Estimating the Dynamics of Individual Opinions in Online Communities

Armin Ashouri Rad, Ph.D. Student, Grado Department of Industrial and Systems Engineering, Virginia Tech, Northern Virginia Center, Falls Church, VA, 22043, USA (ashouri@vt.edu)

Hazhir Rahmandad, Associate Professor, Grado Department of Industrial and Systems Engineering, Virginia Tech, Northern Virginia Center, Falls Church, VA, 22043, USA (hazhir@vt.edu)

Mehdi Yahyanejad, Project Leader, USC Information Sciences Institute, 4676 Admiralty Way, Suite 1001, Marina del Rey, CA 90292 (mehdi@isi.edu)

Abstract

How do opinions change as a result of public interactions and exchange of ideas? How does the proliferation of online media influence these dynamics? While theoretical research provides several hypotheses, empirical analysis of opinion dynamics in online communities is lagging. We develop a unique method for quantifying users' opinions in a social news website and estimate the decision rules that regulate website visit, story posting, voting, and opinion change. We find evidence for significant and nonlinear opinion change as a result of exposure to near-opinions. We also find evidence of learning as people adjust their activity based on the feedback they receive online and strategic reciprocal voting. Incorporating these decision rules in a simulation model we show the propensity of this online community to converge to the majority opinion, and discuss the underlying mechanisms and implications.

1. Introduction

Humans form their opinion by interacting with each other and diverse information sources. Therefore, the evolution of a person's opinion is the result of highly complex feedback process in which one can affect others' opinion and get affected by them. As a result of such dynamics, different trajectories of opinion could emerge in a community (e.g. polarization, plurality, and consensus) (Van Alstyne and Brynjolfsson 2005; Rahmandad and Mahdian 2011)

Several theories have been proposed to explain opinion formation and its dynamics. In general, opinion formation models study the adjustment of individuals' opinion through time based on their interaction with opinions of others. The classical model of opinion dynamics proposed by De Groot (DeGroot 1974) assumes each agent adjust his opinion for the next time step by taking a weighted average on the opinion of others. This model is able to generate consensus or fragmentation of opinion depending on a few conditions. The key feedback mechanism here includes the slow conversion of members to the opinions held by the majority (Figure 1, R1). Friedkin-Johnen model (Friedkin and Johnsen 1990) on the other hand, assumes an individual holds on to his *initial* opinion to a certain degree at any time step while also adjusting towards a weighted average of the community. In the bounded confidence model (Hegselmann and Krause

2002) agents adjust to those opinions that are closer to their own than a threshold. Another theoretical model (Rahmandad and Mahdian 2011) follows a similar idea on the opinion formation and introduces motivation as a separate and relevant component. Individuals' motivation for participating in a community erodes when they are exposed to disagreeable ideas, and that influences the sharing of points of view different from the majority perspective, leading to another reinforcing loop (R2 in Figure 1).

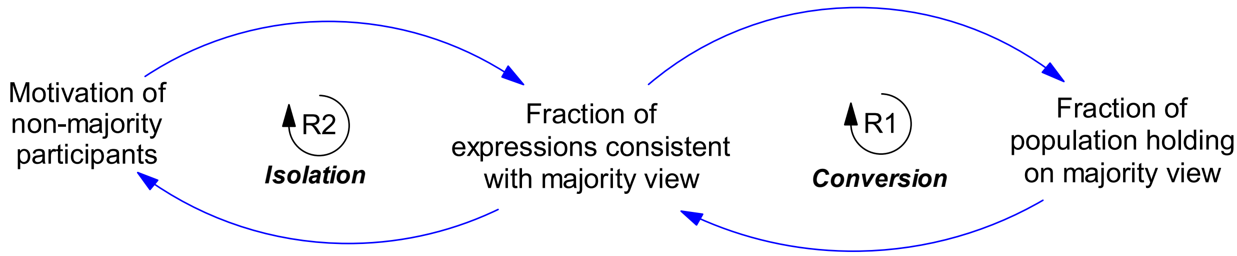

Figure 1. Conversion and motivation feedbacks

Despite these theoretical elaborations, the actual process of opinion change is often very complex, opinions are hard to track, and the influencing factors so many, that empirical assessment of the main contributing factors and the resulting dynamics have not been attempted.

Yet, an increasing fraction of social interactions happen online in social networks, social news websites, forums, and other internet based media through posting, sharing, commenting, and other forms of digital interaction. In such structures, people express themselves and interact using socially generated digital objects such as "like", "vote", "retweet", and "Digg". These online interactions provide us with a valuable source of data on trails of individual choices, and thus opinions, in the online media that can be used to better understand how their opinions change through these interactions.

In this paper, we contribute to this literature by providing a blueprint for empirically modeling the dynamics of online communities. We first propose a novel estimation method to empirically map individuals into an opinion space so that their opinions could be estimated over time. We apply this method to data from a social news website and identify the empirical influence patterns that shape individual opinions over time. We also empirically estimate the decision rules of users in their visiting, posting and voting, the other key decisions that contribute the dynamics of this online community. Finally combining the estimated agent rules with the dynamics of online interactions, we develop a simulation model that projects the opinion dynamics in the study community and uses that to tease out the most likely opinion trajectories possible in this community. By combining simulation modeling and empirical estimation, we bridge a common gap, in which empirical studies of opinion formation only focus on a single aspect of the problem, and the systems level simulation studies are not empirically grounded.

2. Methods and Data

While qualitative coding for assessing opinion is sometimes applied, there is little literature on empirically measuring opinions on a large scale. Opinion space (Faridani, Bitton et al. 2010) and

EU Profiler (Trechsel and Mair 2011) are two tools developed to measure and map opinion on a 2-dimensional space based on answers provided by users to some predefined questions. Both use principal component analysis (PCA) based and reduce the dimension of data collected from users (on a continuous or discontinuous Likert scale) to two dimensions. These methods require original surveys to be administered and therefore time series data collection, which requires multiple measures of the same person over time, is challenging. Here we propose a method for measuring opinions based on matrix factorization and collaborative filtering techniques that could be used on any form of unary, integer or continuous rating data commonly available from online interaction traces of individuals.

In recommendation system literatures, collaborative filtering is defined as “a method of making automatic predictions about the interests of a user by collecting taste information from many users” (Wang 2006). Collaborative filtering starts with a common data structure, in which users and items (e.g. products, movies, stories) are separately identified. It then moves forward based on a simple assumption: users with similar rating on a set of items have close tastes/opinions to each other compared to other users who do not show such similarity in their ratings. In order to find the closeness of opinion, these methods could use matrix factorization methods to form taste vectors for users and items. By estimating the elements of a k-dimensional taste vector assigned to each user and item, matrix factorization collaborative filtering techniques find the taste/opinion values that minimize the difference between the observed and expected ratings. In essence, this procedure estimate both the user, and the item taste vectors so that similarly rated items and corresponding raters (users) have small distance in their taste vectors. As a result, users (and corresponding rated items) with small similarity in rating have larger distance in their taste vector than those with more similarity. Considering that, optimized taste vector of users (acquired from such factorization) could be a representative of users’ position with respect to each other in an opinion space.

We extend the existing matrix factorization methods to the case were ratings are zero-one based (e.g. “like” button in Facebook) and observations are sparse (not every user has seen every story). The details of this method are discussed below, but in short, we use a novel method for identifying the items a user has likely observed (Ashouri-Rad and Rahmandad 2013) and then define a maximum likelihood estimation procedure for estimating user and item opinion vectors. Applying this method on time stamped data (using time based rating matrices as inputs) we are able to map changes in peoples’ opinion through time and study the effect of online interactions on opinion changes. Using our method, we are also able to capture the effect of opinion distance on people’s decision rules in their interaction. Equipped with empirically estimated decision rules we are able to predict the dynamics of opinion change in a community using a well-grounded simulation model.

Next we describe our empirical setting, a social news website (Balatarin.com), and identify the key empirical decision rules used by its users to interact with each other. Then we map users’ opinion on a 2-dimensional space through time and estimate the identified rules based on opinion information and other data we collected from the website. Finally, we discuss the results of simulations that benefit from the statistical estimation work we report in the rest of the paper.

Data

Balatarin is a social news website where users can post stories (i.e. links to different news, websites, videos and other medias; these stories are the "items" we discuss in the methods section), read other users' stories, and vote or comment on the posted stories. Balatarin as the largest social news website for the Persian-speaking community includes over 30000 registered users, a million stories, thirty million votes and several million comments (Bandari, Rahmandad et al. 2013). Popular stories in Balatarin (those with more than a specific threshold) get promoted to its first page and when users visiting the website have the choice of reading first page stories or recently posted stories. We have data only on stories (story ID, posting User's ID, Time of posting, user identified story category (political, economics, sports, social, etc.)) and votes (story ID, time of vote, voting user's ID). Using an algorithm we report elsewhere (Ashouri-Rad and Rahmandad 2013) we rebuild the history of Balatarin in a simulation environment, calculate status of stories and reconstruct users' online behavior over time. Doing that, we are able to collect different forms of data on users and stories, such as the location of each story at any moment, pages users most likely have visited and stories they have read (with some probability) at each point in time, online and offline time of each user, and almost any other key feature of user behavior online. Using this detailed profile for every user from the beginning of this community, we are able to identify and calibrate the empirical decision rules user use in visiting, voting and posting.

3. Analysis and Results

3.1. Mapping users' opinion

We use the collaborative filtering idea to measure and map users' opinion vector on an opinion space through time. We form daily time windows and estimate users and stories opinion vectors on user-story vote data of the day. Doing so, we have enough information to capture the effect of reading stories with similar or diverse opinions on users' opinion change. We assign a 2 dimensional taste vector to each user and story, along with a fixed effect parameter for each user and each story to capture the attractiveness of each published story and the amount of activity each user has on the website regardless of their taste/opinion location. We use the following cost function to factorize the user-story vote matrix to opinion vectors:

$$\text{Argmin}_{U,V}\, F(U,V) = -1 \times \sum_{i=1}^{n} \sum_{j=1}^{m} R_{i,j} \times \log\left(W_{i,j} \times Sigmoid\left(U_{i.}.V_{j.}{}^{T}\right)\right) + \left(1 - R_{i,j}\right) \times \log\left(1 - W_{i,j} + W_{i,j} \times \left(1 - Sigmoid\left(U_{i.}.V_{j.}{}^{T}\right)\right)\right) \quad (1)$$

Where $R_{n\times m}$ is the user-story vote input matrix ($R_{i,j} = 1$ if user i voted for story j and $R_{i,j} = 0$ otherwise), n and m represent the number of users and stories respectively. $W_{n\times m}$ is the exposure probability matrix which is estimated to represent the probability that each user has seen each story. This weight matrix is calculated based on the algorithm we proposed on (Ashouri-Rad and Rahmandad 2013) and details of the algorithm are available in Appendix A and $Sigmoid(x) = 1/(1 + e^{-x})$.

In (1) $U_{i.}$ and $V_{j.}$ are the opinion vectors (for user i and story j respectively) we need to estimate. The cost function defined above has two desirable features. First, it is the log-likelihood function for the observed voting pattern, providing maximum likelihood estimates for U and V vectors. Moreover, the optimization problem is convex, therefore despite the high dimensionality of the problem (with m and n users and stories, we have $k \times m \times n$ parameters to estimate), convergence to the optimum solution is possible with simple gradient descent methods. However, due to the nonlinearity of the problem, with the same input data we may get different optimum values for U and V with the same payoff, all results of rotations and scaling of U and V matrices. In order to track the changes in users' opinions through multiple time windows, we need to have the same scaling and rotation across different time windows. We therefore include a fraction of stories from the previous time window in the next input matrix and fix their taste values to those estimated in the previous window. As a result, the estimated opinion vectors in the current window will follow the rotation and scaling of the previous one, and thus the estimates remain comparable across time.

Using limited memory BFGS optimization algorithm (Liu and Nocedal 1989; Kalhor, Akbarshahi et al. 2013), we calculate opinion vectors of users on a daily basis for a week (longer time horizons are feasible and desirable, but core insights are similar). Figure 2.a shows the opinion values of users on one time window, Figure 2.b represents the change in opinion values for 3 selected users over a one week period (only 100 users mapped on this plot).

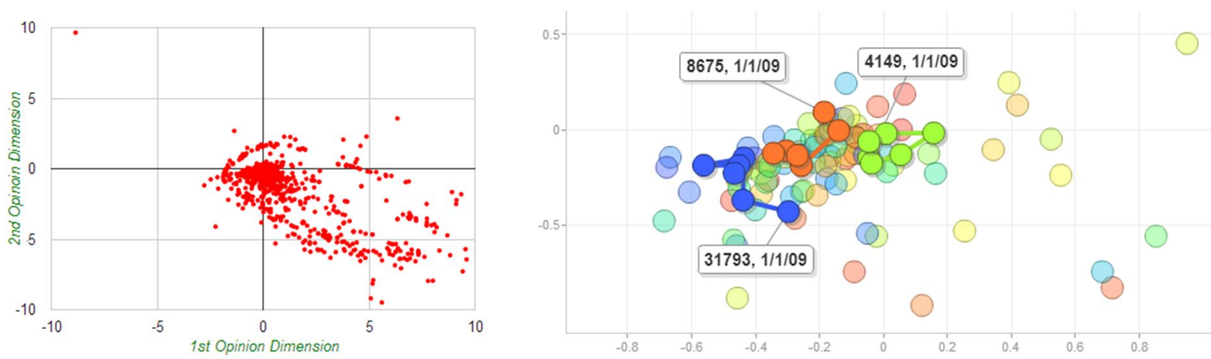

a. Opinion values of users on one time window b. Tracking opinion values for 3 selected users over time

Figure 2. Opinion space

Using optimized opinion vectors, we have the ingredients to estimate how individual opinions interact with story opinions to influence user activity and retention. We also can estimate how opinions change as a result of different factors. In the next part we discuss the statistical models we use to extract these relationships.

3.2. Statistical Estimation

In this section we discuss the key decision rules that shape the basic dynamics of communication and interaction in online communities in general, and Balatarin in particular. These functions identify how users change their opinion due to exposure to different stories, how they decide to show up online, what influences the rate by which they post stories, and how they decide to vote for a story.

3.2.1. How do opinions change?

Using extracted opinion vectors over time, we extract the impact of reading stories (with different embedded opinions) on users' opinion using an adaptation of the bounded confidence model (Hegselmann and Krause 2002). Bounded confidence model proposes that a change in the opinion of an agent may happen due to interactions with other agents "whose opinions differ from his own not more than a certain confidence level". Such model is able to generate different trajectories of opinion in a community (i.e. polarization, plurality, and consensus) using only two parameters: a confidence interval and an opinion dependent bias factor. Using a one dimensional opinion space, bounded confidence model assumes all others in the confidence bound of the agent are equality likely to influence the focal person's opinion. Also, it does not consider any differences among agents in how likely they are to change their opinions (e.g. the heterogeneity of stubbornness).

We model the impact of interacting with other opinions using the same idea of bounded confidence, by studying the change in user's opinion through time caused by reading stories with opinions close to her/his. We estimate this impact by tracking the change in each user's opinion through time (in daily time windows) and predicting that change using the opinions embedded in the stories the user read in the time window. Specifically, consider capturing the impact of reading stories on users' opinions in a 2 dimensional space, where $U_{ik}(t)$ indicates k^{th} dimension of opinion for user i at the time t. $\left(U_{ik}(t-1)-V_{jk}\right)$ is the distance between user and story opinion value for each dimension k. $W_{i,j}$ is the exposure weight we introduced above and T_k is the confidence threshold in the k^{th} dimension. We define a fixed effect, $C_{i,k}$, which captures the changes in each user's opinion during the week of study which is due to factors external to the interactions in the online community. By including this fixed effect we control for the time trends in opinions and stories which are independent of the causal relationship between the two. Finally $\beta_k$represents the impact of reading stories on user's opinion.

$$U_{ik}(t)-U_{ik}(t-1)=$$
$$\beta_k \times \sum_{j=1}^{m}\left\{W_{i,j}\times\left(U_{ik}(t-1)-V_{jk}\right)\middle| for\ j\epsilon\left\{\left|U_{ik}(t-1)-V_{jk}\right|<T_k\right\}\right\}+C_{i,k}\quad k=1,2 \quad (2)$$

The regression above can be estimated using any values of the confidence bound size T, which determines the stories that influence the focal actor's opinion. We select this parameter by running multiple regressions with different values of confidence bound (T_1 and T_2) and finding the combination that provides the best fit based on R-squares. The values 0.25 and 0.5 were chosen as the best fit for T_1 and T_2 respectively, and these values were used in reporting the results. Table 1 provides regression coefficients for the impact of reading on user's opinion in each dimension (R-squared: 0.8176, number of data points: 11507 on 1st dimension and R-squared: 0.8706, number of data points: 12359 on 2nd dimension).

	Estimated Value (Standard Deviation)	p-value
β_1	0.09271 (0.02237)	3.45E-05
β_2	0.03342 (0.01002)	8.59E-04

Table 1. Regression coefficients, standard errors and p-values for variables in model (2)

Both of these estimates are statically significant and the directions of effects are consistent across the two regressions. The magnitudes are different, which is expected because each dimension on the opinion space captures a very different characteristic. Interpreting the coefficients, we can say that when user i read story j ($W_{i,j}=1$) if the taste value of story on dimension k (V_{jk}) falls in user's confidence bound ($\left|U_{ik}-V_{jk}\right|<T_k$), story changes user's opinion (on dimension k) as

β_k times its distance from user's opinion value on that dimension ($\beta_k \times (U_{ik} - V_{jk})$). These effects are rather strong but not unrealistic: they require dozens of slightly different (but not far-fetched) stories to change the opinion of an individual in each dimension.

3.2.2. How often users visit the site?

Users' activity in any online community is directly related to how frequently they go online. Number of times each user visits Balatarin is thus another decision rules we are interested in estimating. In analyzing this measure we build on theories of learning in repeated choice to specify the factors that may influence user visiting pattern and estimate these effects statistically.

In summary, we develop a Poisson regression model to capture the impact of different parameters on daily rate of visiting the site. Setting the number of times each user visits the site in each day as the dependent variable, we include following independent variables in our regression: time spent online in the previous day (to capture both habituation and saturation effects), proximity of opinion between user and stories s/he read (effect of opinion feedback on motivation to visit) in the previous day, average number of votes each of his posted stories receive (effect of reinforcement from peers), and a dummy variable for days of week (to capture weekly cyclical patterns). Details of these factors follow.

A fundamental finding of learning literature is the law of effects: actions that are reinforced by some reward are more likely to be repeated. One such reward in an online community is reading stories that are consistent with one's opinions. Proximity of opinion between user and stories (calculated by equation (4) for each story-user pair) combines the exposure weight (whether the user has read the story) with how favorable the user's opinion is towards the story. Larger (closer to one) values for this function indicate a user who has a strong opinion on some topic and reads stories that support that strong opinion. Values close to zero suggests the stories the user has read are inconsistent with her opinion. Summing over all the stories in the period, this measure provides an estimate of the positive feedback users received because of reading stories that are consistent with their opinions. This measure is in fact close to the total number of stories that the user positively votes for but more precisely measures the impact of reading attractive stories even if they are not voted for. We hypothesize that users will be prone to enjoy reading such stories and thus be more likely to visit the site.

Another rewarding reinforcement for social news website users is when individuals post a story and receive positive feedback (votes) on that story from other community users. We capture this effect using the average of votes received on stories posted by a user in the previous time window. Online time variable is the amount of time user is active on Balatarin in each day and controls for the fact that the number of visits online may change due to spending more/less time online the previous day. Day of the week has a significant effect on users activity in the website (higher activity in weekend compared to working days), which we take into account using dummy variables. We also add fixed effects (C_i) for users to control for variations in individual motivation and availability. Equation (3) specifies the regression model formally:

$$Ln(Rate\ of\ getting\ online\ |X)$$
$$= \beta_0 + C_i + \beta_{1-6} \times Day\ of\ week + \beta_7 \times Online\ time + \beta_8$$
$$\times Total\ proximity\ of opinion\ between\ user\ and\ stories\ he\ read + \beta_9$$
$$\times Average\ number\ of\ votes\ user's\ posted\ stories\ received \quad (3)$$

$$Proximity\ of\ user\ i\ to\ story\ j = W_{i,j} \times Sigmoid\left(U_{i.}.V_{j.}{}^{T}\right) \quad (4)$$

Table 2 provides regression coefficients for model (3) (Conditional R-squared: 0.4014092, number of data points: 8531):

	Estimated Value (Standard Deviation)	p-value
β_0	0.6784 (0.02797)	0.00
β_7	1.78E-05 (1.58E-06)	0.00
β_8	1.26E-04 (4.89E-05)	0.00993
β_9	3.16E-04 (5.55E-04)	0.56928

Table 2. Regression coefficients, standard errors and p-values for variables in model (3)

All coefficients except the one for "average number of votes user's posted stories received" are statically significant. Days of week dummy variables are all statistically significant. In short, individuals are sensitive to the feedback they receive from reading other stories, which increases their likelihood to go online.

3.2.3. How often individuals post stories?

Number of stories users post per day is another decision rules critical to the working of an online community. Building on similar learning arguments as in the previous estimation, we use another Poisson regression to capture the impact of following independent variables on that rate: proximity of opinion between user and stories (equation 4), average number of votes each of posted stories receive, and number of posted stories in the previous time window. Again we add fixed effects (C_i) on users to expel the control for the effect of individual motivation, and dummy variables for days of week. The following equation represents the formal regression model:

$$Ln(Number\ of\ published\ stories\ by\ user\ i\ in\ week|X)$$
$$= \beta_0 + C_i + \beta_{1-6} \times day\ of\ week + \beta_7$$
$$\times Number\ of\ posted\ stories\ in\ the\ previous\ time\ window\ + \beta_8$$
$$\times Total\ proximity\ of opinion\ between\ user\ and\ stories\ he\ read + \ \beta_9$$
$$\times Average\ number\ of\ votes\ user' sposted\ stories\ received \quad (5)$$

Table 3 provides regression coefficients for model (Conditional R-squared: 0.5567877, number of data points: 3767):

	Estimated Value (Standard Deviation)	p-value
β_0	-3.325 (0.08663)	0.00
β_7	-0.068 (0.01555)	1.11E-05
β_8	0.00138 (14E-4)	0.00
β_9	0.0011 (0.0010)	0.2861

Table 3. Regression coefficients, standard errors and p-values for variables in model (5)

Again all the variables except "average number of votes user's posted stories received" are statically significant. Days of week dummy variables are all statistically significant. The results point to a saturation effect (posting many stories in the previous day is likely to reduce the stories an individual will post today) and reading stories that are consistent with an individual's opinion increase the chances of posting more stories.

3.2.4. How do users vote?

Different factors affect users' votes to stories. Given the procedure we used for estimating opinions, it is natural to expect the proximity of user's opinion to a story to be an important explanatory variable for explaining the voting patterns (proximity is calculated using equation 4). Other factors may also play a role. First, much research in diffusion and social influence suggest that the popularity of an item is likely to increase its chances of getting a vote because stories that have already received many votes signal quality and social support. Stories may be from categories inherently more attractive for a user. There is also potential for strategic voting, in which users vote for stories posted by others in a reciprocal tit for tat fashion. Finally, stories location on the website pages could influence their attractiveness. Using the information we obtained from mapping users and stories opinion and reconstruction of user navigation patterns we are able to measure the impact each of these factors on users' voting behavior. We estimate a logistic regression model that predicts the probability of voting for each story if it is exposed to a user ($W_{i,j}$is more than zero (or assumed weight for missing values)):

$$Ln\left(\frac{\pi(x) = Probability\ of\ user\ i\ votes\ for\ story\ j\ published\ by\ user\ l}{1-\pi(x)}\right)$$
$$= \alpha_0 + \alpha_1 \times Number\ of\ votes\ the\ story\ has + \alpha_2$$
$$\times Portion\ of\ user\ i\ stories\ voted\ by\ user\ l\ (publisher\ of\ story\ j) + \alpha_3$$
$$\times If\ the\ story\ belongs\ to\ user's\ (i) favorite\ category + \alpha_4$$
$$\times Place\ of\ story\ in\ the\ website\ + \alpha_5$$
$$\times Proximity\ of\ opinion\ between\ user\ i\ and\ story\ j \qquad (6)$$

Number of votes is a good proxy for stories' popularity at the time that the story is exposed to the user. Place of a story is measured based on which page and row it is placed in and is a good indicator of its visibility. In Balatarin each story is assigned to a category (such as political, social, sport, and art) and since users have different interests and concerns, category of the story is a potentially relevant factor for user's voting. We choose the category that contains the user's most voted stories in the past as her favorite, and use a dummy variable (equals to 1 if the story belongs to that category and 0 otherwise) to capture mentioned impact. There is a subtle complication in using the proximity measures for assessing the closeness of story in the opinion space for the current regression. Specifically, there is circularity in the logic if the opinion values are estimated using the same votes that we attempt to predict in the current regression. To circumvent this problem we first set-aside 20% of user-story vote data (and set the corresponding exposure weights to a small value (0.05)). We estimate the opinion space measures using the remaining data, and then estimate the logistic regression using the 20% of data not used in the estimation of opinion space. Table 4 provides the result (McFadden's R-squared: 0.250039, number of data points: 31372)[1] :

	Estimated Value (standard deviation)	p-value
α_0	-2.26 (0.0137)	0.00
α_1	9.98E-04 (2.33E-04)	1.86E-05
α_2	2.06 (0.0457)	0.00
α_3	0.695 (0.0163)	0.00
α_4	-1.63E-03 (4.50E-05)	0.00
α_5	3.39 (0.0288)	0.00

Table 4. Regression coefficients, standard errors and p-values for variables in model (6)

Fixed effects are embedded in individual and story opinion vectors and thus do not need to be separately captured. All of the effects are statistically significant. Not only the proximity measure is very significant, but also individuals are more likely to vote for stories that already have more votes. Users have a very strong tendency to engage in reciprocal voting. Favorite categories receive significantly more votes, and stories better placed on the page (lower page number and towards the top of the page) are more likely to receive votes.

4. Simulating an Empirically Estimated Online Community

Equipped with estimates for the key actions users take online, we now turn our attention to building and simulating an agent-based model that is grounded in Balatarin data. Specifically, we consider users of the website who visit the site, post stories, read and vote for stories, and change their opinions due to reading different types of stories. Over time opinions and the motivations of individuals change, leading to different levels of activity by different users.

Another key feature of the simulation model replicates the core feature of social news websites, in which the site promotes popular stories (those with votes more than a specified threshold) to its first page. Promoted stories stay on the website for five days but those that couldn't make it to the first page will be removed after only one day. Therefore, we define two states for stories in the model (New Page and First Page) with different life cycles (Figure 3.a). User visiting the website could choose between reading and voting on first page stories or the recently posted stories (New Page). Figure 3.b shows the different states and transition defined for users.

[1] Due to collinearity, we residualized proximity over other variables before using it to fit the model.

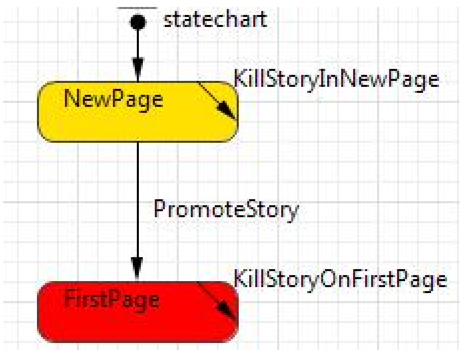

a. Stories' state charts and transitions

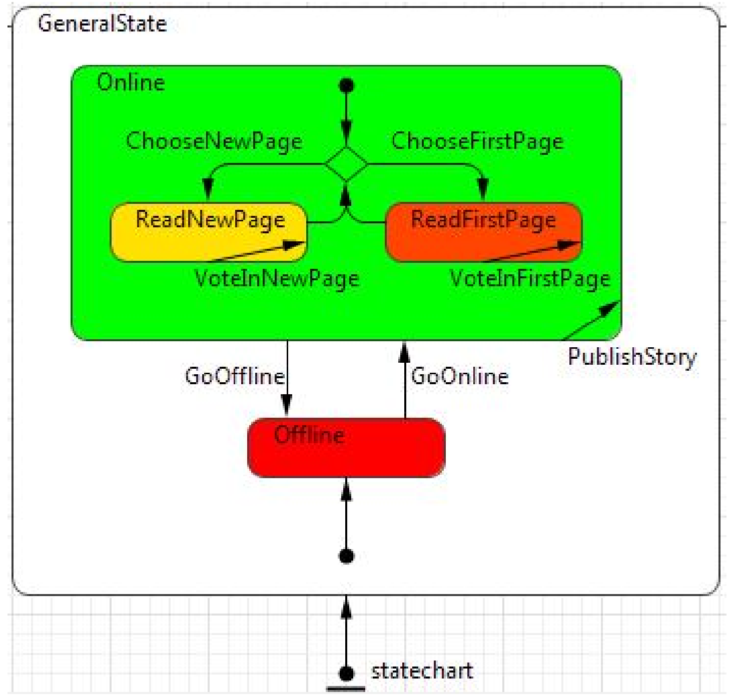

b. Users' state charts and transitions

Figure 3. Simulation states and transitions

We simulate a population of 1500 users (close to the core group of active users on Balatarin) and do not consider new entries into the system for simplicity. Incidentally, due to technical issues that Balatarin faced which lead to stopping the acceptance of new users for a long time, this scenario is close to Balatarin's experience, but different entry rates could be easily incorporated. We consider a two dimensional opinion space which is consistent with our statistical analysis. Besides the decision rules specified in the previous section, several other parameters are directly estimated from the data or from the regression results and residual distributions. These are summarized in Table 5 and include the distribution of users' opinion and their activity, diversity of story opinion from its publisher, distribution of stories' attractiveness, average time users spend on the website in each online session, rate of reading stories, probability of reading first page stories compared to recently published ones, and number of active users per day. We then run the simulation to project the opinion dynamics in Balatarin based on the empirically estimated decision rules.

Parameter	Estimated value
Distribution of users' opinion on 1st dimension	Normal(mean=0.26,sd= 2.23)
Distribution of users' opinion on 2nd dimension	Normal(mean=-0.24,sd=2.38)
Distribution of users' activity	Normal(mean=3.24,sd=5.53)
Deviation of story from its publisher's opinion in 1st dimension	Normal(mean=0,sd=1.15)
Deviation of story from its publisher's opinion in 2nd dimension	Normal(mean=0,sd=1.36)
Distribution of stories' attractiveness	Normal(mean=-6.71,sd=1.23)
Average time of online session	34 minutes
Rate of reading stories	44/hour
Probability of reading first page	41%

Table 5. Parameter values estimated from data

Figure 4 presents the timeline of changes in agents' opinion in each dimension for a single simulation run over 10 years. The individual metrics are color coded based on initial opinions of agents in each dimension separately.

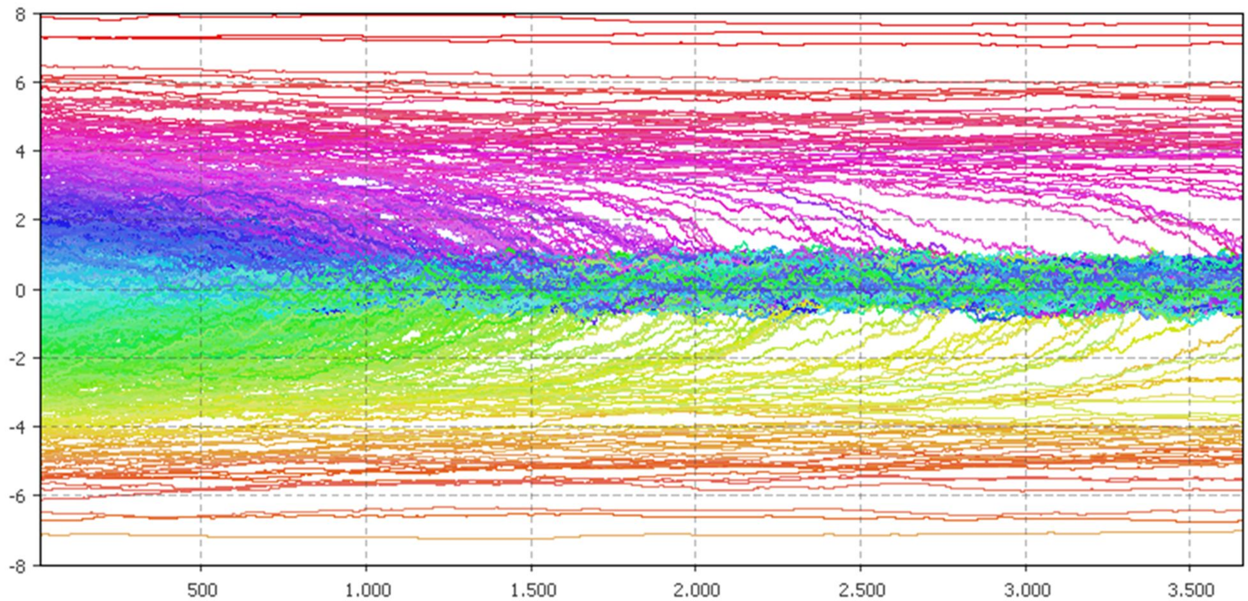

a. Distribution of agents' opinion on 1st dimension over 3650 days (10 years)

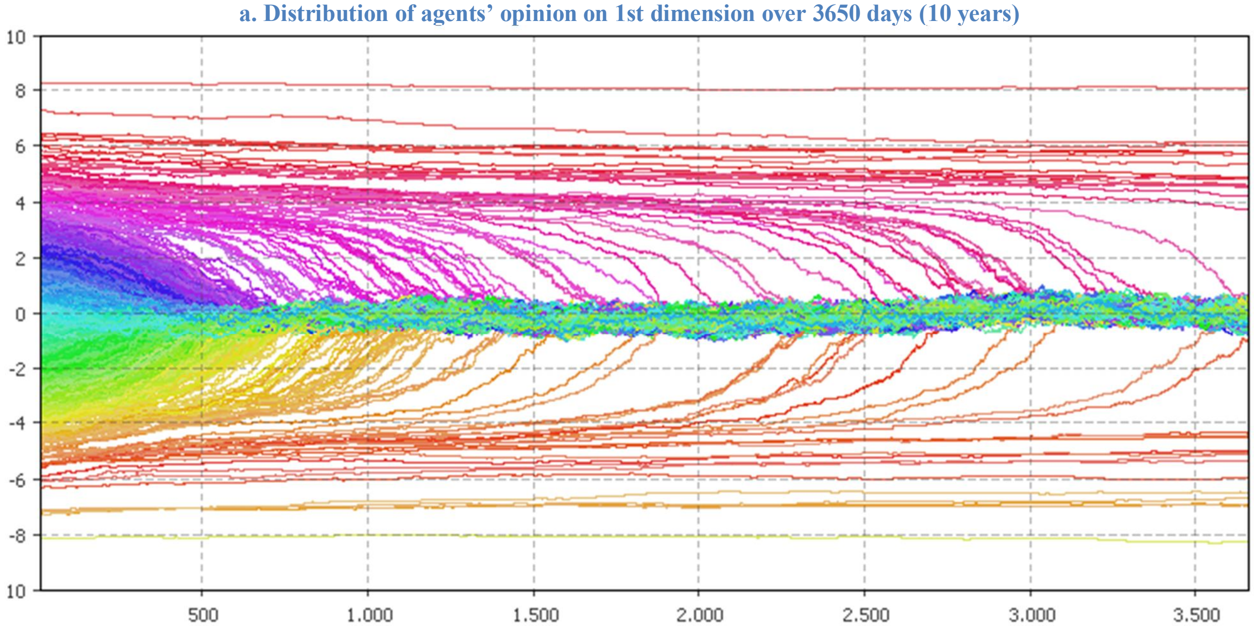

b. Distribution of agents' opinion on 2nd dimension over 3650 days (10 years)

Figure 4. Distribution of agents' opinion

Simulation results show majority consensus shapes in the 1st dimension after nearly 5 years and in 2nd dimension after only 2 years. In both opinion dimensions people close to the majority opinion (with middle-ground initial opinions) converge quickly as they post stories that support each others' opinion and reinforce a single majority position. The convergence process includes two stages, early on the individuals change their opinion only very slowly because the stories in their confidence bound are not many. As they get closer to the population mean, they are exposed to an increasing pressure towards conformity, and thus their speed of convergence increases. The convergence speed ultimately goes down due to the incremental opinion updating process. Those with minority views (at very low or high levels of initial opinion) may take a very long time to converge to the majority. For one thing, exposed to contrary items, their motivation is likely to go down leading to fewer visits to the website and thus a reduction in their exposure. Moreover, being in the extreme positions, they are likely not to read many stories that they find credible enough to change their opinion (i.e. few stories in their confidence bound). In fact a handful of users will not adjust their opinion to the majority during the 10 years of simulation as they find very few credible stories and are generally inactive.

We include the possibility of users posting stories not embedding the exact opinion as the individual holds. This is different from previous theoretical models (Friedkin and Johnsen 1990; Hegselmann and Krause 2002; Van Alstyne and Brynjolfsson 2005) but is motivated by a couple of observations. First, in many cases individuals may not find stories that they fully agree with. Moreover, they may perceive the inherent opinion of that story different from how others perceive it. In fact we find clear evidence from our estimation of opinion space that confirms this intuition. The resulting deviations are captured in the parameters "Deviation of story from its publisher's opinion" in Table 5 and calculated by empirically comparing the opinions of posted stories and the users who posted them.

The impact of this natural noise is potentially important. On the one hand, the community will never converge to a single opinion because there is always a random noise around the mean community opinion sustained by randomness in expression of opinions. Moreover, no user can sustain an opinion very different from the core group for a very long time horizon, because even the extremely distant individuals are infrequently exposed, by chance, to a few stories that are not far from their opinion but are biased towards the community consensus. These perspective over the long run change the distant users' opinions towards the consensus, albeit very slowly, until they get close enough that the majority falls close to their confidence bound and then they are quickly attracted to the consensus group.

In fact, a similar mechanism makes it harder to observe the emergence of competitively polarized communities, where the members converge to two or more smaller sub-communities. The randomness in posted stories' opinions creates a mechanism for communication between these sub-communities and triggers a slow migration from smaller to larger communities. Figure 5 shows this effect: if there is no deviation in the opinions expressed in the stories, compared to those of the posting user, the community fragments to many small sub-communities, each with slightly different, but fully homogenous opinions. In the absence of randomness in expressed opinions, individuals converge to the exact same opinion values and thus stop hearing from other sub-communities which are outside of groups' confidence bound. Disconnected from each other, each sub-community in this hypothetical case will continue to post and be impacted by stories that are fully consistent with their opinion and ignoring the rest of the community.

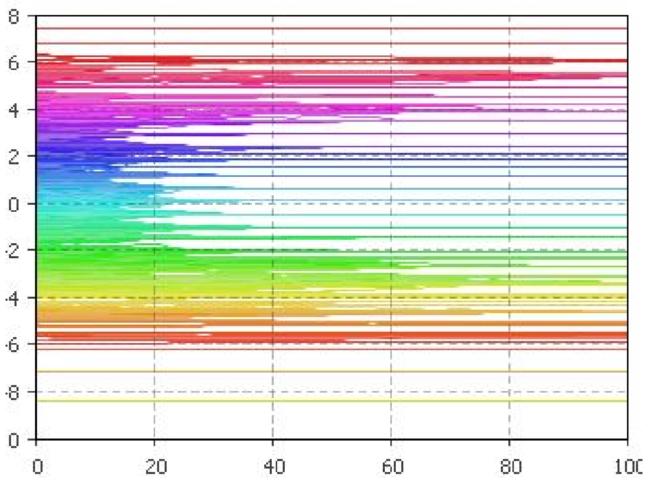

a. Distribution of agents' opinion on 1st dimension over 100 days without deviation of stories

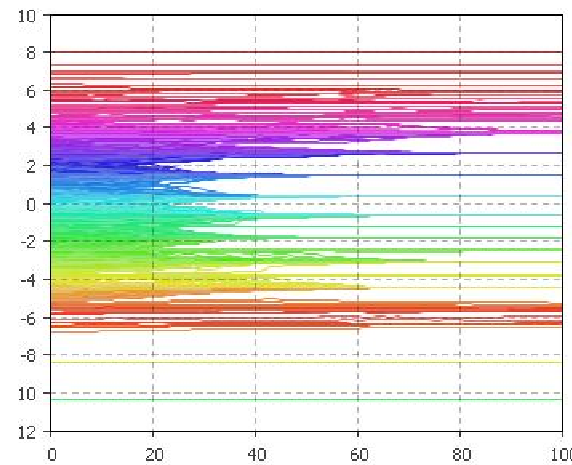

b. Distribution of agents' opinion on 2nd dimension over 100 days without deviation of stories

Figure 5. Distribution of agents' opinion without deviation of stories

In fact, within the parameter ranges specified for motivation, posting, and other effects, the model is quite sensitive to the deviation distribution, and is able to generate different trajectories by only altering the standard deviation value. Figure 6 show different formation including polarization (6-a), plurality (6-b), diversity (6-c), and full consensus (6-d) predicted by the model with different values of standard deviation in deviation distribution. Such sensitivity shows the importance of misunderstanding in interaction on future opinion formation.

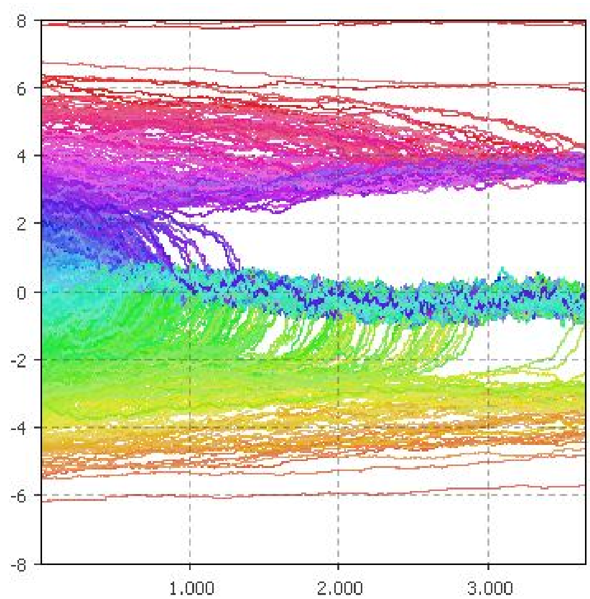

a. Polarity (1st dimension SD=0.65)

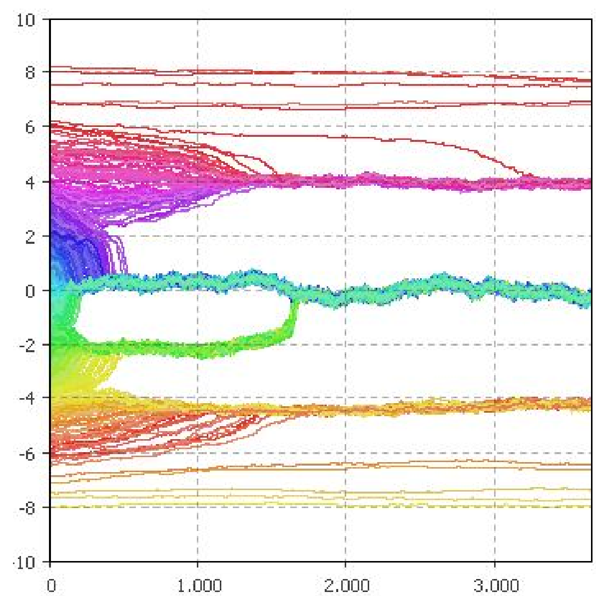

b. Plurality (2nd dimension SD=0.55)

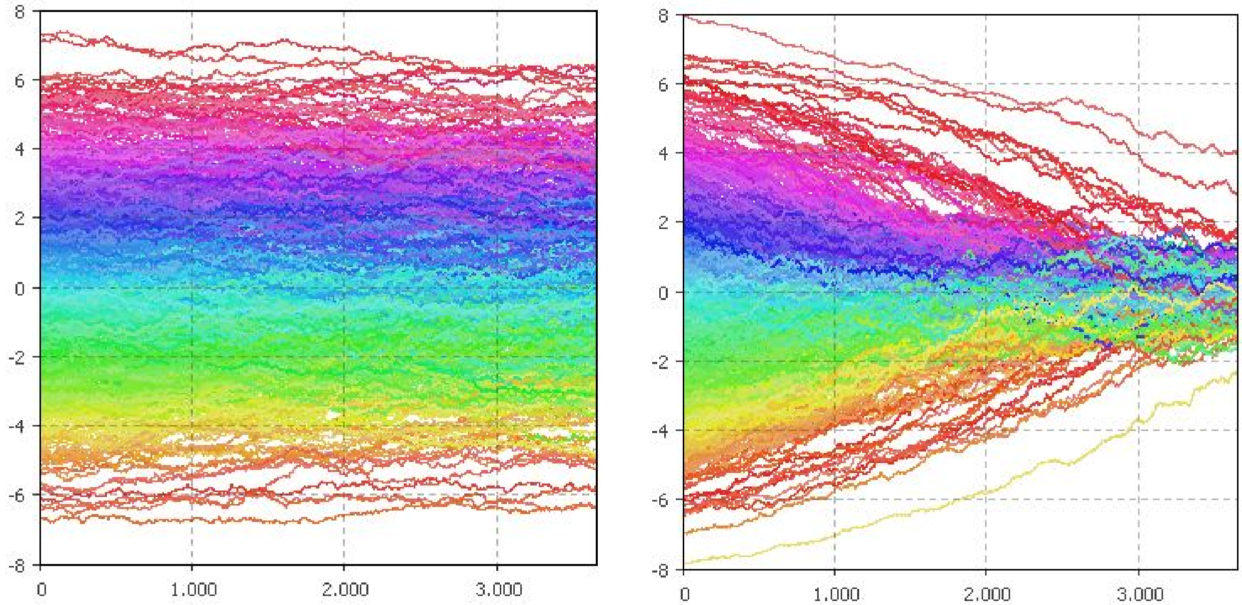

c. Diversity (1st dimension SD=5) d. Consensus (2nd dimension SD=3)

Figure 6. Different opinion formation generated by deviation distribution

5. Discussion and conclusion

The current paper combines statistical estimation and dynamic modeling to better understand the dynamics of online communities. Specifically, we provided an automated method for extracting user opinions in online communities based on their interaction pattern, used that method to estimate the online opinion changes, and showed how individual opinions change as a result of exposure to stories not too far from them. We also estimated the underlying decision rules that guide individual participation in online communities including visiting the website, posting stories, and voting for stories. Using these findings we then built an agent-based model that is fully specified based on realistic parameter ranges. Analysis of this model showed that the dominant mode of operation in the social news website we analyzed is majority domination, where the majority of participants converge to a single region of the opinion space, and the outliers become relatively inactive in the system. We also show that other modes of behavior, including competitively polarized, diversity, and complete consensus are also feasible. An important parameter in determining the dominant mode of behavior is the distance between the posted stories embedded opinion and the opinion held by the individual who posts the story.

The core feedback mechanisms we empirically estimate are simple, but powerful. First, individuals change their opinions as they consume media. In turn, they produce media for the consumption of others (in our case through posting stories they find from the internet) which is closely related to their own opinion. As a result the majority tends to convert more of the community members, further reducing the diversity of opinions expressed in this community. A second mechanism relates to erosion of motivation among the outliers of the community. Observing few attractive stories and getting little support, these individuals are more likely to leave the community, which will further reduce the heterogeneity of the opinions expressed. Qualitative evidence in support of these mechanisms are abound, however, measuring them empirically has not been simple. Empirically we find the first feedback to be pretty strong: individual opinions change in time scales that can be measured within one week of exposure to different stories. The change in motivation was also significant, but less so, and more likely unfolds over months rather than days. We expect to extend the analysis to a longer time horizon that captures these longer-term effects.

From the emergence of cultural norms to evolution of public opinion, the reinforcement of prejudices and the construction of tolerance, opinion dynamics are central to many questions researchers and policy makers care about. With the increasing proliferation of digital interactions, new questions and opportunities for research emerge. Mapping the opinions in online communities using the readily available digital traces of user actions provides a unique opportunity to empirically assess how people change their opinions, how they express themselves, and how their motivations change. One can map the path to formation of new norms in a community, and see how opinions which were initially considered extreme manage to become the norm in a large parts of the society. While our analysis did not focus on the role of specific individuals, much heterogeneity may exist in the opinion adjustment speeds, and that may create novel perspectives into the role of opinion leaders and outliers in changing community dynamics.

Current online community structures leverage two alternative designs to prioritize information sharing and avoid overload: Filtering and ranking systems. Ranking systems promote the majority's point of view (like Balatarin.com or Reddit.com) and leads to majority dominant online environment. Filtering methods create a bubble around similar opinions by personalized filtering (e.g. Netflix). This research can also be utilized to design online communities that are better with respect to some social utility measures. For example the filtering of stories presented to individuals in a social news website could be personalized and take into account the impact on the polarization, motivation, and long-term diversity of the community. In fact experiments insides a simulation model that is well calibrated can be used for filtering the best design modifications, which could then be tested empirically and assessed using the metrics and methods we demonstrated. This approach can extend the current filtering and ranking methods significantly and provide opportunities not currently available.

Yet, converging on a single preferred dynamic pattern is hard and likely premature. Studies on opinion formation indicate that each reference mode has its own advantages and pitfalls. For example, consensus decreases the conflict but leads to lack of variation in opinions and encouraged group-think, bias, and discrimination of minority. On the other hand, polarized communities increase the interaction among competing groups and force them to challenge beliefs and assumptions, but increase the conflict which reduces user motivation, derives out the less outspoken members, and is likely not stable over long horizons. The dynamics within a single platform, such as Balatarin, would also interact with the other platforms. For example individuals who are dissatisfied with one platform may leave it for one which is more consistent with their opinions or design preferences. Therefore the sustainability of any new design is also constrained by the economics of the media platform.

6. References

Ashouri-Rad, A. and H. Rahmandad (2013). Reconstructing online behaviors by effort minimization. Social Computing, Behavioral-Cultural Modeling and Prediction, College Park, MD, Springer Berlin Heidelberg

Balali, V., B. Zahraie, et al. (2012). "Integration of ELECTRE III and PROMETHEEII decision making methods with interval approach: Application in selection of appropriate structural systems." Journal of Computing in Civil Engineering.

Bandari, R., H. Rahmandad, et al. (2013). "Blind Men and the Elephant: Detecting Evolving Groups In Social News." ArXiv e-prints.

DeGroot, M. H. (1974). "Reaching a consensus." Journal of the American Statistical Association **69**(345): 118-121.

Faridani, S., E. Bitton, et al. (2010). Opinion space: a scalable tool for browsing online comments. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM.

Friedkin, N. E. and E. C. Johnsen (1990). "Social influence and opinions." Journal of Mathematical Sociology **15**(3-4): 193-206.

Hegselmann, R. and U. Krause (2002). "Opinion dynamics and bounded confidence models, analysis, and simulation." Journal of Artificial Societies and Social Simulation **5**(3).

Jalali, H. R. M. S. and H. Ghoddusi (2013). Industrial and Systems Engineering, Virginia Tech, Falls Church, VA 22043, USA. Simulation Conference (WSC), 2013 Winter, IEEE.

Kalhor, R., H. Akbarshahi, et al. (2013). Multi-Objective Optimization of Axial Crush Performance of Square Metal–Composite Hybrid Tubes. ASME 2013 International Mechanical Engineering Congress and Exposition, American Society of Mechanical Engineers.

Liu, D. C. and J. Nocedal (1989). "On the limited memory BFGS method for large scale optimization." Mathematical programming **45**(1-3): 503-528.

Pan, R., Y. Zhou, et al. (2008). One-class collaborative filtering. Data Mining, 2008. ICDM'08. Eighth IEEE International Conference on, IEEE.

Rahmandad, H. and M. Mahdian (2011). Modeling polarization dynamics in online communities. Proceedings of the 29th International Conference of the System Dynamics Society.

Trechsel, A. H. and P. Mair (2011). "When parties (also) position themselves: An introduction to the EU Profiler." Journal of Information Technology & Politics **8**(1): 1-20.

Van Alstyne, M. and E. Brynjolfsson (2005). "Global village or cyber-balkans? Modeling and measuring the integration of electronic communities." Management Science **51**(6): 851-868.

Wang, J. (2006). Encyclopedia of data warehousing and mining, IGI Global.

7. Appendix A

Exposure weighting formula:

$$W_{i,j} = \frac{d_c \times e^{-\theta_1 {d_p}^2} + d_p \times e^{-\theta_2 {d_c}^2}}{|X_{previous\ voted\ story} - X_{current\ voted\ story}|} \tag{7}$$

$$d_p = |X_{previous\ voted\ story} - X_j|$$

$$d_c = |X_j - X_{current\ voted\ story}|$$

Where $W_{i,j}$ is the corresponding probability of story j being exposed to user i. In our case formula (7) means how likely the user seen the stories in between each pairs he voted on a same page. Here θ_1 and θ_2 define the slope of decrease in weight by getting far from pair of voted stories. We set $\theta_1 = \theta_2 = 0.01$ and consider a small weight (0.05) for missing values. We put $W_{ij} = 1$ for each pair of (i, j) that have $R_{ij} = 1$, (voted stories definitely exposed to the voter). d_p is the distance (number of stories) between story j and previous voted story in that page, and d_c has the same definition for current voted story. There are number of alternative ways to deal with missing data (Pan, Zhou et al. 2008; Balali, Zahraie et al. 2012; Jalali and Ghoddusi 2013), but based on our available data on Balatarin we proposed a behavior reconstruction process that gives us a good estimation on such binary missing data.