

Opinion Formation Under Costly Expression

FANG WU and BERNARDO A. HUBERMAN

HP Labs, Palo Alto

Opinions play an important role in trust building and the creation of consensus about issues and products and a number of studies have focused on the design, evaluation, and utilization of online opinion systems. However, little effort has been spent on the dynamic aspects of online opinion formation. In this article, we study the dynamics of online opinion expression by analyzing the temporal evolution of very large sets of user views and determine that in the course of time, later opinions tend to show a big difference with earlier opinions, which moderates the average opinion to the less extreme. Online posters also tend to disagree with previous opinions when the cost of expression is high.

Categories and Subject Descriptors: J.4 [Social and Behavioral Science]— *Sociology*

General Terms: Theory, Human Factors

Additional Key Words and Phrases: Opinion formation, costly expression

ACM Reference Format:

Wu, F. and Huberman, B. A. 2010. Opinion formation under costly expression. *ACM Trans. Intell. Syst. Technol.* 1, 1, Article 5 (October 2010), 13 pages.

DOI = 10.1145/1858948.1858953 <http://doi.acm.org/10.1145/1858948.1858953>

5

1. INTRODUCTION

Since opinions play such an important role in trust building and the creation of consensus about many issues and products, there have been a number of recent studies focused on the design, evaluation and utilization of online opinion systems [Dellarocas 2000, 2006; Cosely et al. 2003; Gao et al. 2006; Hu et al. 2006; Danescu-Niculescu-Mizil et al. 2009] (for a survey, see Dellarocas [2003]). Given the importance of group opinions to collective social processes such as group polarization [Asch 1955; Myers 1982; Sunstein 2002] and information cascades [Banerjee 1992; Bikhchandani et al. 1992; Salganik et al. 2006; Chevalier and Mayzlin 2006], it is surprising that with the exception of one study [Li and Hitt 2004], little research has been done on

Author's address: B. A. Huberman, HP Labs, 1501 Page Mill Road, Palo Alto, CA 94301; email: bernardo.huberman@hp.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2010 ACM 2157-6904/10-ART5 \$10.00
DOI 10.1145/1858948.1858953 <http://doi.acm.org/10.1145/1858948.1858953>

ACM Transactions on Intelligent Systems and Technology, Vol. 1, No. 1, Article 5, Publication date: October 2010.

the dynamic aspects of online opinion formation. It remains unclear, for example, whether the opinions about books, movies, or societal views fluctuate a long time before reaching a final consensus or they undergo systematic changes as time goes on. Thus the need to understand how online opinions are created and evolve in time in order to draw accurate conclusions from that data.

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 little effort, like clicking on a button of a Web site, it is more puzzling in situations where it is costly, such as composing a review. (e.g., when an Amazon user 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 [Ketzan 2002], so the cost of opinion expression is indeed high.) 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 not. This is reminiscent of the well analyzed voter's paradox [Downs 1957; Riker and Ordeshook 1968] 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. 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.

Recently, Schlosser [2005] has pointed out that posters (those who post opinions) tend to disagree with previous ones because they do not want to be perceived as indiscriminate, while lurkers (those who do not post opinions) are less affected by others' negative opinions [Amabile 1983]. In an empirical study involving a few hundred students, she found out that posters presented more than one side when publicly explaining their attitudes despite their favorable product experiences and commitment to these attitudes. Interesting as these findings are, they have not been extended to very large scale (i.e., millions of opinions) in a real-world setting.

Within this context, we studied the dynamics of online opinion expression by analyzing the temporal evolution of a very large set of user views, ranging from online reviews of the best-selling books at Amazon.com, and movie reviews at the Internet Movie Database IMDB. On the massive scale that the web offers, we observed that later opinions in the course of time tend to show a large difference with previous ones, which moderates the average opinion to the less extreme. Furthermore, posters tend to disagree with previous ones when the cost of expression is higher. This is a robust and quantitative observation which conforms with Schlosser's hypotheses.

2. STUDY 1: AMAZON.COM

2.1 Data

We will use the following notations throughout this article: X_n denotes the quantified value of the n th opinion of one particular item; $\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X_i$ denotes the average value of the first n opinions of one particular item; EX_n denotes the sample mean of X_n for all items in a dataset; $E\bar{X}_n$ denotes the sample mean of \bar{X}_n for all items in a dataset.


In order to perform our study on online opinions, we first analyzed book ratings posted on Amazon.com. Our sample consisted of 1,729,830 book ratings of the top 4,000 best-selling titles of Amazon in each of the following 12 categories, as of July 1, 2007: arts and photography, biographies and memoirs, history, literature and fiction, mystery and thrillers, reference, religion and spirituality, sports, travel, nonfiction, science, and entertainment. For each of the 48,000 books, a series of user ratings was collected in time order, where each rating is an integer between 1 and 5. Among the 48,000 books, 16,454 books have no less than 20 ratings, 7,385 books have no less than 50 ratings, and 3,495 books have no less than 100 ratings.

2.2 Results


We first checked the average rating of books (EX_n) in each of these three sets of books as a function of the rating index (n). As can be seen from Figure 1(a), the ratings are generally favorable (the average of all 1,729,830 ratings is $\bar{X}_{\text{all}} = 4.04$). There is a clear decreasing trend, indicating that later users tend to write different reviews from those of earlier users. This is consistent with the empirical finding of Ariely and Levav [2000] that when group members voice their opinions sequentially, an opinion shift may be induced. However, as opposed to the widely observed group-polarization effect, the overall opinion on Amazon tends to decrease away from the extreme ones.

Since the book ratings in our dataset are generally favorable, one might wonder whether the moderating effect in Figure 1(a) is merely an artifact of high initial ratings. It would be helpful if one could locate a set of books with low initial ratings and see if their ratings also exhibit a moderating effect over time. To this end, we divided the books into 6 tiers based on the average value of their first 3 ratings (\bar{X}_3) and plotted the average rating as a function of the rating index for each tier. As one can see in Figure 1(b), all 6 tiers tend to move closer to the overall mean $\bar{X}_{\text{all}} = 4.04$ as opposed again to the divergent pattern commonly observed in group polarization experiments.

In order to better compare our results with existing studies of group polarization, which usually measure the difference between the final and initial average opinion, we next measured the change of ratings at the level of each individual book. For each book with $l \geq 2n$ ratings, we calculated the mean of its first n ratings ($\bar{X}_{1,n} = \frac{1}{n} \sum_{i=1}^n X_i$) and the mean of its last n ratings ($\bar{X}_{l-n+1,l} = \frac{1}{n} \sum_{i=l-n+1}^l X_i$). We then performed a series of statistical tests on the alternative hypothesis " $\bar{X}_{l-n+1,l} < \bar{X}_{1,n}$ " for books in each tier, n ranging from



(a)



(b)

Fig. 1. (a) Average rating of Amazon book reviews as a function of the rating index. EX_n is the sample average rating of all the n th ratings of books with no less than (1) 20 ratings, (2) 50 ratings, and (3) 100 ratings. As one can see from the figure, EX_n decreases by 0.4 stars in 20 steps. (b) We divided the 16,454 books which have no less than 20 ratings into 6 tiers. From top to bottom, respectively, they contain 11,393 books with $\bar{X}_3 \in (4.5, 5]$, 1,645 books with $\bar{X}_3 \in (4, 4.5]$, 2,252 books with $\bar{X}_3 \in (3.5, 4]$, 500 books with $\bar{X}_3 \in (3, 3.5]$, 562 books with $\bar{X}_3 \in (2, 3]$, and 102 books with $\bar{X}_3 \in [1, 2]$. The horizontal dashed line $X_{\text{all}} = 4.04$ is the overall mean of all 1,729,830 ratings in our full dataset. We see that the average book rating in each tier shifts towards the dashed line.

Table I. Test of the Null Hypothesis $\bar{X}_{l-n+1,l} > \bar{X}_{1,n}$ for Books in 6 Tiers. The Average Rating of Books in Tier 1 and 2 Goes Down While the Average Ratings of Books in Tier 4, 5, 6 Goes Up

Tier	\bar{X}_3	Result
1	(4.5, 5]	Rejects null with $p < 0.001$ for $n = 1, \dots, 50$
2	(4, 4.5]	Rejects null with $p < 0.001$ for $n = 1, \dots, 50$
3	(3.5, 4]	Insignificant
4	(3, 3.5]	Supports null with $p < 0.05$ for $n = 1, \dots, 17$
5	(2, 3]	Supports null with $p < 0.05$ for $n = 1, \dots, 33$
6	(1, 2]	Supports null with $p < 0.05$ for $n = 1, \dots, 50$

1 to 50. The results are listed in Table I (see also Figure 2). As can be seen, the conclusions drawn from Figure 1(b) are robust.

While these empirical findings reveal interesting trends in online opinion formation in contrast to group polarization, we stress that the convergence pattern shown in Figure 1(b) does not necessarily imply that people come to an agreement in the end. In fact, when we measured the standard deviation $\sigma_{1,n}$ of the first n ratings and $\sigma_{l-n+1,l}$ of the last n ratings for all books with no less than $2n$ ratings, all 49 tests supported the alternative hypothesis “ $\sigma_{l-n+1,l} > \sigma_{1,n}$ ” with p -value less than 0.001 for $n = 2, \dots, 50$, suggesting diversification of opinions over time. We also measured the fraction of each star rating as a function of the rating index. Let $f_n(r)$ denote the fraction of r -star ratings among all the n th ratings in our dataset for $r = 1, 2, 3, 4, 5$. Figure 3 plots the growth of fractions for $r = 1, 2, 3, 4$. As one can see, the fraction of 1-star ratings grows fastest, indicating that instead of ending with a group of people with moderate opinions, one observes a variety of diverse views. This shows that public discourse on Amazon typically intensifies over time and does not reach a consensus in the end in contrast to the imitative behavior seen in herding and information cascades.

In Figure 1(b), we could only compare the ratings of two groups of subjects (the first n and the last n), for Amazon does not allow its users to leave multiple reviews. Hence our results do not necessarily imply that as time goes on the average opinion of the whole population changes for the late reviewers might come from a different group than the earlier ones and need not be representative of the whole population. This is seen when plotting the average “helpful ratio” as a function of star rating in Figure 4 for users of Amazon. As can be seen, the whole population finds high ratings in general more helpful than low ratings. If one assumes that agreeing with a review is correlated with finding it useful, this implies that the majority of the population does not necessarily agree with the low ratings. This additional data then suggests that rather than indicating a real choice shift in the whole population, the observed dynamic trend on Amazon might be more of a *selection bias*. By this we mean that the more recent ratings are generated by a subgroup of users whose opinions deviate considerably from the average opinion of the whole group, a point that we will elaborate below.





Fig. 2. Sample histograms of statistical tests listed in Table I. (a) Histogram of $\bar{X}_{l-49,l} - \bar{X}_{1,50}$ for books in the top tier. Notice the shift towards the left of the origin, signaling a downward trend in the average opinion. (b) Histogram of $\bar{X}_{l-19,l} - \bar{X}_{1,20}$ for books in the bottom tier. Notice that the majority of the data is to the right of the origin, indicating an upward trend in the average opinion.




Fig. 3. The change in the fraction of 4-star, 3-star, 2-star, and 1-star reviews on Amazon.

2.3 Discussion

One possible explanation for these results is that in cases like Amazon, people will derive more utility the more they can influence the overall rating as in the voter’s paradox. To be precise, in cases where users’ opinions can be quantified and aggregated into an average value (when a user opens up a book page on Amazon, she sees the average rating at the top of the page right under the title), the influence of an online opinion can be measured by how much its expression will change the average opinion [Osborne et al. 2000]. Suppose that n users have expressed their opinions, X_1, \dots, X_n , on a given topic at a Web site. If the $(n + 1)$ th person expresses a new opinion X_{n+1} , it will move the average rating to $\bar{X}_{n+1} = (n\bar{X}_n + X_{n+1})/(n + 1)$, and the absolute change in the average rating is given by $|\bar{X}_{n+1} - \bar{X}_n| = |X_{n+1} - \bar{X}_n|/(n + 1)$. Thus 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 similarities to Schlosser’s hypotheses [2005], and has also been made within the “brag-and-moan” model [Hu et al. 2006; Dellarocas and Narayan 2006], 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 to the observed dynamical trends.

In order to test our hypothesis, we measured directly how much one’s rating deviates from the observed average rating. We plot the expected deviation




Fig. 4. (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 best-selling 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.

$Ed_n = E|X_n - \bar{X}_{n-1}|$ as a function of n in Figure 5(a), where X_n is the rating left by the n th user, and \bar{X}_{n-1} is the average rating the n th user observes. As can be seen, Ed_n increases with n . Since the expected deviation Ed_n of an independent and identically distributed sequence normally decreases with n , this increasing trend is indeed significant.





Fig. 5. (a) The average deviation of Amazon ratings increases with the number of reviewers. (b) The average deviation of X_n from $(X_{n-3} + X_{n-2} + X_{n-1})/3$ increases with the number of reviewers.

It is known in the group polarization literature that exposure to the group average is sufficient to stimulate a more polarized response and that those exposed to the full distribution of others' choice are not significantly more polarized than those who merely witnessed the group average [Brown 1974; Myers 1978]. Because Amazon also shows its users a histogram of all the past

ratings and a list of the most recent ratings near the bottom of the page, we also measured the deviation of each rating from its recent three predecessors. Figure 5(b) plots $Ed'_n = E|X_n - (X_{n-3} + X_{n-2} + X_{n-1})/3|$ as a function of n . Not only the increasing trend remains but the deviation actually exceeds Ed_n by 0.05 stars. This again supports our conjecture that those users who disagree with the existing average will be more willing to express themselves and thus soften the overall opinion of a given book.

3. STUDY 2: IMDB.COM

3.1 Data

While our hypothesis seems to explain the softening of opinions observed in Amazon, it would be more conclusive if one could conduct a test that directly compares people's opinions expressed at different cost levels. In order to address this issue we conducted a study of IMDB.com (The Internet Movie Database). Unlike users of Amazon who are required to write a review when rating a book, users of IMDB are free to choose the effort level when reviewing a movie. Specifically, after observing the current average rating of a movie, a user can either submit a quick rating by clicking on a scale of 10 stars or make the extra effort involved in writing a comment between 10 and 1,000 words.

Our dataset consists of all 1,275 USA movie titles released after year 2000 which have no less than 5,000 ratings, as of August 1, 2008. The list of titles were obtained using IMDB's Power Search engine. A constraint on the number of ratings is necessary because the proportion of commented rating among all ratings is very low (1.17% in our dataset). For each movie, we know its average rating (taken among all ratings with or without a comment) as well as the value and date-stamp of each of its commented ratings, but we do not have any specific information about each uncommented rating. There are 407,557 commented ratings in total. We divided the movies into 3 tiers based on their ratings: the top tier contains 545 "good movies" whose ratings are at least 7, the bottom tier contains 34 "bad movies" whose ratings are at most 4, and the middle tier contains the rest.

3.2 Results

We observed a clear decreasing trend in the ratings associated with comments in the top tier, as shown in Figure 6. We did not observe any clear trend in the bottom tier. Similar to the Amazon study, we tested the alternative hypothesis " $\bar{X}_{l-n+1,l} > \bar{X}_{1,n}$ " in each of the three tiers, for $n = 1, \dots, 50$. The 100 tests in the top tier and the middle tier all support the alternative with p -value less than 0.001. None of the 50 tests in the bottom tier yield statistically significant results (i.e., the alternative can neither be supported nor rejected), possibly due to its small size. While it is not reliable to say that bad movies tend to receive higher ratings, it is safer to conclude that good movies accumulate lower ratings as time goes on.

We also examined the difference between the overall average rating (with or without a comment) and the average rating associated with a comment for




Fig. 6. Average rating associated with a comment of movies with average rating no less than 7 as a function of the number of existing ratings. It can be seen that good movies tend to receive lower ratings as time goes on.

each movie, and the result is shown in Figure 7. Assuming as we do that the benefit of opinion expression diminishes as the number of opinions grows, it follows that it is those who are willing to incur a high cost that will continue to express opinions at very late times. Since the data in the figure shows that the latter deviate opposite from the population average, this offers a tentative (but not conclusive) explanation for the moderating effect in the average opinion.

An alternative explanation is that people are reluctant to be the first one to dissent so after the first bad review appears, a series of bad reviews follow up. To test the effect of pioneering bad reviews, we counted the number of 1-star ratings among the first n ratings of each Amazon title (denoted by Y_n), and then performed the linear regression $X_{n+1} \sim a + b\bar{X}_n + cY_n$ for $n = 2, 3, 4, 5$. All the fittings yield a statistically significant coefficient c between -0.002 and -0.001 , suggesting a weak herding effect.

4. CONCLUSION

In this article, we studied the dynamics of online opinion expression by analyzing the temporal evolution of a very large set of user views, ranging from online reviews of the best selling books at Amazon.com, and movie reviews at the Internet Movie Database IMDB. Our observations revealed that later opinions tend to show a large difference with previous ones, which in turn moderates the average opinion to the less extreme. When opinion expression is costly, we found that people tend to project a perception of being discriminate by




Fig. 7. Selection bias induced by commented ratings. Each cross in this figure corresponds to one movie title. The horizontal coordinate represents the movie's overall average rating (\bar{r}) taken over both commented and uncommented ratings. The vertical coordinate represents the movie's average rating taken over only commented ratings (\bar{r}_c). The two vertical lines separate the data points into three tiers. Clearly, points in the bottom tier tend to fall above the line, and points in the top tier tends to fall below. That is to say, those users who spend the additional cost to write a comment tend to express opinions opposite to those of the majority. A t -test of the alternative hypothesis that $\bar{r}_c < \bar{r}$ for good movies and a similar t -test of $\bar{r}_c > \bar{r}$ for bad movies both yield a p -value less than 0.001.

disagreeing with previous ones. This is a robust and quantitative observation in conformity with Schollosser's hypotheses [2005].

These results also show that in the process of articulating and expressing their views online, people tend to follow a different pattern from that observed in information cascades or group polarization. This might be due to the fact that both at Amazon and IMDB, the setting and measurements differ from the standard ones in group polarization experiments and information cascades. We point out, however, that it is opinions like the ones we studied that are increasingly used by millions of people to navigate the sea of information they encounter every day. Thus these results might be of practical value.

We close with a cautionary note on the interpretation of online 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

AMABLE, T. M. 1983. Brilliant but cruel: Perceptions of negative evaluators. *J. Experim. Psych.* 19, 146–56.

- ARIELY, D. AND LEVAV, J. 2000. Sequential choice in group settings: Taking the road less traveled and less enjoyed. *J. Consumer Resear.* 27, 279–290.
- ASCH, S. E. 1955. Opinions and social pressure. *Scientific Amer.* 193, 5, 31–35.
- BANERJEE, A. V. 1992. A simple model of herd behavior. *Quart. J. Econ.* 107, 3, 797–817.
- BIKHCHANDANI, S., HIRSHLEIFER, D., AND WELCH, I. 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Political Eco.* 100, 5, 992–1026.
- BROWN, R. 1974. Further comment on the risky shift. *Amer. Psychol.* 29, 468–470.
- CHEVALIER, J. A. AND MAYZLIN, D. 2006. The effect of word of mouth online: Online book reviews. *J. Market. Resear.* 43, 3, 345–354.
- COSLEY, D., LAM, S. K., ALBERT, I., KONSTAN, J. A., AND RIEDL, J. 2003. Is seeing believing? How recommender interfaces affect users opinions. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI'03)*.
- DANESCU-NICULESCU-MIZIL, C., KOSSINETS, G., KLEINBERG, J., AND LEE, L. 2009. How opinions are received by online communities: A case study on Amazon.com helpfulness votes. In *Proceedings of the World Wide Web Conference (WWW'09)*.
- DELLAROCAS, C. 2000. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *Proceedings of the ACM Conference on Electronic Commerce (EC'00)*.
- DELLAROCAS, C. 2003. The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Manag. Science* 49, 10, 1407–1424.
- DELLAROCAS, C. 2006. Strategic manipulation of internet opinion forums: Implications for consumers and firms. *Manag. Science* 52, 10, 1577–1593.
- DELLAROCAS, C. AND NARAYAN, R. 2006. What motivates consumers to review a product online? A study of the product-specific antecedents of online movie reviews. In *Proceedings of the International Conference on Web Information Systems Engineering (WISE'06)*.
- DOWNES, A. 1957. *An Economic Theory of Democracy*. Harper & Row.
- GAO, G., GU, B., AND LIN, M. 2006. The dynamics of online consumer reviews. In *Proceedings of the International Conference on Web Information Systems Engineering (WISE'06)*.
- HU, N., PAVLOU, P. A., AND ZHANG, J. 2006. Can online reviews reveal a product's true quality? Empirical findings and analytical modeling of online word-of-mouth communication. In *Proceedings of the ACM Conference on Electronic Commerce (EC'06)*.
- KETZAN, E. 2002. Amazon.com and the new democracy of opinion.
http://www.themodernword.com/pynchon/papers_ketzan1.html.
- LI, X. AND HITT, L. M. 2004. Self selection and information role of online product reviews. In *Proceedings of the International Conference on Web Information Systems Engineering (WISE'04)*.
- MYERS, D. G. 1978. The polarizing effects of social comparison. *J. Experim. Social Psych.* 14, 554–563.
- MYERS, D. G. 1982. Polarizing effects of social interaction. In *Group Decision Making*, J. H. D. H. Brandstätter and G. Stocker-Kreichgauer, Eds. Academic Press, 125–161.
- OSBORNE, M. J., ROSENTHAL, J. S., AND TURNER, M. A. 2000. Meetings with costly participation. *Amer. Econ. Rev.* 90, 4, 927–943.
- RIKER, W. H. AND ORDESHOOK, P. C. 1968. A theory of the calculus of voting. *Amer. Political Science Rev.* 62, 25–42.
- SALGANIK, M. J., DODDS, P. S., AND WATTS, D. J. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311, 5762, 854–856.
- SCHLOSSER, A. E. 2005. Posting versus lurking: Communicating in a multiple audience context. *J. Consumer Resear.* 32.
- SUNSTEIN, C. R. 2002. The law of group polarization. *J. Political Philos.* 10, 2, 175–195.

Received January 2010; revised June 2010; accepted June 2010