Social Computing

Online Product Rating Manipulation and Market Performance[1]

Wenqi Zhou, Duquesne University

Yuhong Liu, Santa Clara University

Manipulating online ratings of a product, in terms of both volume and value, can substantially influence its market performance, but the benefits of a particular strategy can vary across products and might not be maximized by the highest rating values.

With the rapid growth of e-commerce and social media, online rating systems that let users post reviews of products and services are playing an increasingly important role in consumers' Web-based purchasing decisions.[1] According to a 2013 survey conducted by Dimensional Research, 90 percent of respondents said that positive online reviews influenced their decision to buy a product, while 86 percent said that their decision not to buy a product was influenced by negative online reviews (www.zendesk.com/resources/customer-service-and-lifetime-customer-value).

Given this trend, the incentive for companies to manipulate such ratings is rapidly rising, and there's ample evidence of online vendors using "reputation service" companies to inflate ratings of products to boost sales. For example, for $35, IncreaseYouTubeViews.com will provide a vendor with 30 "I like this" ratings of a promotional video clip to help it go viral.

However, the economic benefits *vs.* costs to online vendors of rating manipulation have never been quantified. To address this deficiency in the literature, we applied statistical analysis to publicly available software download data to assess the impact of rating manipulation on market performance. Building on these results, we assessed the effectiveness of the same manipulation strategy for different products with distinct ratings. Finally, while employing a well-known defensive scheme against rating manipulation, we compared the utility of various rating manipulation strategies.

Does Rating Manipulation Boost Market Performance?

We first quantified the potential impact of online rating manipulation on market performance. Specifically, we sought to determine whether and to what degree a product's market performance can be enhanced by an increase in average rating and an increase in rating volume, all else being the same.

To assess these effects scientifically, we performed a statistical analysis on a large set of robust real-market data. Due to the difficulties of accessing proprietary sales figures, we obtained our data from CNET Download.com (CNETD), a leading and representative platform hosting more than 150,000 free-to-try legal downloads ranging from security patches to mobile apps to games. In particular, the dataset consisted of the weekly number of downloads of all Web browsers over the course of six months and included each program's average user rating (out of five stars) and rating volume.[2]

We regressed the number of weekly software downloads on average rating and rating volume while controlling for software characteristics, network effect, product diffusion, and product intrinsic quality. The increase rate of downloads as a result of rating manipulation can be measured by $10\^(\alpha_1 \times \Delta\text{Rating} + \alpha_2 \times \Delta\text{Volume}) - 1$, where ΔRating and ΔVolume are the increment in average rating and rating volume, respectively, and α_1 and α_2 are the coefficients to be estimated.

Using our regression model, we estimated α_1 to be 0.0328 and α_2 to be 0.0005 (model details and the complete estimation report are available upon request). These values indicate that weekly downloads of Web browsers on CNETD were boosted 7 percent by each one-star increase in average rating ($10\^(0.0328) - 1$) and 0.12 percent by

[1] Please cite this paper as Zhou, W and Y. Liu "Online Product Rating Manipulation and Market Performance," *IEEE Computer*, 48 (5), 2015, pp. 72-75.

the addition of each review (10^(0.0005) −1).

The results show that both average rating value and rating volume can substantially impact a product's market performance. The underlying reasons for this effect are easily surmised: a high average rating value conveys a positive image of a product by users familiar with it and thereby instills confidence among other users in its desirability or quality; a high rating volume signals the product's deep marketplace penetration and popularity (even in the case of a large number of negative reviews) and thus attracts more attention from potential consumers. Given these results, online vendors clearly have a strong economic motivation to manipulate product ratings.

Do Existing Ratings Affect Manipulation Effectiveness?

Online vendors must consider the economic costs as well as the benefits of manipulating ratings of their products. (Nonmonetary costs, such as loss of reputation, are beyond the scope of this article). Given that reputation services generally charge by volume—more positive reviews, tweets, Facebook likes, and so on—producing more false ratings will incur more costs. To maximize profits, vendors thus have to determine the point at which the benefits exceed the costs.

In this context, we wanted to know whether adding the same number of manipulated ratings for two different products would yield the same monetary benefit. To answer this question, we first compared the impact of adding 300 false five-star ratings for two hypothetical products, A and B, with different existing rating values. For simplicity, we assumed that the rating system had no security scheme in place to prevent manipulation and thus accepted all new ratings.

Product A served as a benchmark and had 200 existing ratings with an average value of 1.5 stars before manipulation. Inserting the false ratings increased the rating volume by 300 and the average value by 2.1 stars. According to our statistical analysis, this rating manipulation would increase the number of product A's weekly downloads on CNETD by nearly 70 percent.

Product B had the same initial rating volume of 200 but an average rating value of three stars. We estimated that adding 300 false five-star ratings would increase the number of product B's weekly downloads by 55.73 percent, about 20 percent lower than that for product A. Clearly, the average value of existing ratings significantly affects manipulation effectiveness.

To isolate the effect of rating volume, we compared the impact of adding 300 false five-star ratings for product B and a product C with the same average rating value of 1.5 stars but only half of the number of existing ratings (100). In this case, the number of weekly downloads of product C on CNETD would increase by 43.33 percent, far below the percent increase for product A.

In summary, the effectiveness of a given online rating manipulation strategy can vary dramatically for different products depending on the average rating value and existing rating volume. For products with a large amount of very low ratings, adding false ratings might have a limited impact on market performance and might not be worth the investment.

What's the Best Rating Manipulation Strategy for a Given Product?

Assuming an online vendor has determined that inflating its product's ratings is economically beneficial, what's the optimum strategy? Given that market performance will monotonically increase when either the average rating value or rating volume is boosted, the best strategy would seem to be to maximize the number of false ratings with the highest value. This is impractical in the real world, however, as many online rating systems have security mechanisms that can easily detect and remove ratings with suspiciously high values, leading to a waste of manipulation effort.

In fact, designing the best manipulation strategy is an intricate and complex process. To illustrate some facets of the problem using CNETD data, we once again consider product A and assume that 100 of its existing ratings are one star and the other 100 ratings are two stars. Moreover, product A's rating system employs the beta reputation scheme,[3] a representative defensive mechanism that assumes normal ratings follow beta distributions and thus removes the outlier ratings.

Figure 1 shows the distribution of increased weekly software downloads for product A, indexed by color, caused by varying both the volume and average value of false ratings. For simplicity, we assume that each manipulation strategy assigns all false ratings an identical value. Note that the rating system only allows integer values ranging from one to five stars.

Figure 1. Effect of different online rating manipulation strategies on a hypothetical software product available on CNET Download. A monotonic relationship exists between the number of false ratings and download increases when the former is between 60 and 400. When the rating volume exceeds 400, downloads will be maximized by adding false four-star rather than five-star ratings.

In this figure, we can see that adding a fixed number of false five-star ratings might not maximize downloads. In fact, as the top graph shows, when the volume of false ratings is small (<60), five-star ratings clearly aren't the best choice: adding 10 false two-star ratings will increase the number of downloads by the same small amount, 3.44 percent, as 50 false five-star ratings.

This effect can be attributed to the beta reputation mechanism, which will identify extremely high ratings as false and remove them from the system when the volume of false ratings after manipulation is small compared to the original rating volume. On the contrary, low-value ratings (for example, two stars), won't be purged and thus will contribute to the rating volume increase, even though they're not maximally effective at increasing the overall average rating value. In short, choosing the appropriate false rating value is critical to determining net manipulation profits.

As Figure 1 shows, a monotonic relationship exists between the number of false ratings and download increases when the former is between 60 and 400. However, when the rating volume exceeds 400, surprisingly, downloads will be maximized by adding false four-star rather than five-star ratings. The reason is that, when a large number of false ratings predominate, providing false ratings with extremely high values (for example, five stars) will push the rating distribution toward the high end. In such a scenario, the beta reputation mechanism might consider existing low-value ratings as false and delete them, leading to a decrease in the rating volume that in turn inhibits the increase in downloads.

User ratings have emerged for almost category of product and service on the Web. However, media reports and online forums abound with complaints about suspected false ratings, and reputation service companies that charge vendors to inflate ratings of products have proliferated. So just how trustworthy are these ratings? To help answer this question, researchers are exploring whether manipulating such ratings is profitable and, if so, what the best strategy might be.

In considering this issue, we applied statistical analysis to software download data that combined aspects of marketing, information systems, and cybersecurity. Our analysis revealed that online rating manipulation, in terms of both rating volume and value, can substantially influence a product's market performance. In addition, the benefits of a particular manipulation strategy can vary dramatically across products. Finally, contrary to what many might assume, providing the highest rating values is a questionable strategy and in many cases definitely not the best one.

References

1. W. Duan, B. Gu, and A.B. Whinston, "The Dynamics of Online Word-of-Mouth and Product Sales—An Empirical Investigation of the Movie Industry," *J. Retailing*, vol. 84, no. 2, 2008, pp. 233–242.
2. W. Zhou and W. Duan, "Online User Reviews, Product Variety, and the Long Tail: An Empirical Investigation on Online Software Downloads," *Electronic Commerce Research and Applications*, vol. 11, no. 3, 2012, pp. 275–289.
3. A. Jøsang and R. Ismail, "The Beta Reputation System," *Proc. 15th Bled Electronic Commerce Conf.* (Bled 02), 2002; http://folk.uio.no/josang/papers/JI2002-Bled.pdf.

Wenqi Zhou *is an assistant professor of information systems management in the Palumbo-Donahue School of Business at Duquesne University. Contact her at zhouw@duq.edu.*

Yuhong Liu *is is an assistant professor of computer engineering in the School of Engineering at Santa Clara University. Contact her at liuyuhongbupt@gmail.com.*