How public opinion forms

Fang Wu and Bernardo A. Huberman

Social Computing Lab, HP Labs, Palo Alto, CA 94304

Abstract. No aspect of the massive participation in content creation that the web enables is more evident than in the countless number of opinions, news and product reviews that are constantly posted on the Internet. Given their importance we have analyzed their temporal evolution in a number of scenarios. We have found that while ignorance of previous views leads to a uniform sampling of the range of opinions among a community, exposure of previous opinions to potential reviewers induces a trend following process which leads to the expression of increasingly extreme views. Moreover, when the expression of an opinion is costly and previous views are known, a selection bias softens the extreme views, as people exhibit a tendency to speak out differently from previous opinions. These findings are not only robust but also suggest simple procedures to extract given types of opinions from the population at large.

On reflection, it is rather surprising that people contribute opinions and reviews of topics which have already been extensively covered by others. While posting views is easy to understand when it involves no effort, like clicking on a button of a website, it is more puzzling in situations where it is costly, such as composing a review. If the opportunity to affect the overall opinion or rating diminishes with the number of published ones, why does anyone bother to incur the cost of contributing yet another review? From a rational choice theory point of view, if the utility to be gained does not outweigh the cost, people would refrain from expressing their views. And yet they do. This is reminiscent of the well analyzed voter's paradox [5, 12, 13], where a rational calculation of their success probability at determining the outcome of an election would make people stay home rather than vote, and yet they show up at the polls with high turnout rates (for a review, see [7]). In contrast to a political election, there is no concept of winning in online opinion systems. Rather, by contributing her own opinion to an existing opinion pool, a person affects the average or the distribution of opinions by a marginal amount that diminishes with the size of that pool.

Since user opinions play such an important role in trust building and the creation of consensus about many issues, there have been a number of recent of studies focused on the design, evaluation and utilization of online opinion systems [1, 2, 6, 8] (for a survey, see [3]). It is surprising that with the exception of one study [10], little research has been done on the dynamic aspects of online opinion formation. It remains unclear, for example, whether reviews undergo any systematic changes as time goes on, or whether the opinions about given political or societal views fluctuate a long time before reaching a final consensus. Thus the need to understand how online opinions are created and evolve in time in order to draw accurate conclusions from that data.

In this context we analyzed the dynamics of online opinion expression by analyzing the temporal evolution of a very large set of user views, ranging from 1.8 million online reviews of the 48,000 best selling books at `Amazon.com`, to thousands of political resolutions voted on `Essembly.com` and the many arbitrary opinions offered for voting on `Jyte.com`.

To start with, a forum where no historical data is available should exhibit no polarization of views as they are expressed over time. In order to test this as calibration for our study, we analyzed the votes of 16,660 resolves posted on the website `Essembly.com` from August 2, 2005 to December 12, 2006, among which 14,171 resolves received more than 21 votes. `Essembly.com` is a website that lets its users post and vote on political resolves by selecting one of the four choices: "agree", "lean agree", "lean against", and "against". A user does *not* see the voting results until she submits her own vote. When a user posts a new resolve, she is required to vote on it. The four voting options, from "agree" to "against", are represented by $-1, -0.5, 0.5, 1$ in our database, respectively. When a user posts a new resolve, she tends to formulate it in a tone that sounds positive to her. As a consequence, 96.0% of all first votes are "agree", while only 2.7% of all first votes are "against". To remove this artificial bias, we discarded the first vote of each resolve and replaced all the remaining votes $\{X_n\}$ by $\{-X_n\}$ if $\sum X_n > 0$, where X_n denotes the quantified value of the n'th vote. This way every vote is "agreed" by the majority.[1] This formed our final data set, which consists of 14,171 resolves, each having more than 20 votes.

We observed no clear trend in the expected value of EX_n. For each resolve we calculated a series of average votes: $\bar{X}_1, \ldots, \bar{X}_{20}$, where $\bar{X}_i$ is the average vote of the first i votes. We performed a linear regression of $\bar{X}_n$ over n: $\bar{X}_n \sim kn + b$. The slope k reflects the overall trend of $\bar{X}_i$: if $k > 0$ the votes increase with time, if $k < 0$ they decrease. A histogram of the 14,171 slopes is shown in Fig. 1. A t-test of the null hypothesis "$k = 0$" yields a p-value 0.064, which is not enough to reject the null. This confirms the absence of an overall dynamical trend.

To see the effects of other people's opinions on the overall public opinion formation we studied `Jyte.com`, a website that allows its users to make any claim they wish and let the community vote on it at no cost. The claims are wide ranging, from "Ocean exploration has more potential to benefit the human race than space exploration" to "Homeopathy shouldn't be available on the NHS". The web interface is simple and intuitive. Each claim is flanked by a positive button and a negative button and the numbers of total positive and negative votes are shown on the face of the buttons. Each user sees the numbers, makes up her mind, and submits her vote by clicking on one of the two buttons. Then the numbers get updated instantly.

---

[1] Another way to remove the bias is to randomly flip the sign of all votes for each resolve. The conclusion would be the same. Notice that the lack of normalization would lead to an artificially high initial value that would decay over time.

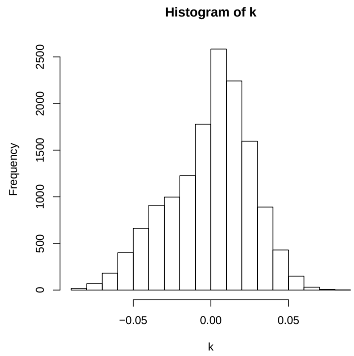

Fig. 1. The histogram of 14,171 slopes calculated from the `Essembly` data. It is not very clear whether most slopes are positive or negative.

We tracked the voting dynamics of all claims made in July 2007, among which 1,208 claims received no less than 10 votes. For each vote, we recorded $X_n = 1$ if the n'th vote is positive, or $X_n = 0$ if it is negative. The quantity $\bar{X}_n = \frac{1}{n}\sum_{i=1}^{n} X_i$ then represents the fraction of positive votes among the first n votes.

We constructed two data subsets from the 1,208 claims. The first set contains all the claims with less than say, 3 positive votes, among the first 10 (i.e. $\bar{X}_{10} < 0.3$). We call it the "negative" set. The second set contains all the claims with more than say, 7 positive votes, among the first 10 (i.e. $\bar{X}_{10} > 0.7$). We call it the "positive" set. These two sets contain the claims that are generally regarded as "very negative" and "very positive", respectively. The negative set consists of 405 claims, and the positive set consists of 521 claims. The average $\bar{X}_n$ for the two data sets are shown in Fig. 2(a) and (b). As we can see from the figure, the negative claims tend to become more negative as the voting goes on, and the positive claims tend to become more positive.[2] This shows that when people observe previous opinions before they express their own, they tend to follow the trend. As a result of this trend following, extreme views get reinforced and become increasingly more extreme.

Our last study focused on the situation where it is costly to express an opinion. We thus considered the online rating data for a large number of books collected from `Amazon.com`. On `Amazon`, a user observes the average rating of a book when she visits a book page (usually shown at the top, right under the title). If she decides to review a book, she is required to write a short paragraph of review in addition to a simple star rating. The average word count of `Amazon` reviews is 181.5 words [9]. Thus, the cost of opinion expression is high for `Amazon`

[2] This is independent of our choice of the values 3 and 7.

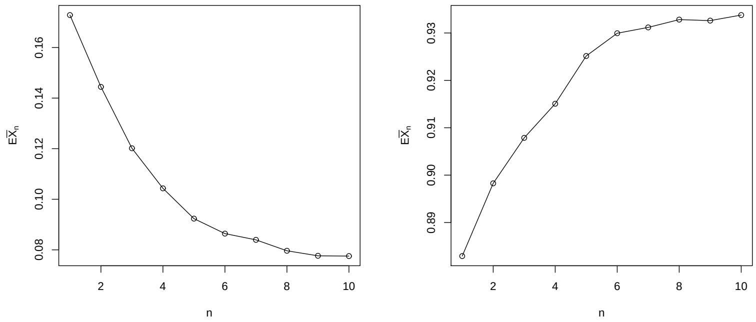

(a) (b)

Fig. 2. The sample average fraction of positive votes $E\bar{X}_n$ as a function of the number of votes n, collected from `Jyte.com`. (a) Negative claims become more negative as time goes on. The average fraction decreases by an absolute amount 9% after 10 votes. (b) Positive claims become more positive as time goes on. The average fraction increases by an absolute amount 5% after 10 votes.

compared with `Essembly` and `Jyte`, and the person will only contribute if the gain from expressing an opinion is higher than the cost.

In the spirit of the voter's paradox, which assumes that the voter is more likely to vote when his vote is more probable to affect the outcome, we speculate that in cases like `Amazon`, people will derive more utility the more they can influence the overall rating. To be precise, in cases where users' opinions can be quantified and aggregated into an average value, the influence of an online opinion can be measured by how much its expression will change the average opinion [11]. Suppose that n users have expressed their opinions, $X_1, \ldots, X_n$, on a given topic at a website, with X_i denoting the quantified value of the i'th opinion. If the $(n+1)$'th person expresses a new opinion X_{n+1}, it will move the average rating to

$$\bar{X}_{n+1} = \frac{n\bar{X}_n + X_{n+1}}{n+1}, \tag{1}$$

and the absolute change in the average rating is given by

$$|\bar{X}_{n+1} - \bar{X}_n| = \frac{|X_{n+1} - \bar{X}_n|}{n+1}. \tag{2}$$

We thus conjecture that a person is more likely to express her opinion whenever $|X_{n+1} - \bar{X}_n|$ is large — an opinion is likely to be expressed if it deviates by a significant amount from those already stated. Indeed, what is the point of leaving another 5-star review after one hundred people have already done so? This point has also been made within the "brag-and-moan" model [4, 8] which assumes that consumers only choose to write reviews when they are very satisfied with the products they purchased (brag), or very disgruntled (moan). Note however, that the brag-and-moan model is static and thus predicts that $\bar{X}_n$ is constant over time, in contradiction with the observed dynamical trends.

Our sample consisted of the top 4,000 best-selling titles in each of the following 12 categories, as of July 1, 2007: arts & photography, biographies & memoirs, history, literature & fiction, mystery & thrillers, reference, religion & spirituality, sports, travel, nonfiction, science, and entertainment. For each of the 48,000 books, a series of ratings was collected in time order, where each rating is an integer between 1 and 5 (number of stars). Among the 48,000 books, 16,454 books have no less than 20 ratings, and 11,920 have an average rating above 4.

We first checked the average rating of the 16,454 books as a function of n. As can be seen from Fig. 3, EX_n decreases almost linearly with n, so there is a clear dynamical trend in the ratings, which corroborates the observation reported in [10]. Later users indeed tend to write different reviews from those of earlier users. As opposed to what we observed in `Jyte.com`, the overall opinion tends to decrease away from the extreme ones.[3]

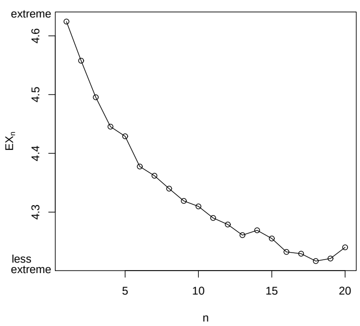

Fig. 3. The average rating of 16,454 books on `Amazon.com` with more than 20 reviews. EX_n is the sample average rating of all the 16,454 n'th ratings. As one can see from the figure, EX_n decreases by 0.4 stars in 20 steps. We did not obtain enough data from low selling books to show the opposite trend.

Next we examined whether this dynamical trend is still prominent at the level of each individual book. Similar to our `Essembly` study, we performed a linear regression of $\bar{X}_n$ over n: $\bar{X}_n \sim kn + b$. The histogram of 16,454 slopes (k) are shown in Fig. 4. As can be seen, most of the slopes are below zero. A t-test of the alternative hypothesis "$k < 0$" yields a p-value less than 2.2×10^{-16}, which further confirms the declining trend.

[3] In a forthcoming paper we show that conditional expectations of `Amazon` ratings (given the average previous ones) follow the same trend.

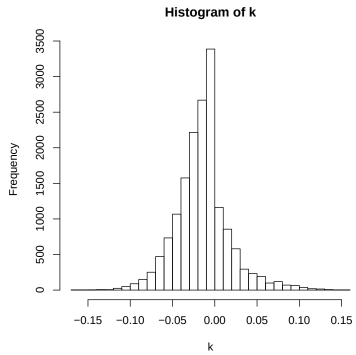

Fig. 4. Histogram of the slopes of average book ratings on `Amazon.com`. Most of the slopes are negative, testifying a declining trend in the average ratings.

Finally we measured directly how much one's rating deviates from the observed average rating. We plot the expected *deviation*

$$Ed_n = E|X_n - \bar{X}_{n-1}|. \tag{3}$$

as a function of n in Fig. 5. As can be seen, Ed_n increases with n. Since the expected deviation Ed_n of an i.i.d. sequence normally *decreases* with n, this increasing trend is indeed significant. This again supports our conjecture that those users who disagree from the public opinion will be more willing to express themselves and thus soften the overall opinion of a given book.

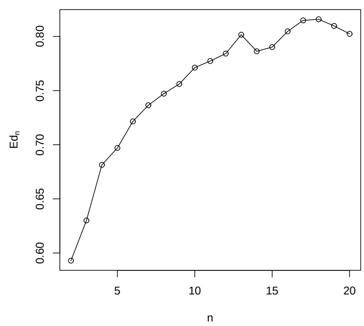

Fig. 5. The average deviation of `Amazon` ratings increases with the number of people.

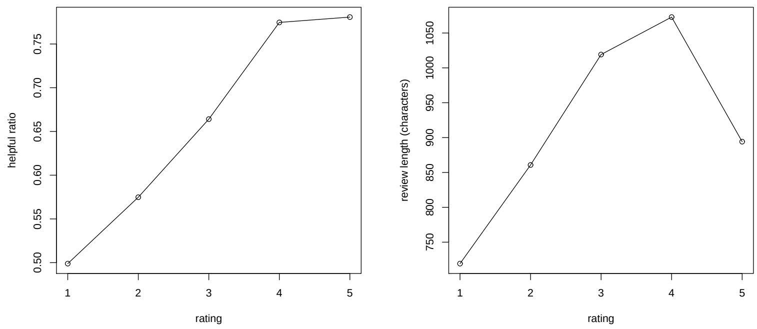

(a) (b)

Fig. 6. (a) The average helpful ratio of five different star ratings. (b) The average review length of five different star ratings in the number of characters. The data is calculated for 4,000 bestselling mystery books. By comparing the two figures it is clear that people find high ratings more helpful not just because they are longer. For instance, 5-star reviews are on average shorter than 4-star and 3-star reviews but are nevertheless more helpful.

One point to be stressed is that the results do not imply that as time goes on the average perception of the book changes. Rather, from a large pool of readers it is only those that want to make a difference with the prevailing opinion that choose to express themselves. This is seen when plotting the average "helpful ratio" as a function of star rating in Fig. 6 for users of `Amazon`. It can be seen that people find high ratings more helpful than low ratings, implying that the majority does not agree with this expression bias.

These results, made possible by the fact that the web presents a natural laboratory to study millions of opinions, show that in the process of expressing their views, people tend to follow different but regular patterns. When no information of previous views is available, the opinions expressed are drawn from a uniform distribution within the community. In cases where previous opinions are made known and it is painless to post a view, one observes either neutral opinions or a polarized consensus which reflect trend following by the group. In the latter case, opinions tend to reinforce previous ones and thus become more extreme. Finally there are many cases where expressing a view is costly, like when writing a book review. In this case people will tend to do so whenever they perceive they can offset the current view by presenting a differing one. Since the impact decreases with the number of posted opinions, the larger the pool, the more extreme the difference expressed. As a consequence one sees a softening of the prevailing view.

Besides explaining the observed data, these results show a cautionary note on the interpretation of public opinion. This is because a simple change in the order or frequency of given sets of views can change the ongoing expression in the community, and thus the perceived collective wisdom that new users will find when accessing that information.

References

1. Dan Cosley, Shyong K. Lam, Istvan Albert, Joseph A. Konstan, and John Riedl. Is seeing believing? How recommender interfaces affect users' opinions. In *CHI 2003*, Ft. Lauderdale, Florida, USA, April 2003.
2. Chrysanthos Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *ACM EC'00*, Minneapolis, Minnesota, USA, October 2000.
3. Chrysanthos Dellarocas. The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49(10):1407–1424, October 2003.
4. Chrysanthos Dellarocas and Ritu Narayan. What motivates consumers to review a product online? A study of the product-specific antecedents of online movie reviews. In *WISE 2006*, Evanston, Illinois, USA, December 2006.
5. Anthony Downs. *An economic theory of democracy.* Harper & Row, New York, 1957.
6. Guodong Gao, Bin Gu, and Mingfeng Lin. The dynamics of online consumer reviews. In *WISE 2006*, Evanston, Illinois, USA, December 2006.
7. Benny Geys. 'Rational' theories of voter turnout: A review. *Political Studies Review*, 4:16–35, 2006.
8. Nan Hu, Paul A. Pavlou, and Jennifer Zhang. Can online reviews reveal a product's true quality? Empirical findings and analytical modeling of online word-of-mouth communication. In *ACM EC'06*, Ann Arbor, Michigan, USA, June 2006.
9. Erik Ketzan. `Amazon.com` and the new democracy of opinion. `http://www.themodernword.com/pynchon/papers_ketzan1.html`, November 2002.
10. Xinxin Li and Lorin M. Hitt. Self selection and information role of online product reviews. In *WISE 2004*, Brisbane, Australia, November 2004.
11. Martin J. Osborne, Jeffrey S. Rosenthal, and Matthew A. Turner. Meetings with costly participation. *American Economic Review*, 90(4):927–943, 2000.
12. William H. Riker and Peter C. Ordeshook. A theory of the calculus of voting. *American Political Science Review*, 62:25–42, 1968.
13. Alexander A. Schuessler. Expressive voting. *Rationality and Society*, 12(1):87–119, 2000.