

Quantifying Social Influence in Epinions

Akshay Patil

Golnaz Ghasemiefteh

Roosbeh Ebrahimi

Jie Gao

Department of Computer Science, Stony Brook University {akshay, gghasemiesfe, rebrahimi, jgao}@cs.stonybrook.edu

Abstract—In many eCommerce websites and consumer review websites, users can review products they purchased as well as the reviews others wrote. Users can also rate each other as trusted or untrusted relationships. By studying a data set from Epinions, we examine and quantify the correlation between trust/distrust relationships among the users and their ratings of the reviews. We discover that there is a strong alignment between the opinions of one’s friends and his/her ratings. Our findings also suggest that there is a strong alignment between the collective opinion of a user’s friends and the formation of his/her future relationships.

I. INTRODUCTION

Many eCommerce platforms such as Amazon and Ebay actively engage their users in providing their opinions or ratings of the products they have bought. Consumer review websites such as Epinions, Yelp, Angie’s List¹, etc allow users to review products and also rate reviews written by other users. These websites have come to play an important role in guiding people’s opinions on products and in many cases also influence people’s decisions in buying or not buying the product. As shown in the global Nielsen survey of 26,486 Internet users in 47 markets, consumer recommendations are the most credible form of advertising among 78% of the study’s respondents [1]. A recommendation from a trusted source goes a long way in shaping others’ opinions.

Some websites incorporated social network structures into their rating systems. Users can rate each other as trusted or distrusted relationships. In such environments, we are dealing with two types of data: “rating system” data and “social structure” data. The incorporation of the social network structure into the rating system leads to an interplay between the two that affects both of them. Users may rate each other as trusted/distrusted based on reviews they write. The social structure might influence users in rating the reviews of others. Furthermore, the reviews in the system and the trust/distrust relationships based on those reviews might influence the formation of new relationships.

To carry out an investigation of social influence we use data from a widely used consumer review website Epinions [2], [3]. Users write reviews on products or services that may earn them money or recognition. Users also rate reviews written by other users and can also express trust or distrust towards other users. While trust relationships are publicly visible, distrust relations are not visible (except to the user making the relationship). When rating reviews, users have an option of making their rating public or keeping it private.

The trust and distrust relationships for a user are combined to determine the *web of trust* for that user; the user is shown reviews written by people within this web of trust more prominently as opposed to other reviews.

In this paper we investigate the existence of a correlation between the opinions of a user’s current trustees/friends and formation of his/her future relationships. Specifically, if user A ’s friends collectively have an opinion (trust/distrust) about user B , would user A ’s future relation with B have a correlation with the collective opinion of his/her friends regarding B ? For example, if more of A ’s friends tend to *trust* B (rather than *distrust* B), would A be more likely to trust B ? In how many cases would A make a decision in contrast to what his friends think of B ? How are the choices made by A related to trustworthiness of both A and B ?² To carry out such an investigation, we must be very careful about issues like innate differences in *trustworthiness* of different users.

Results: Our findings suggest that there is a strong *alignment* between the collective opinion of a user’s friends and the formation of his/her future relationships, after we factor out the innate biases of trustworthiness of different users.

We also investigate the existence of a correlation between the opinions of a user A ’s current trustees/friends regarding another user B and the score of rating that A would assign to reviews written by B . We are specifically interested in the case where A and B do not have a direct trust/distrust relationship but rather an indirect one (A is not friend/foe of B , but one of A ’s friends/foes is a friend/foe of B). We refer to this problem as studying *friend-of-friend dynamics*. If there is no correlation between friend-of-friend dynamics and a user’s ratings, then we can say that the social structure does NOT provide any advantage to the ecology of the rating system. On the other hand, if there exists such a correlation, we could use this finding to recommend better content to the user. It would also imply that the social structure supports or improves the overall user experience by helping him/her identify relevant content. Again, we take into consideration the innate differences in *rating habits* of different users.

Results: Our analysis leads us to conclude that in cases where user A ’s friends have expressed approval or disapproval of another user B , there exists an alignment between the A ’s rating of B ’s reviews and his friends opinions regarding B . On the other hand, approval or disapproval expressed by foes

¹The websites are <http://www.amazon.com>, <http://www.ebay.com>, <http://www.epinions.com/>, <http://www.yelp.com/>, and <http://www.angieslist.com>.

²While Leskovec *et al.* in [4] have looked at triadic closure: i.e. A ’s relationship with B depending on both A ’s and B ’s relationship with an additional nodes X , we take an aggregate view of the relationship formation process; i.e. we intend to understand the collective role played by all of A ’s friends in guiding his future relations.

seem to have no correlation with user’s rating whatsoever.

II. BACKGROUND AND RELATED WORK

The interplay of eCommerce and social networks has been studied in a number of prior work, which studied how to identify, propagate and predict trust (and distrust in some cases) relationships.

Signed Relationships. A relationship between two individuals can be signed, i.e., positive or negative. A positive relationship means friends or a trusted relationship. A negative relationship means foes or a distrusted relationship. The first theory on signed relationships are known as the *structural balance theory* (formulated by Heider in 1946 [5] and generalized by Cartwright & Harary in 1956 [6]). The theory considers the possible ways in which triangles on three individuals can be signed and states that triangles with exactly odd (one & three) number of positive signs are more prevalent in real networks as opposed to triangles with even (zero & two) number of positive signs. This is also termed the *strong structural balance*. Davis’ theory on *weak structural balance* [7] states that all types of triangles except the ones with exactly two positive signs are plausible in real networks.

Leskovec *et al.* [4] observed that real datasets from Epinions, Slashdot & Wikipedia on ratings are aligned with the theory of weak structural balance when ignoring the directed nature of links in ratings. Their findings though are at odds with Heider’s structural balance theory in two out of the three datasets. In the case of directed links, the authors proposed the *status theory* which states that both the sign and direction of the link convey a notion of status. The theory was tested and used for link prediction in a followup work [8].

Predicting Signed Relationships In real data sets it is often the case that only some of the edges are explicitly labeled with a sign. Thus one interesting problem is to predict the sign of any two individuals. Guha *et al.* consider propagation of trust or distrust in a network, thus predict the sign of an edge from the signs of other (neighboring) edges [9]. They build on the “path-algebra” model of trust propagation put forward by Richardson *et al.* [10]. [3], [11] build further on this idea. In [2], [3], Massa & Avesani, come to conclusion that local trust metrics computed via trust propagation are more effective in recommending “suitable” content as compared to global trust metrics and collaborative filtering techniques. Similarly, O’Donovan & Smyth in [11] demonstrate the usefulness of incorporating trust values into a collaborative filtering algorithm. We refer the reader to [12], [13], [14], [15] for works on propagating trust in other related settings.

A separate line of research considered a diverse set of features, at both a network level and an individual level, for “signed” link prediction in online social networks [16], [8]. They demonstrate high accuracy (in excess of 90% for 2 out the 3 datasets) in predicting links and their signs based on localized features [8]. More recent work by DuBois *et al.* [17] developed a spring-embedding algorithm that they use in conjunction with their path-probability technique [18] to infer trust or distrust.

III. DATASET

We conduct our experiments on the community of users of the popular consumer review website Epinions. The website allows users to write reviews about products and services. It also allows users to rate reviews written by other users. In our dataset, each review can be rated on an integer scale of [1 – 5] with a rating of 5 signifying a very good review. Epinions also allows users to define their *web of trust*, i.e. “reviewers whose reviews and ratings have been consistently found to be valuable” and their *block list*, i.e. “reviewers whose reviews are found to be consistently offensive, inaccurate or not valuable”. The *web of trust* and the corresponding *block list* for each user can be used to construct the directed, signed social network of trust and distrust.

TABLE I
OVERALL STATISTICS FOR EPINIIONS

Statistic for Epinions	
Number of Persons	131,828
Number of Reviews	1,197,816
Number of Trust Edges	717,667
Number of Distrust Edges	123,705
Number of Ratings	12,943,546
Time Range	Jan-2001 to Aug-2003

Table I provides statistics for the Epinions dataset. The dataset contains 132k users who issued 840k statements of trust and distrust, wrote 1.2 million reviews and assigned 13 million ratings to the reviews. The dataset spans 2.5 years and it also contains timestamp information on trust/distrust edge formation as well as time of rating a review. About 85k users have received at least one statement of trust or distrust in the dataset. Trust (positive) edges dominate the social network, i.e. over 85% of the edges are trust ones.

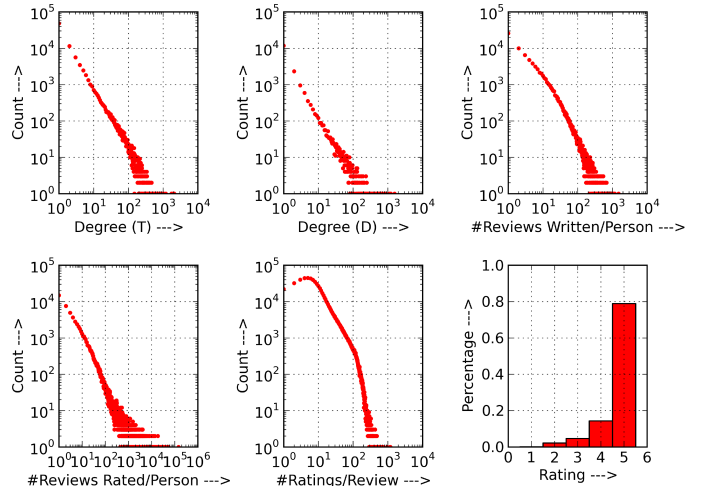


Fig. 1. Distributions for Epinions Data. The plots (from left to right, top to bottom) indicate, (1) Degree Distribution of Trust Edges, (2) Degree Distribution of Distrust Edges, (3) Distribution of the Number of Reviews Written per Person, (4) Distribution of the Number of Reviews Rated per Person, (5) Distribution of the Number of Ratings per Reviews, and (6) Distribution of Ratings in the Dataset.

Figure 1 details various distributions that arise in the dataset.

Summarizing the findings, we observe:

- A power-law degree distribution for both trust and distrust edges, i.e. which indicates that most users do not issue more than a handful of trust/distrust statements.
- A large number of users are passive users, i.e. users who write few or no reviews and rarely rate reviews.
- The bulk of content comes from a handful of users who take efforts on writing and rating reviews. We refer to these users as **active users**.
- A rating of 5 is by far the most prevalent rating in the dataset. Over 78% of all ratings have a score of 5; on the other hand only 0.01% and 2.13% of ratings have scores of 1 and 2 respectively.




Fig. 2. Relationship formation scenario. A is about to decide on whether to trust or distrust B at time t . $F_1, F_2 \dots F_n$ are friends of A (i.e. people trusted by A), some of whom have already expressed trust or distrust of B . Thick links indicate distrust, others indicate trust.

IV. RELATIONSHIP FORMATION

We investigate the correlation between the opinions of current friends and future relation formations under a simple scenario (see Figure 2). User A is about to trust or distrust user B at time t . At the moment, A has n friends/users, $F_1, F_2 \dots F_n$, whom he trusts. Among these n friends/users, b of them trust B and r of them distrust him. It is important to note that distrust relationships are hidden, i.e. A does not know that “ r ” of his friends distrust B .

We summarize here the results of this section:

- 1) First, we present raw observations based on the above scenario. Though informative, these observations should pass some processing to be statistically meaningful.
- 2) Then, we employ a random shuffling approach utilized in [4] to gauge whether the cases we observe are *over* or *under-represented* in the data relative to mere chance. Therefore, we compare our observed results with results achieved after randomly shuffling of the edge signs (trust/distrust). Besides over/under-representation, this approach allows us to determine statistical significance of our observed results.
- 3) Finally, our analysis should consider the fact that users in the data set exhibit a diverse spectrum of linking habits. Again, we use the approach put forward by [4]. We capture and gauge *trustfulness* and *trustworthiness* of users and employ them to get a better picture of the correlation between relationship formation and the opinions of friends.

A. Relationship Formation Scenario and Raw Observations

The following definition captures all the scenarios arising from Figure 2 in more detailed terminology. Depending on values of b and r , there are four distinct possible cases.

Definition 1. We categorize the cases as follows:

- 1) $b > r$ i.e. A 's friends **collectively trust** user B .
- 2) $b < r$ i.e. A 's friends **collectively distrust** user B .
- 3) $0 < b = r$ i.e. A 's friends are **neutral** on whether to trust or distrust B .
- 4) $0 = b = r$ i.e. A 's friends are **unopinionated** about user B .

We say that A 's friends are **opinionated** about his decision (trust/distrust) toward B if either (1) or (2) happens. On the other hand, if either (3) or (4) happens, we say that A 's friends are **neutral**. Depending on the cases, we can figure out if A **agrees/disagrees** with his friends in trusting B .

Raw Observations: Our observations show that in majorities of cases there is an agreement between the opinions of a user's friends and his/her future relationships for majority of users. The detailed results are presented in Table II. In more than 56% of the cases, A 's decision to trust or distrust B is aligned with his/her friends trust/distrust of B . Only in 3.69% of the cases, A 's decision differs from his friend's trust/distrust of B ³.

In 40% of the cases, A 's friends are unopinionated/neutral. The data show that a case of *neutral* (3) is very rare ($\approx 1\%$). We do not seek to test any correlation in the unopinionated/neutral cases.

Table III summarizes the data in Table II in terms of agreement rate. If A 's friends are *opinionated* about another user, then there is a strong correlation between A 's friends opinions and his/her decision (93.85%); both the *collectively trust* and *collectively distrust* categories show very high agreement rates (95.09% and 85.93% respectively). The *collectively distrust* case is specifically interesting because distrust links are *private* to other users.

B. Statistical Significance

We use a similar approach as in [4] to test the statistical significance of raw observations. In such an analysis we intend to gauge whether the cases we observe are *over* or *under-represented* in the data relative to mere chance. Therefore, we compare our observed results with results achieved after randomly shuffling of the edge signs (trust/distrust). In other words, if edge signs were produced at random (keeping the same proportion of trust and distrust edges), what would be the correlation between user's current relationships and his/her future ones?

Definition 2 (From [4]). We define the notion of **surprise**, s , as the number of standard deviations by which the actual quantity differs from the expected number under the random-shuffling model. Formally, for a given scenario S , an actual quantity attached to the scenario Q , expected quantity $E[Q]$

³In our dataset, we observe that about 69% of trust/distrust edges are observed on the first day of the crawl, i.e. they bear the same timestamp. Since there is no way of resolving the order in which these 69% of the edges were formed, we treat the network induced by these edges as the snapshot of the network at time t_0 . We conducted our relationship formation analysis for the remaining subset of edges (31%) that are formed at time $t > t_0$.

TABLE II

EVALUATION OF ROLE PLAYED BY FRIENDS IN DETERMINING RELATIONSHIPS RELATIVE TO THE RANDOM-SHUFFLING MODEL. CASES MARKED IN GREEN ($surprise > 0$) INDICATE OVER-REPRESENTATION AND CASES MARKED IN YELLOW ($surprise < 0$) INDICATE UNDER-REPRESENTATION.

Category	Case	Friends Choose To	A's Decision	Count	Surprise	
Agreeing with friends (56.27%)	$b > r$ $r > b$	Collectively Trust Collectively Distrust	Trust	129,229 (49.25%)	11.72	78.23
			Distrust	18,420 (7.02%)	457.88	
Disagreeing with friends (3.69%)	$b > r$ $r > b$	Collectively Trust Collectively Distrust	Distrust	6,656 (2.54%)	-101.98	-110.97
			Trust	3,016 (1.15%)	-40.59	
No role by friends (40.04%)	$0 < b = r$	Neutral	Trust	1,852 (0.71%)	-53.56	
			Distrust	878 (0.33%)	-9.49	-11.04
	$0 = b = r$	Unopinionated	Trust	76,636 (29.21%)	-36.95	
			Distrust	25,689 (9.79%)	87.75	

and prior probability of the scenario p_0 under the random shuffling model, surprise s is defined as follows:

$$s = \frac{Q - E[Q]}{\sqrt{E[Q](1 - p_0)}} \quad (1)$$

A surprise value of $s > 0$ indicates over-representation whereas a value $s < 0$ indicates under-representation. A surprise value of 6 yields a p-value of $\approx 10^{-8}$. In our experiments, we consistently encounter large surprise values, thereby making all our observations statistically significant.

TABLE III

THE AGREEMENT RATE OF USERS IN FORMING TRUST/DISTRUST RELATIONSHIPS IN SCENARIOS WHERE THEIR FRIENDS ARE OPINIONATED.

Category	Agreement Rate
Collectively Trust	95.09%
Collectively Distrust	85.93%
Opinionated	93.85%

Table II summarizes the results of the relationship formation analysis relative to the random shuffling model. One can clearly see over-representation in the ‘‘Agreeing with friends’’ scenario and under-representation in the other two scenarios. These results further support our findings:

- 1) There exists a strong correlation between a user’s friends opinion and formation of his/her future relationships.
- 2) Such correlation exists even considering the distrust links which are the hidden links.

C. Validity of Observations with respect to Linking Habits

The data set include users with a diverse spectrum of linking habits, meaning that they might have great differences in the number of trustees/trusted links. We would like to see whether our observations (in previous subsections) hold even when we take these different habits into consideration. Again, we use the approach put forward by Leskovec *et al.* in [4].

Definition 3 (From [4]). *We define generative and receptive baselines for a user as the fraction of positive links created and received by the user respectively. The baselines for a group of users are derived by summation of individual users’ baselines in the group.*

The generative baseline for user A is a measure of A ’s trustfulness towards other users. Whereas, the receptive baseline for user B is a measure of B ’s trustworthiness.

We also need the notions of *generative/receptive surprise* values in order to check the validity of observations with respect to linking habits.

Definition 4 (From [4]). *The generative surprise for a case is the number of standard deviations by which the actual number of positive $A \rightarrow B$ edges in the data differs above or below the expected number. The receptive surprise is defined along similar lines. If there was no correlation between friend of A ’s opinions and his decision to trust or distrust B , we would expect surprise values of 0.*

If A made a decision in forming a relation with B solely based on his linking habits (which is given by generative baseline), then *generative surprise* = 0. If B is trusted/distrusted by people based on his receptive baseline, then *receptive surprise* = 0.

TABLE IV

GENERATIVE & RECEPTIVE SURPRISE VALUES FOR ALL 4 SCENARIOS.

Case	Percentage of $A \rightarrow B$ Trust Edges	Generative Surprise	Receptive Surprise
Collectively Trust	95.10%	96.76	34.99
Collectively Distrust	14.07%	-104.15	-56.31
Neutral	67.54%	-6.73	-17.38
Unopinionated	74.89%	-26.87	-27.94

Table IV provides the generative and receptive surprise values for all the cases of Definition 1. The observations show that A ’s relationship with B has a strong correlation with his friends opinions regarding B . This alignment might be due to homophily/heterophobia or social influence and our data set can not conclusively show which is the case. However, the existence of such alignments has consequences and applications for websites like Epinions.

- In the *collectively trust* case, both the generative and receptive surprise values are significantly positive. This means that A exceeds generative baseline in trusting B (which can be attributed to the fact that his friends have collectively trusted B or just *homophily*). Also, B exceeds it’s receptive baseline in being trusted by A .
- In the *collectively distrust* case, both the generative and receptive surprise values are significantly negative. This means that A falls behind his generative baseline in distrusting B (which can be attributed to the fact that his friends have collectively distrusted B or *heterophobia*). B falls behind it’s receptive baseline in being distrusted by A .

V. RATINGS AND FRIEND OF FRIEND DYNAMICS

We extend the analysis to investigate the existence of any correlation between the opinions of the web of trust of a user and the opinions he/she expresses towards other users. As mentioned before, users of Epinions rate reviews of other users. We seek to determine if there is any correlation between the friendship dynamics and the ratings of a user. We will refer to the problem as analyzing *friend-of-friend (FoF) dynamics*, though it will also include the remaining three cases, namely, *friend-of-enemy (FoE)*, *enemy-of-friend (EoF)* and *enemy-of-enemy (EoE)*.

Definition 5 (friend-of-friend (FoF) dynamics). We define the problem using Figure 3. Consider three individuals A , B and C . A has expressed trust/distrust of B at time t_1 and B has expressed trust/distrust of C at time t_2 (without loss of generality we can assume that $t_1 < t_2$). A chooses to rate a review written by C at time t_3 ($t_1 < t_2 < t_3$).

We would like to quantify the correlation between A 's rating and the opinion of B regarding C . Notice that A has **NOT** expressed any opinion of trust or distrust about C at time t_3 .

The existence of a correlation might be used to improve recommendation systems (i.e. show content that would be useful to the user)⁴.




Fig. 3. Scenarios in friend-of-friend dynamics. In scenario (a), A rates a review written by C such that C is a friend-of-friend of A . Also A hasn't expressed any trust or distrust in C at the time of the rating. The remaining scenarios deal with enemy-of-friend, friend-of-enemy and enemy-of-enemy cases respectively. Trust & distrust edges are colored blue & red respectively.

There are four possible scenarios as shown in Figure 3. These four scenarios are as follows: FoF, EoF, FoE and EoE dynamics. In our dataset, 55% (≈ 7.2 million) of all ratings fall under at least one of these four scenarios (a rating can fall under multiple scenarios depending on A 's friends/foes being friends or foes with C). Table V provides the raw observations for each scenario. As illustrated in the table, the FoF scenario dominates (over 92%) the other three scenarios, while the EoE (under 1%) scenario being the most rare scenario.

⁴It is important to point out that this analysis differs from the one carried out by Leskovec *et al.* in [4]. They investigate signed triads in online social networks and analyze the over or under-representation of various triadic cases by applying theories from social psychology. The reviews and subsequent ratings by users do not appear in their analysis.

TABLE V

RANDOM-SHUFFLING MODEL ANALYSIS. R[1-5] REFERS TO THE FIVE RATINGS. SURPRISE VALUES INDICATE OVER/UNDER REPRESENTATION RELATIVE TO CHANCE UNDER THE RANDOM-SHUFFLING MODEL [4]. COLORS INDICATE WHETHER THE SURPRISE VALUES ARE GREATER (GREEN) OR LESS (YELLOW) THAN 0. WE SEE A SHIFT TOWARDS ASSIGNING HIGHER RATINGS (R5) IN THE FoF SCENARIO AND A SHIFT TOWARDS ASSIGNING LOWER RATINGS (R1 & R2) IN THE EoF SCENARIO.

Dynamics		FoF	EoF	FoE	EoE
Count (c)		77,801,592	2,225,299	3,069,957	746,361
%age (p)		92.79%	2.65%	3.66%	0.89%
Surprise (s)		4163.27	-2760.75	-2466.35	-826.46
R1	c	1585	1411	1053	134
	p	0.002%	0.06%	0.03%	0.02%
	s	-87.42	67.24	30.72	3.05
R2	c	109,193	216,251	75,253	28,927
	p	0.14%	9.72%	2.45%	3.88%
	s	-1214.14	782.66	36.66	103.25
R3	c	6,242,221	88,495	53,398	14,742
	p	8.02%	3.98%	1.74%	1.98%
	s	1440.39	-45.06	-244.93	-111.09
R4	c	8,825,021	240,785	247,666	55,783
	p	11.34%	10.82%	8.07%	7.47%
	s	-755.62	-147.63	-309.12	-165.92
R5	c	62,623,572	1,678,357	2,692,587	646,775
	p	80.49%	75.42%	87.71%	86.66%
	s	341.32	-129.48	377.85	162.89

Again we investigate the statistical significance of raw observations by using random shuffling model [4]. We also consider the different rating/rate-ability habits of users by incorporating the generative/receptive baselines approach of [4].

- 1) We shuffle the edge signs and compute the surprise values for each of the four scenarios. The fourth row in Table V illustrates the observations. These values indicate that the FoF scenario is over-represented in the dataset relative to chance, whereas the other three scenarios (EoF, FoE, EoE) are under-represented.
- 2) We keep the edge signs intact while shuffling the ratings around (keeping the same distribution of the ratings shown in figure 1). The random shuffling of the ratings allows us to study the over or under-representation of each rating in each scenario relative to chance (i.e. rating being assigned at random). Table V provides the raw counts, percentages and surprise values associated with ratings (as numbers in $\{1, 2, 3, 4, 5\}$) under each scenario.
- 3) To get a measure of the rating habits of a user, we extend the definitions of generative and receptive baselines defined in definition 3. Each user is associated with five **generative/receptive rating baselines** based on five categories of ratings. A **generative rating baseline** of 0.80 associated with a user A for rating of score 5 would imply that the user A assigns a rating of score 5 in 80% of times. Similarly, a **receptive rating baseline** of 0.20 associated with a user C for rating of score 4 would suggest that user C receives a rating of score 4 for 20% of his reviews. The generative rating baselines for user A is a measure of A 's opinions towards reviews written by other users, whereas

the receptive rating baselines for user C are a measure of other users's opinions of reviews written by C . We can now compute the generative rating surprise and receptive rating surprise values associated with each rating under each scenario. Table VI provides the *generative rating surprise* and *receptive rating surprise* values associated with each rating under each scenario.

TABLE VI

RATING HABITS ANALYSIS. SURPRISE VALUES INDICATE OVER/UNDER REPRESENTATION RELATIVE TO THE RATING HABITS OF THE USERS. s_g & s_r INDICATE THE GENERATIVE/ RECEPTIVE RATING SURPRISE VALUES RESPECTIVELY. WE AGAIN OBSERVE A SHIFT TOWARDS ASSIGNING HIGHER RATINGS (R5) IN THE FoF SCENARIO AND A SHIFT TOWARDS ASSIGNING LOWER RATINGS (R1, R2 & R3) IN THE EoF SCENARIO.

Dynamics		FoF	EoF	FoE	EoE
R1	s_g	-43.77	66.03	-8.36	-1.43
	s_r	-10.91	19.80	57.13	2.76
R2	s_g	-627.54	789.36	-108.83	26.54
	s_r	-527.77	89.58	206.16	4.03
R3	s_g	-360.72	2.01	-304.42	-124.23
	s_r	-181.17	10.16	65.53	5.89
R4	s_g	-847.21	-115.23	-381.94	-190.69
	s_r	-370.22	-3.57	81.06	-4.27
R5	s_g	1065.09	-189.93	531.88	214.75
	s_r	519.91	-61.03	-173.88	-1.36

A. Findings

We find the following by computing the under/over representation of each rating in each scenario (please refer to Tables V & VI for the surprise values): i) Low ratings (ratings of scores 1 and 2) are over-represented in all scenarios except the FoF one; ii) Rating of score 4 is under-represented across all scenarios; iii) Rating of score 5 (the highest and most frequent rating) is over-represented in all scenarios except the EoF one (in which it is under-represented). Based on our findings, we observe the following trends:

1) We see a clear trend of alignment between ratings and opinions of friends/foes in the FoF and EoF scenarios:

FoF: We see a shift towards higher ratings (rating of score 5 is over-represented and ratings of scores 1 and 2 are under-represented in this scenario). These results suggest that user A is more likely to assign higher ratings to user C 's reviews when C happens to be a friend of a friend of A . The generative and receptive surprise values also indicate such a correlation.

EoF: We see a trend towards assigning lower ratings (ratings of scores 1 and 2 are over-represented whereas all other ratings are under-represented). This indicates that A would be more likely to assign lower ratings to C 's reviews when C happens to be an enemy of a friend of A . Again, the generative and receptive surprise values support such a correlation.

2) In the remaining two scenarios of FoE and EoE, we see a divided picture. Both low ratings (scores of 1 and 2) and high rating (score of 5) are over-represented whereas the ratings in the middle (scores of 3 and 4) are under-represented. For these two scenarios, we are not able to conclusively show signs of correlation from this analysis.

VI. CONCLUSION

Studying correlation is very important in design and analysis of rating systems. We have showed that such a study can be utilized to improve the user experience in similar systems or websites. From the sociology point of view, both of the alignments discovered in this paper can be explained by social influence or homophily/heterophobia. However, we should mention that our analysis is not conclusive in favor of either explanation.

ACKNOWLEDGEMENTS

This research is funded in part by DARPA/ADAMS program under contract W911NF-11-C-0216. Any opinions, findings, and conclusions or recommendations in this material are those of the authors and do not necessarily reflect the views of the government funding agencies.

REFERENCES

- [1] "Word-of-mouth the most powerful selling tool: Nielsen global survey," http://nz.nielsen.com/news/Advertising_Oct07.shtml, 2007.
- [2] P. Massa and P. Avesani, "Controversial users demand local trust metrics: An experimental study on epinions.com community," in *AAAI*, 2005, pp. 121–126.
- [3] —, "Trust-aware recommender systems," in *Proceedings of the 2007 ACM conference on Recommender systems*. ACM, 2007, pp. 17–24.
- [4] J. Leskovec, D. P. Huttenlocher, and J. M. Kleinberg, "Signed networks in social media," in *CHI*, 2010, pp. 1361–1370.
- [5] F. Heider, "Attitudes and cognitive organization," *The Journal of psychology*, vol. 21, no. 1, pp. 107–112, 1946.
- [6] D. Cartwright and F. Harary, "Structural balance: a generalization of heider's theory," *Psychological review*, vol. 63, no. 5, p. 277, 1956.
- [7] J. Davis, "Structural balance, mechanical solidarity, and interpersonal relations," *American Journal of Sociology*, pp. 444–462, 1963.
- [8] J. Leskovec, D. Huttenlocher, and J. Kleinberg, "Predicting positive and negative links in online social networks," in *Proceedings of the 19th international conference on World wide web*. ACM, 2010, pp. 641–650.
- [9] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins, "Propagation of trust and distrust," in *Proceedings of the 13th international conference on World Wide Web*. ACM, 2004, pp. 403–412.
- [10] M. Richardson, R. Agrawal, and P. Domingos, "Trust management for the semantic web," *The Semantic Web-ISWC 2003*, pp. 351–368, 2003.
- [11] J. O'Donovan and B. Smyth, "Trust in recommender systems," in *Proceedings of the 10th international conference on Intelligent user interfaces*. ACM, 2005, pp. 167–174.
- [12] J. Parreira, D. Donato, C. Castillo, and G. Weikum, "Computing trusted authority scores in peer-to-peer web search networks," in *Proceedings of the 3rd international workshop on Adversarial information retrieval on the web*. ACM, 2007, pp. 73–80.
- [13] S. Kamvar, M. Schlosser, and H. Garcia-Molina, "The eigentrust algorithm for reputation management in p2p networks," in *Proceedings of the 12th international conference on World Wide Web*. ACM, 2003, pp. 640–651.
- [14] D. Quercia, S. Hailes, and L. Capra, "Lightweight distributed trust propagation," in *Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on*. IEEE, 2007, pp. 282–291.
- [15] J. Golbeck and J. Hendler, "Inferring binary trust relationships in web-based social networks," *ACM Transactions on Internet Technology (TOIT)*, vol. 6, no. 4, pp. 497–529, 2006.
- [16] H. Liu, E. Lim, H. Lauw, M. Le, A. Sun, J. Srivastava, and Y. Kim, "Predicting trusts among users of online communities: an epinions case study," in *Proceedings of the 9th ACM conference on Electronic commerce*. ACM, 2008, pp. 310–319.
- [17] T. DuBois, J. Golbeck, and A. Srinivasan, "Predicting trust and distrust in social networks," in *SocialCom/PASSAT*, 2011, pp. 418–424.
- [18] —, "Rigorous probabilistic trust-inference with applications to clustering," in *Web Intelligence and Intelligent Agent Technologies, 2009. WI-IAT'09. IEEE/WIC/ACM International Joint Conferences on*, vol. 1. IET, 2009, pp. 655–658.