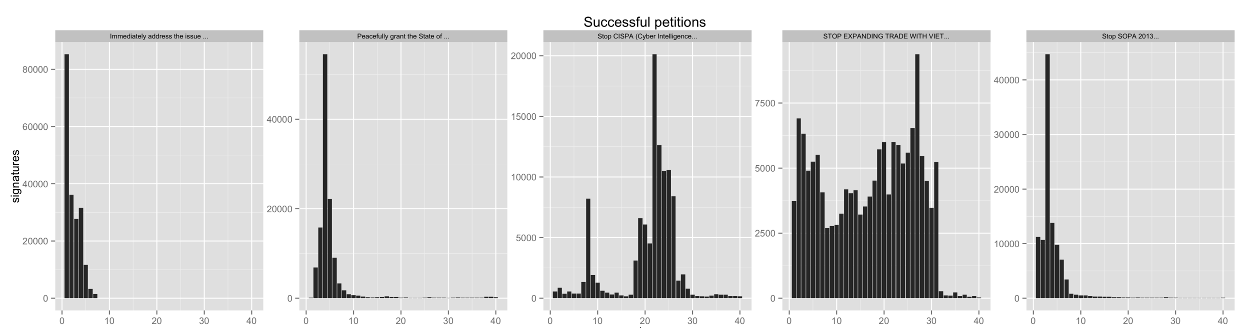

Figure 11: Temporal distributions of 5 randomly chosen petitions which reached the 100,000 signature mark.

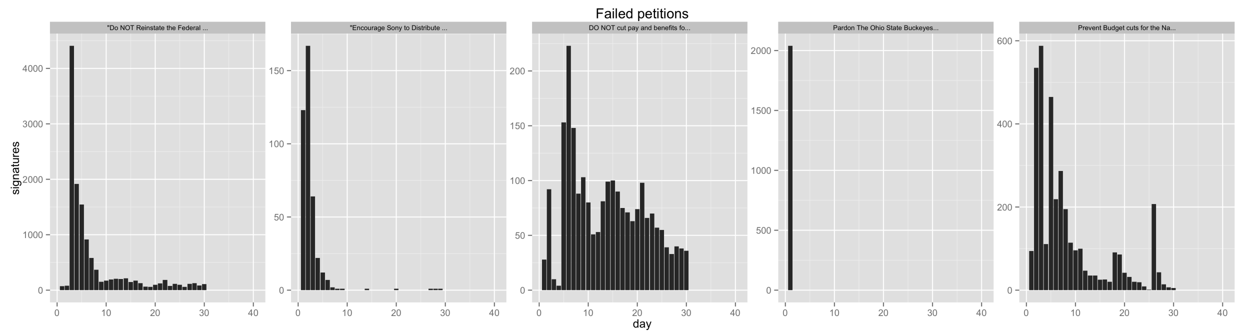

Figure 12: Temporal distribution of 5 randomly chosen petitions which did not reach the $100,000$ signature mark.

If so, such measures of shape would be useful in predicting a petition's future performance. Moreover, understanding how measures of shape relate to petition success would provide hints for better understanding the processes underlying petition spread.

We begin by testing Hypotheses 3 and 4: that less successful petitions have more peaked distributions and fewer signatures in their right tail, and that more successful petitions take longer to reach their peak.

5.5.1 Measures

The measures of interest are described in Table 1. We will refer to a petition's 'global peak' as the day at which it received the most signatures (where day 1 is the day the petition was introduced). The dependent variable, *total*, is the total number of signatures a petition acquires over the 60 day period.

Table 1: Measures of shape examined in this section

Measure of Shape	**Interpretation**
Skewness	whether distribution has larger 'tails' extending to the left or right
Kurtosis	how peaked a distribution is
Location of global peak	the day on which the petition received the most signatures
Number of local peaks	number of days on which the petition received more signatures than on adjacent days

5.5.2 Regression results

We perform a linear regression of number of signatures against our measures in Table 1; the results are shown in Table 2.

Table 2: Linear regression shows a highly significant positive relationship between peak day (peak_day) and final total number of signatures (total). Model (2) shows that this relationship remains significant when controlling for the number of peaks, variance, skewness, and kurtosis. Model (3) shows that this remains significant when log transforming the dependent variable (number of signatures).

	Dependent variable:			
		total		log(total)
	(1)	(2)	(3)	(4)
skewness	5,009.992***		5,875.782***	0.558***
	p = 0.00003		p = 0.00000	p = 0.000
kurtosis	−585.658***		−532.246***	−0.062***
	p = 0.00001		p = 0.00003	p = 0.000
global_peak_day		262.890***	227.812***	0.008**
		p = 0.000	p = 0.00000	p = 0.013
num_local_peaks			2,103.878***	0.150***
			p = 0.000	p = 0.000
Constant	−342.411	4,985.854***	−22,660.150***	5.455***
	p = 0.860	p = 0.000	p = 0.000	p = 0.000
Observations	3,682	3,682	3,682	3,682
R^2	0.006	0.010	0.087	0.089
Adjusted R^2	0.005	0.010	0.086	0.088
Residual Std. Error	19,316.340 (df = 3679)	19,271.360 (df = 3680)	18,518.850 (df = 3677)	1.364 (df = 3677)
F Statistic	10.355*** (df = 2; 3679)	36.997*** (df = 1; 3680)	87.054*** (df = 4; 3677)	90.216*** (df = 4; 3677)
Note:				*p<0.1; **p<0.05; ***p<0.01

As shown in Columns 1, 3 and 4 of Table 2, we find that skewness and kurtosis are also significantly correlated with signature count. Under all three model specifications, petitions with right skewed distributions (i.e. larger right tails) tend to end up with more signatures, and petitions with lower kurtosis (i.e. having less sharp peaks) also tend to end up with more signatures. This confirms Hypothesis 3, suggesting that more successful petitions do have systematically different signature distributions than less successful petitions.

As shown in column 2 of Table 2, linear regression suggests that on average, a petition that peaks 1 day later ends up with 262.89 more signatures ($p \approx 1.3 \times 10^{-9}$). The relationship remains about as strong, and still highly significant, when we control for the number of local peaks, and the skewness and kurtosis of the petition's temporal distribution (column 3) as well as when we replace the dependent variable by its logarithm (to ensure that petitions with large signature counts do not excessively influence the fitted coefficients).

Hence, we find that petitions with later global peak days tend to end up with more signatures than petitions with earlier peaks, confirming Hypothesis 4. This finding is also illustrated in Figure 13, in which petitions are separated into those with global peaks on day 1, 2, and so on; and we observe, the mean numbers of signatures seems to broadly increase as the peak gets later, which agrees with the regression findings.

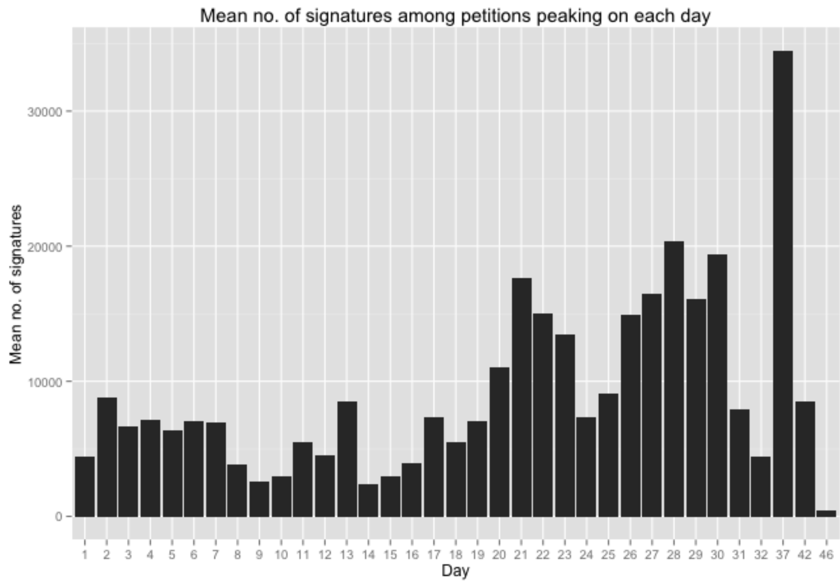

Figure 13: Petitions which peak on later days tend to end up with higher number of signatures than those which peak on earlier days.

Finally, we observe positive and also significant coefficients for the number of local peaks; that is, more successful petitions have more local peaks. However, further analysis reveals that this phenomenon occurs mainly because less successful petitions receive very few signatures after day 30 (as discussed in our threshold section) particularly when they do not reach the $100,000$ mark, and hence have fewer local peaks. Indeed, when we consider only days 1 to 30 and perform another regression of $\log(total)$ against num_local_peaks, the regression coefficient becomes much smaller and no longer statistically significant (coefficient of $-0.009, p = 0.437$).

In summary, there are highly significant correlations between the shape of a petition's temporal distribution and the number of signatures it ends up with. High-signature petitions tend to have:

- a later global peak;
- longer right tails;
- a less peaked distribution;

confirming Hypotheses 3 and 4.

All of these findings are in fact fairly intuitive: a common pattern for petitions that do poorly is that they receive a sharp spike of petitions early on (e.g. in days 1 or 2), but then lose momentum very quickly. Such petitions would have few signatures while having an early global peak, short right tails and a highly peaked distribution, which serves to explain the relationships we observe. For example, in Figure 12, the fourth petition from the left exemplifies this pattern, and the first and second petitions from the left do as well albeit to a lesser extent.

5.5.3 Predictive measures of petition success

So far, we have observed a number of differences in the shape of temporal distributions between more and less successful petitions, and we have hypothesized that petitions that do poorly are characterized by early spikes. This finding can be used as a predictive measure of petition success: we would expect to observe a drop in momentum for unsuccessful petitions fairly early on, for example, even on day 1 or 2. To test if this is true, we construct a new variable, indicating whether each petition had more signatures in its first day or second day.

Among the 59 successful petitions, 68% had more signatures in their second day than their first day, but among the 3623 unsuccessful petitions, this percentage was only 38%. Thus, there is a clear correlation between a petition's success and having more signatures in its second than its first day; this relationship is also statistically significant (χ^2 test p-value $< 10^{-5}$).

This finding agrees with our earlier intuition that poorly performing petitions tend to receive an initial burst of signatures but then decay quickly. Moreover, this measure (whether the petition had more signatures in its first or second day) is particularly interesting as it only relies on the first two days, and thus can be used as an early indicator of whether a petition is likely to succeed.

6 Simulations

So far, in our empirical analysis of petition adoption curves in Section 5, we have observed the following threshold effects:

1. Goal-gradient threshold effect: petitions start to receive fewer signatures after reaching the $100,000$ signatures mark.
2. Temporal threshold effect: if a petition becomes 30 days old without receiving $100,000$ signatures, it starts to receive fewer signatures.

We also observed that more successful petitions have:

1. later global peaks;
2. lower kurtosis (i.e. less peaked signature distribution);
3. more right skew (i.e. a larger fraction of signatures occuring in the 'right tail' of the signature distribution).

In this section, we perform simulations with various models. Our two main goals are 1) to replicate the patterns in the adoptions curves (such as in Figures 11 and 12) as accurately as possible, and 2) to investigate which of our empirical findings can be explained by our simulations. We start by using the Bass diffusion model - a standard model used for product adoption - before considering how it can be modified to fit the data better. Finally, we introduce a 'broadcast and viral' model that better accounts for patterns in the data.

6.1 Bass diffusion model

The Bass diffusion model is a standard model for product adoption (Bass 1969). In this model, we have a fixed population. During each time unit, some fraction of the population adopts the product on their own ('innovation'), while other people are influenced by previous adopters ('imitation'). At any time, the probability of a particular person adopting the product through imitation is proportional to the number of existing adopters of the product.

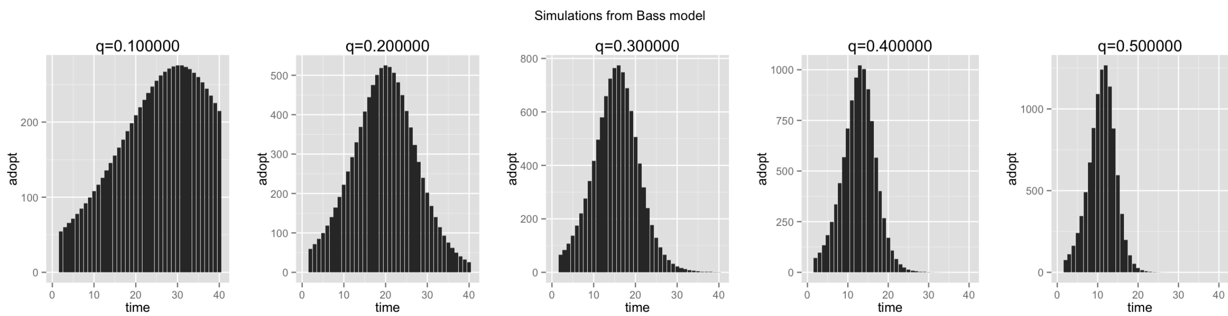

Figure 14: Simulated petitions using the Bass model, for various values of q, the strength of imitation: the higher q is, the stronger the influence that previous adopters have making new users adopt. For all values of q, the Bass model leads to an unrealistically gradual increase and decrease in the adoption curve with at most a single peak.

The Bass model gives rise to an S-shaped curve for the cumulative number of adopters over time; in terms of number of *new* adopters per time step, this is a smoothly increasing then decreasing curve. Figure 14 shows several curves for the number of new adoptions (or signatures) over time, for various values of the parameter for the strength of imitation, q, which affects how strongly previous adopters influence new people to adopt. As Figure 14 shows, for any value of q, the number of new adopters smoothly rises, then peaks, then gradually falls to zero. Indeed, under the Bass model, this is the only pattern possible - effectively, under this model, previous adopters continually exert an influence on new users to adopt, and this influence does not diminish over time, so the product keeps spreading until everyone has adopted it.

As a model for the petition data, therefore, the Bass model is unrealistic - comparing the curves in Figure 14 with the actual petition adoption curves in Figures 11 and 12 we find obvious differences. The patterns in the real data are much more irregular, with often a fairly large number of peaks, not just one, and show much sharper peaks with much less gradual build-up than predicted by the Bass model.

Another problematic feature of the Bass model for our purposes is that once a product (or petition) starts spreading, it cannot 'die out' until the entire population has adopted it. This happens because previous adopters exert a constant pressure on each user to adopt the product; and this pressure does not diminish over time, so the process inevitably continues until the entire population has adopted. This is intuitively unrealistic, as we expect many petitions to lose momentum over time and die out prematurely. Moreover, this property ensures that in simulations, the final number of signatures for a petition is almost exactly equal to the total population in the simulation, making it impossible to study how various parameters related to petition spread affect the final total number of signatures.

6.2 Independent Cascade model

The Independent Cascade model is another model proposed for explaining the spread of a a process or information in a network (Goldenberg et. al. 2001). As a model for petition spread, each person newly signing the petition has a fixed probability of spreading the petition to each person who has not signed, who may then sign the petition at the next time step. Unlike the Bass model, each person only spreads the petition immediately after signing it, rather than continually spreading it at every time points after they sign it. Hence, this makes it possible for petitions to die out before reaching the entire population.

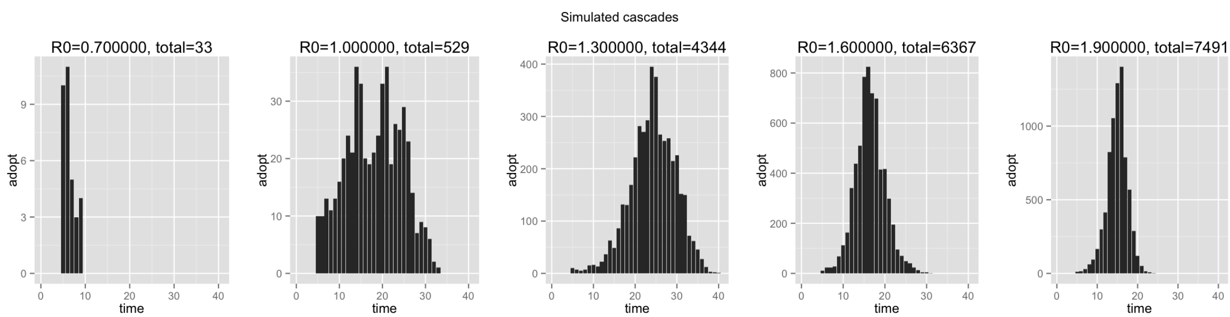

Figure 15: Simulated petitions using the Independent Cascade model, for various values of the basic reproduction number R_0, which is the average number of people that each person signing the petition spreads it to, assuming that no one else has signed yet. When $R_0 < 1$, the petition dies out quickly. *total* is the total number of signatures that the petition ends up with.

Figure 15 shows simulated petitions from the independent cascade model, for different values of the parameter R_0. R_0. The basic reproduction number is the average number of people that a single user spreads the petition to, assuming the rest of the population has not signed the petition yet. When $R_0 < 1$, each person signing the petition spreads it to less than 1 person on average; hence, the petition is very likely to die out before reaching most of the population, which explains the sharply decreasing shape of the leftmost cascade, and the small number of total signatures in this case (33). When R_0 is above 1, each person signing the petition on average brings in 1 or more further signatures, at least initially, when most of the population has not signed the petition. However, as the fraction of people who have already signed the petition grows, each new user has the chance to spread the petition to fewer and fewer subsequent cases, eventually reaching the point where a new user does not replace himself or herself on average, and the petition dies out.

Compared to the actual petition curves in Figures 11 and 12, this model still produces too limited a variety of patterns: when $R_0 < 1$, the adoption curve is a sharply decreasing spike, while when $R_0 \geq 1$ it resembles the Bass model adoption curves, and cannot accomodate the multiple peaks or irregular patterns in Figures 11 and 12. In contrast to the Bass model, however, the cascade model has the more desirable property that the total number of signatures (labeled as 'total' in Figure 15) is higher for higher values of R_0; hence, in contrast to the Bass model, the final total number of signatures is not completely determined by the population size. This happens because the higher the value of R_0, the larger proportion of users have to sign before we reach the point where each user no longer replaces himself or herself.

6.3 Broadcast and viral model

To correct for the independent cascade model not producing the irregular patterns and multiple peaks found in actual data, we next modify the independent cascade model to a 'broadcast and viral' model. This model explains the observed signatures as a mixture of broadcasts (which induces a group of users to sign the petition, e.g. a news broadcast) and viral spreading (spreading of the petition from people who have just signed it to new users).

Under this model, at each time step, a petition has a small probability of a broadcast occurring. Each broadcast brings in a number of users, where the number is drawn from a lognormal distribution, which allows for large variance in broadcast sizes similar to what we observe in actual data. Naturally, this distribution can be replaced by any appropriate distribution in other applications depending on the researcher's prior beliefs.

At the same time, viral spread is happening constantly - similar to in the Independent Cascade model, each user who has just signed the petition has a small probability of spreading the petition to each other user who has not signed it yet. As in the Independent Cascade model, the strength of the viral spread can be parametrized by the basic reproduction number R_0. In addition, since we observe in the real data a fairly constant and low 'background' level at which users sign the petition, we also add a similar low background probability for each user in the population to sign the petition at each time step, independent from the existing broadcast and viral mechanism.

Figure 16 shows simulated petitions from this model for various values of R_0; broadcasts are marked in red vertical dotted lines. For our simulations, the probability of broadcasts is chosen to give an average of 3 broadcasts per petition, with broadcast size following a lognormal distribution: if X is a broadcast size, then $\log X \sim \mathcal{N}(\mu, \sigma^2)$, where we use $\mu = 5, \sigma = 1.5$. There is always at least one broadcast, and the first broadcast occurs on day 1. For viral spread, R_0 is chosen between 0.7 and 1.9 (in fixed intervals for Figure

16, and from a uniform distribution over this interval for the regression in Section 6.3.1). The background level is set such that each user who has not signed the petition has a 0.002 chance of signing it at each time step.

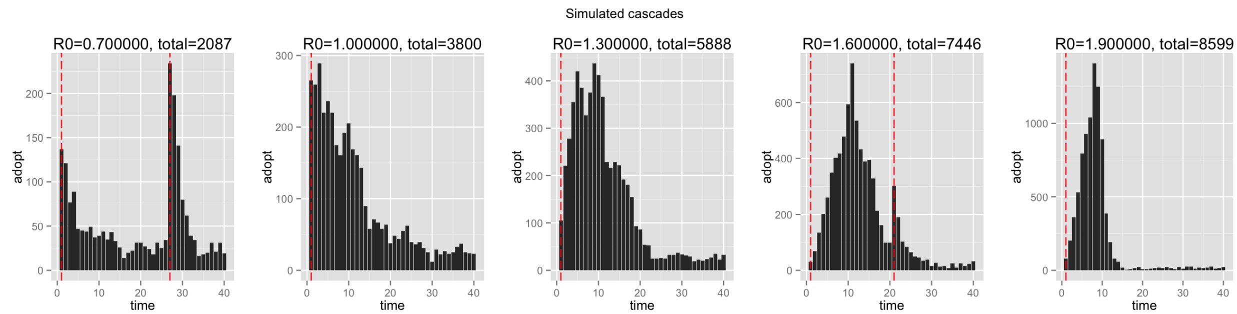

Figure 16: Simulated petitions using the broadcast and diffusions model, for various values of the basic reproduction number R_0. Broadcasts are marked in red dotted lines.

6.3.1 Replicating empirical findings on shape and petition success

A key question we may now consider is the following: are the empirical findings on the correlation between petition shape and petition success also present in the simulations? If they are, this provides a possible explanation for the empirical findings; if they are not, they suggest a way of improving the model.

To answer this question, we simulate 5000 petitions using the broadcast and viral model, and do a linear regression of logarithm of the total number of signatures received by a petition against measures of shape as we did earlier (in Table 2). As before, we use the logarithm of the number of signatures as the response variable to prevent outliers from excessively influencing the fit. Table 3 shows the regression coefficients when using these simulations (left column) compared to the original regression coefficients for the actual data (right column).

Table 3: Comparison between regression coefficients when using simulated data (left column) and actual data (right column). The regression coefficients are all significant and match in terms of sign; however, there is a stronger effect of num_local_peaks for the actual data. Variables: global_peak_day: which day the most signatures were received. num_local_peaks: number of days at which more signatures were received than the previous and next day.

	Dependent variable:	
	log(total)	
	Simulated	Actual Data
	(1)	(2)
global_peak_day	0.007***	0.008**
	p = 0.000	p = 0.013
num_local_peaks	0.024***	0.150***
	p = 0.000	p = 0.000
skewness	0.453***	0.558***
	p = 0.000	p = 0.000
kurtosis	−0.028***	−0.062***
	p = 0.000	p = 0.000
Constant	5.991***	5.455***
	p = 0.000	p = 0.000
Observations	5,000	3,682
R^2	0.298	0.089
Adjusted R^2	0.298	0.088
Residual Std. Error	0.403 (df = 4995)	1.364 (df = 3677)
F Statistic	530.431*** (df = 4; 4995)	90.216*** (df = 4; 3677)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 3 shows that all the coefficients for the regression on simulated data are significant with the same sign as in the original petition data. Most are of fairly similar magnitudes, with the exception of num_local_peaks, which has a stronger effect in the actual data. However, we observed earlier that this variable is significant in the actual data largely as an artifact of the long runs of zeros for less successful petitions.

Since the simulations are based on a simple model, we can explain these regression findings. Petitions with low R_0 peak early when they receive an initial broadcast but then lose momentum extremely quickly due to the lack of strong viral spread; hence, they have earlier global peaks, short right tails, and a highly peaked distribution. Petitions with high R_0 accumulate signatures more gradually due to having stronger viral spread, then lose momentum gradually as the population runs out of users who have not signed the petition; hence having later global peaks, larger right tails, and a less peaked distribution. Since petitions with high R_0 end up with more total signatures, these account for the regression coefficients.

Certainly, this does not necessarily imply that the same effects are present in the actual data. However, our simulations do provide a plausible explanation for the empirical findings that more successful petitions have later global peaks, more skewed and less peaked distributions. In particular, the simulation results suggest that under a simple model of viral spread, as long as different petitions have varying values of R_0 (i.e. rate of viral spread), we should expect correlations between total number of signatures and these measures of petition shape. As we have observed in our simulations, this is because higher R_0 petitions have different characteristic shapes than low R_0 petitions, and also end up with more signatures.

7 Discussion

This paper observes temporal dynamics of adoption and diffusion patterns in online petitions-signing, in order to understand what makes petitions gain traction and growth momentum. In this final section, we return to the motivating questions and hypotheses that motivated this study (as outlined in Section 3) and discuss theoretical and practical implications of our observations.

Our first finding pertains to the growth patterns that characterize online collection mobilizations. Some investigations have identified an S-shaped curve as a model for growth in mobilization (Gonzalez-Bailon et al., 2011), while others find contrary evidence (Hale et al. 2013). The former follows conventional models of offline collective action in positing that the joining rate will be slow in the initial stages of a mobilization until a certain participation threshold is reached (Granovetter, 1978; Centola, 2007). Our findings are more congruent with the latter, since we do not find evidence of S-curve patterns among the petitions on WeThePeople.

The more accurate picture of diffusion patterns appears be an initial growth momentum occurring at the very outset of a petition's introduction, which then decays over time. The growth curve is punctuated by short bursts of growth in support. Such evidence of early growth is commensurate with findings of others that social contagion (which we refer to as "diffusion" in this paper) requires exposure to a diversity of sources (Centola, 2007). Intuitively, it also makes sense that diffusion via online platforms and social networking sites encourage bursts, because a user who does not receive the information about a petition in their social networking platform is unlikely to come into contact with it unless he or she actively seeks it out, which evidence from previous research suggests is atypical. (Lin et al., 2013; Gleeson et al., 2013). All in all, the preponderance of the early growth characteristic highlights the importance of gaining early traction. Experimental research shows that the willingness of individuals to sign a petition varies with the social information provided on how many other individuals have signed the petition already (Margetts, H. 2011; Margetts, H., 2013). The early growth of petitions reflects a similar feedback loop as the petitions with the most signatures get further signatures.

Our second question, relating to Hypotheses 1 and 2, pertains to threshold effects observed in petitions data. Specifically, we identified two thresholds – the temporal threshold at the 30-day mark and the goal of reaching 100,000 signatures required for an official response - and observed effects that differed across petitions with varying degrees of success. The first suggestive evidence of signature thresholds being significant with respect to petitioning activity comes from discontinuities observed in overall adoption patterns. From there, we tested four hypotheses on both temporal and signature thresholds, and observed the varying extent to which these hypotheses hold across successful and unsuccessful petitions.

We confirm that diffusion momentum grows as petitions approach a given temporal threshold through the observation that number of signatures increase at a significantly faster rate right before the 30-day mark, and falls drastically once it crosses the day 30 threshold. This suggests that diffusion dynamics in online petitioning are responsive to thresholds. We further observe variations in responsiveness to temporal thresholds across petitions: highly successful petitions appear to be significantly less sensitive to the 30-day mark. This might be due to the fact that more successful petitions tend to have stronger momentum so support takes a longer time to decay, and that successful petitions are seen to benefit from more signatures (from which it gains a stronger signal as a social call-to-action) even after it has reached its goal. In any case, that the overall signature count of a petition is in no small part a function of temporal thresholds would justify careful consideration of the appropriate time frame for gathering the required number of signatures.

Our observation that the number of signature increases at a faster rate as petitions approach the signature threshold supports the goal-gradient hypothesis, which would predict gains in diffusion momentum as petitions approach a signature threshold. This is congruent with what we would expect to see if goal-gradient behavior is at work, given previous documentation of a similar effect (Kivetz, Urminsky and Zheng, 2006; Cheema and Bahachi, 2011). Why might people invest greater efforts as they approximate a goal? One reason that goal gradient patterns occur is that people judge late-stage events to have greater value than equivalent early-stage events. The ratio of benefit to (remaining) cost increases as one approaches a goal (Koo and Fishbach, 2012). Another reason may have to do with last minute broadcasts by campaigners to tip the petitions over the threshold. A possible test for the latter would be to observe whether broadcast effects are more prominent right before reaching a temporal threshold, so we may better identify whether the last minute gain in momentum is being driven primarily by campaigners or signatories. We would expect the psychological mechanism behind goal-gradient behavior among campaigners and signatories to be quite different, though understanding the precise mechanism would demand further experimental research.

What happens after a goal has been reached? Our findings demonstrate that petitions lose momentum for increasing the number of signatures once the signature threshold has been crossed. It is known in social psychology literature on motivation processes that goals serve as reference points that systematically alter outcomes (Heath et al. 1999) - primarily, they serve as a directive function that directs attention and effort toward goal-relevant activities (Rothkopf and Billington, 1989), and have an energizing function (Bandura and Cervone, 1983; Bryan and Locke, 1967). However, less has been said about what happens *after* a goal has been attained. Our study makes the observation that in the context of online petitions the momentum driving online collective action does not persist beyond the target threshold. Any discussion on adoption and diffusion patterns would have to account for goals and thresholds and their effects on behavior. A more practical implication for a campaigner whose aim is to sustain the momentum of support for their causes is that upward adjustments to goals would be necessary.

The third key question outlined in Section 3 concerns the relationship between temporal diffusion patterns and petition success, as formalized in Hypotheses 3 and 4. In examining systematic correlations between shapes of petitions and their eventual success, we find that more successful petitions tend to exhibit three features: 1) a later global peak; 2) a longer right tail; 3) a less peaked distribution. Petitions that under-perform often experience early bursts of momentum at the outset but the decay of such spikes are usually rapid. Based on our simulations, we find that a simple model combining broadcasts and viral diffusion, in which different petitions have different strength of viral diffusion (or R_0) can account for these three findings, primarily due to the different characteristic shapes between high and low R_0 petitions.

Based on statistical correlates our study also finds an early indicator of a petition's eventual success: a petition is likely to fail if the number of signatures gathered on its second day is lower than its first day. Understanding the causal mechanism driving this relationship will demand further research, but one might venture to say that the first two days are strongly indicative of the reach and effectiveness of a campaigner's initial broadcast. Absent an effective broadcast, it is highly unlikely that a viral effect will set in and bring about the necessary momentum for growth in support.

8 References

Adar, E. and Adamic, L. (2005). Tracking information epidemics in blogspace. In IEEE/WIC/ACM International Conference on Web Intelligence. IEEE Computer Society,Compiegne University of Technology, France

Bandura, A., & Cervone, D. (1986). Differential engagement of self-reactive influences in cognitive motivation. Organizational behavior and human decision processes, 38(1), 92-113.

Bakshy, E. Hosfman, J., W., and Watts, D. (2011). Everyones an influencer: quantifying influence on twitter. In Proceedings of the fourth ACM international conference

Bakshy, E., Karrer, B., and Adamic, L. (2009). Social influence and the diffusion of user-created content. In Proceedings of the tenth ACM conference on Electronic commerce. Association of Computing Machinery, 325334.

Bass. (1969). A new product growth for model consumer durables. Management Science 15, 5, 215227.

Coleman, J., Katz, E., & Menzel, H. (1957). The diffusion of an innovation among physicians. Sociometry, 253-270.

Centola, D., Macy, M. (2007). Complex contagions and the weakness of long ties. American Journal of Sociology 113(3), 702734

Cheema, A., & Bagchi, R. (2011). The effect of goal visualization on goal pursuit: Implications for consumers and managers. Journal of Marketing, 75(2), 109-123.

Cryder, C. E., Loewenstein, G., & Seltman, H. (2013). Goal gradient in helping behavior. Journal of Experimental Social Psychology, 49(6), 1078-1083.

Dodds, P. S., & Watts, D. J. (2005). A generalized model of social and biological contagion. Journal of Theoretical Biology, 232(4), 587-604.

Gladwell, M. 2002. The tipping point: How little things can make a big difference. Little, Brown and Company

Gleeson, J.P., Cellai, D., Onnela, J.-P., Porter, M.A., Reed-Tsochas, F. 2013. A simple generative model of collective online behaviour. arXiv preprint arXiv:1305.7440

Goel, S., Anderson, A., Hofman, J., & Watts, D. (2013). The structural virality of online diffusion. Preprint, 22, 26.

Goel, S., Watts, D. J., & Goldstein, D. G. (2012). The structure of online diffusion networks. In Proceedings of the 13th ACM conference on electronic commerce (pp. 623-638). ACM.

Goldenberg, Jacob, Barak Libai, and Eitan Muller. "Talk of the network: A complex systems look at the underlying process of word-of-mouth." Marketing letters 12.3 (2001): 211-223.

Gonz alez-Bail on, S., Borge-Holthoefer, J., Rivero, A., Moreno, (2011). The dynamics of protest recruitment through an online network. Scientific Reports 1(197)

Granovetter, M. (1978). Threshold models of collective behavior. American Journal of Sociology 83, 6, 14201443

Hale, S. A., John, P., Margetts, H. Z., & Yasseri, T. (2014, August). Investigating Political Participation and Social Information Using Big Data and a Natural Experiment. In APSA 2014 Annual Meeting Paper.

Hale, S. A., Margetts, H., & Yasseri, T. (2013, May). Petition growth and success rates on the UK No. 10 Downing Street website. In Proceedings of the 5th Annual ACM Web Science Conference (pp. 132-138). ACM.

Heath, C., Larrick, R. P., & Wu, G. (1999). Goals as reference points. Cognitive psychology, 38(1), 79-109.

Hull, C. L. (1934). The rat's speed-of-locomotion gradient in the approach to food. Journal of Comparative Psychology, 17(3), 393.

Iyengar, R., Van Den Bulte, C., and Valente, T. W. (2010). Opinion leadership and social contagion in new product diffusion. Marketing Science.

Jones, B. D., & Baumgartner, F. R. (2005). The politics of attention: How government prioritizes problems. University of Chicago Press.

Jungherr, A., & Jrgens, P. (2010). The Political Click: Political Participation through EPetitions in Germany. Policy & Internet, 2(4), 131-165.

Kempe, D., Kleinberg, J., and Tardos, E. (2003). Maximizing the spread of influence through a social network. In 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. Association of Computing Machinery

Kitsak, M., Gallos, L. K., Havlin, S., Liljeros, F., Muchnik, L., Stanley, H. E., & Makse, H. A. (2010). Identification of influential spreaders in complex networks. Nature Physics, 6(11), 888-893.

Kivetz, R., Urminsky, O., & Zheng, Y. (2006). The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention. Journal of Marketing Research, 43(1), 39-58.

Koo, M., & Fishbach, A. (2012). The small-area hypothesis: Effects of progress monitoring on goal adherence. Journal of Consumer Research, 39(3), 493-509.

Leskovec, J., Singh, A. and Kleinberg, J. (2006). Patterns of influence in a recommendation network. Advances in Knowledge Discovery and Data Mining,380389.

Lin, Y.-R., Margolin, D., Keegan, B., Baronchelli, A., Lazer, D.2013. : bigbirds never die: Understanding social dynamics of emergent hashtags Locke, E. A. (1968). Toward a theory of task motivation and incentives.Organizational behavior and human performance, 3(2), 157-189.

Lopez-Pintado, D. and Watts, D. (2008). Social influence, binary decisions and col- lective dynamics. Rationality and Society 20, 4, 399443

Margetts, H.Z, John, P., Escher, T., Reissfelder, S. (2011). Social information and political participation on the Internet: An experiment. European Political Science Review 3(3), 321344. doi:10.1017/S1755773911000129

Margetts, H.Z., John, P., Hale, S.A., Reissfelder, S. (2013): Leadership without leaders? starters and followers in online collective action. Political Studies

Rogers, E. (1962). Diffusion of innovations. Free Press

Rothkopf, E. Z., & Billington, M. J. (1979). Goal-guided learning from text: inferring a descriptive processing model from inspection times and eye movements. Journal of educational psychology, 71(3), 310.

Staab, S., Domingos, P., Golbeck, J., Ding, L., Finin, T., Joshi, A., & Nowak, A. (2005). Social networks applied. Intelligent Systems, IEEE, 20(1), 80-93.

Sun, E., Rosenn, I., Marlow, C., and Lento, T. (2009). Gesundheit! modeling contagion through facebook news feed. In Proc. of International AAAI Conference on Weblogs and Social Media.

Valente, T. W. (1995). Network models of the diffusion of innovations (Vol. 2, No. 2). Cresskill, NJ: Hampton Press.

Yasseri, T., Hale, S. A., & Margetts, H. (2013). Modeling the rise in internet-based petitions. arXiv preprint arXiv:1308.0239.

Young, H.P. (2009). Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. American Economic Review 99, 5, 18991924