

Revoking and Moral Hazard on eBay: An Empirical Investigation

Extended Abstract: 2677

Shun Ye Gordon Gao Siva Viswanathan

Robert H. Smith School of Business, University of Maryland, College Park

e-mail: {sye, ggao, sviswana}@rhsmith.umd.edu

Abstract

This study examines a crucial aspect of the reputation mechanism design in electronic markets – the ability of buyers and sellers to revoke or mutually withdraw negative feedback and ratings. We find that the two-way reputation system – where both buyers as well as sellers could provide mutual feedback - enabled certain sellers to behave opportunistically by getting buyers to revoke their negative feedbacks, making the reputation system less effective in discerning the quality of sellers. We also find that recent changes to the two-way reputation system have a significant influence on these sellers’ behavior. After the ban on revoking, we find that sellers exert more effort in improving the quality of their transactions and their reputation. Our findings support the moral hazard, rather than the adverse selection, assumption regarding seller behavior.

1. Introduction

Reputation systems are useful tools to reduce information asymmetry in online markets. Arguably, the most popular example is eBay. A buyer can leave either positive, neutral, or negative ratings for a seller after a transaction. Based on these feedbacks, the seller’s reputation is calculated and characterized by the total number of feedbacks and the percentage of positive ratings. Although it has been argued that the reputation system works well for eBay, it is not without drawbacks. One particular problem is the potential for “gaming” the system. For instance, it has been suggested that due to the threat of retaliation from sellers, buyers with a bad experience might prefer to remain silent instead of leaving negative ratings (Dellarocas and Wood 2008).

Despite the increasing awareness of users’ strategic gaming behavior of online reputation mechanisms and their potential adverse impacts, there are still significant gaps in our understanding of how users react to reputation systems and how reputation systems could be appropriately designed (Mascllet and Penard 2007). In this paper, we investigate a crucial aspect of the reputation mechanism design – the ability of buyers and sellers to revoke or mutually withdraw their feedbacks. On eBay, a negative feedback can be “revoked” if both the seller and the buyer mutually agree to do so. Such revoking may happen if a seller “corrects” his mistake by either replacing a previous low quality product with a better one or refunding the buyer. Alternatively, a seller could retaliate against a buyer who provides negative feedback and “force” the buyer to revoke the negative feedback. While in the former case, the revoking behavior reflects a seller’s responsibility and honesty to some extent, in the latter, the ability to revoke negative feedbacks enables bad sellers to disguise their dark pasts. In other words, a seller could be an inherently “good seller” and revoking could be a sign of a genuine mistake. Alternatively, a seller could be an inherently “bad seller” and revoking could be a sign of the seller trying to masquerade as a high quality seller. Either of these characterize a situation with adverse selection where the seller type, good or bad, is a given – a common assumption in models of reputation mechanisms. A third possibility is one of moral hazard: where the seller is able to behave

honestly, but chooses not to because of self-interest – the case of the “*ugly seller*”. We are particularly interested in understanding which of these behaviors drive revoking behaviors.

Our study is among the first to examine the strategic feedback-revoking behavior in online markets. It is also among the first study to empirically examine if such revoking behavior is due to moral hazard or adverse selection. Empirically, we take advantage of a recent significant change in eBay’s reputation system and the resulting strike. Although there have been numerous studies showing that reputation matters to sales in eBay auctions (Dewan and Hsu 2004, Lucking-Reily et al. 2007), there is little research on sellers’ reactions to changes in reputation systems. We provide the first empirical evidence that sellers do indeed respond to the design of reputation system. In the two-way reputation system (prior to the recent changes), certain sellers exhibit opportunistic behavior by revoking negative feedbacks they receive. We further find these sellers are the ones who are more likely to participate in an online strike to protest against the recent reputation system changes on eBay. Interestingly, changes to the reputation system have a significant influence on these sellers’ behavior. Most interestingly, we find that after the changes to the reputation system design (banning revoking), sellers who were more likely to revoke feedbacks earlier exert more effort into improving their reputation scores.

2. Problems with eBay’s Reputation System and the Recent Changes

One major problem with eBay’s reputation system is retaliation. eBay allows sellers and buyers to independently leave feedbacks and the feedback is available immediately to the other party. This creates incentive for one party to strategically hold back its feedback as a way of retaliation if the other party provides a negative feedback. Resnick and Zeckhauser (2002) find evidence that half of the time sellers hold their feedback to buyers, even if the sellers receive payment from buyers before buyers receive the items. Cabral and Hortacsu (2006) also find that a buyer who leaves a negative feedback for a seller has a 40% chance of getting a negative feedback in return from the seller. Consequently, due to fear of retaliation, buyers with bad transaction experience are much less likely to leave negative feedbacks for sellers (Reichling 2004; Dellarocas and Wood 2008).

The second problem with eBay’s feedback mechanism – and one that is the focus of our study - is revoking. eBay’s revoking policy states that the two parties in a transaction are allowed to withdraw their feedbacks based on mutual agreement. Researchers have argued that such a revoking policy is important in reputation system design as it helps prevent a breakdown of trust (Vasalou et al. 2008). Although the intention of this policy is to facilitate the reconciliation of any dispute in a transaction and correct errors in ratings, in practice, revoking creates a way for sellers to manipulate their reputation. A seller who receives a negative feedback has an incentive to strategically retaliate with a negative feedback to induce a mutual agreement on withdrawal (Bolton et al 2007). If indeed revoking helps convert negative ratings to nullified ratings, revoking can help bad sellers disguise their dark pasts and send misleading signals to buyers on the marketplaces. Because negative ratings are very rare (typically less than 1% of total ratings), by revoking negative ratings, a bad seller can effectively manage its reputation to be as good as, or even better than truly good sellers.

Retaliation and revoking has made eBay reputation system less effective in signaling the quality of sellers. In an attempt to address the above potential problems inherent in the design of its reputation mechanism, eBay announced dramatic changes to its reputation system in January 2008. Among other minor changes, the most significant change involved limiting sellers’ strategic gaming behavior. Starting May 2008, sellers are no longer able to leave negative

feedbacks to buyers. Furthermore, no mutual withdrawal (revoking) of the feedbacks is allowed. The changes led to several claims of the new system being unfair to sellers. As a result, a proposal to boycott eBay was initiated by sellers. A discussion thread on eBay's own forums with the title "Sign the pledge: no sales Feb 18-25!" had received thousands of posts, many expressing intentions to join the boycott. Due to this strike, eBay's number of listings was reported to decrease by more than 13% (USA Today 2008).

Undoubtedly the changes to eBay's reputation system and the online strike by infuriated sellers are two significant events. The online strike initiated by sellers serves as a natural field experiment that separates different types of sellers. The change in eBay's reputation system serves as a unique exogenous event that allows us to investigate how different types of sellers respond to the different reputations systems.

3. Data

The data was collected from April 2008 to February 2009. Table 1 presents our sample. The "general sellers" sample and "forum sellers" sample are stratified random samples based on the product category distribution of the strikers to control for product categories sold by these eBay members. To allow enough time for the new reputation system to have an impact, we define July 2008 to January 2009 as the post-change period and correspondingly July 2007 to January 2008 as the pre-change period. We collected the lifetime feedback history data for all three groups of sellers. While our primary focus is on the change to eBay's reputation system, eBay also instituted changes to its fee structure: lower listings fees and higher final value fees. To control for the potential impact of this change, we collected sellers' listing histories from January 2008 to March 2008 and January 2009 to March 2009.

Table 1. Sample Definition

	Definition/brief description	# of obs.
Strikers	Sellers who signed on the strike thread and participated in the strike	431
General Sellers	Sellers who neither had activities in the forum nor participated in the strike	3037
Forum Sellers	Sellers who had activities in the forum but not participated in the strike	2479

4. Preliminary Findings

4.1 Can Strategic Revoking Behavior Explain the Strike?

Table 2 presents the reputation profile for the three categories of sellers in the pre-change period. If we only count the positive and negative feedbacks as eBay does when displaying overall reputation profile, strikers are similar to general sellers and forum sellers. However, a close inspection of all types of feedbacks reveals that strikers have a lower positive feedback percentage and much higher revoked feedback percentage than both general sellers and forum sellers. Considering the original negative value of revoked feedbacks, we find that strikers actually have significantly more negative feedbacks ($0.21\% + 0.66\% = 0.87\%$) than both general sellers ($0.31\% + 0.13\% = 0.44\%$) and forum sellers ($0.22\% + 0.11\% = 0.33\%$). This implies that

strikers have strategically negotiated with buyers to revoke negative feedbacks so that they are similar to other sellers¹.

Table 2. Pre-Change Reputation Profile Comparison for All Feedbacks

	Score	Distribution Displayed		Hidden Detailed Distribution				eBay-Withdrawn
		Positive	Negative	Positive	Negative	Neutral	Revoked	
Strikers	267.11	99.78%	0.22%	98.72%	0.21%	0.35%	0.66%	0.06%
General Sellers (Forum Sellers)	436.03 (341.49)	99.65% (99.78%)	0.35% (0.22%)	98.96% (99.22%)	0.31% (0.22%)	0.51% (0.38%)	0.13% (0.11%)	0.09% (0.07%)
T-test	-3.41*** (-1.45)	1.40 (0.25)	-1.40 (-0.25)	-2.00* (-6.40***)	-2.38* (-0.22)	-3.97*** (-0.82)	8.37*** (12.24***)	-0.67 (1.10)

⁺p<0.10 *p<0.05, **p<0.01, ***p<0.001

Table 3. Logit Regression Analyses

	Model I		Model II	
	coefficient	p-value	coefficient	p-value
Intercept	-2.980***	0.000	-3.059***	0.000
# of Listings	1.770-e05	0.919	-2.202-e04	0.192
Powerseller Status	0.083	0.459	0.055	0.631
# of Months on eBay	0.007***	0.000	0.007***	0.000
Fee Difference	0.001*	0.012	0.001*	0.018
Reputation Score	-3.506-e04*	0.027	-4.085-e04*	0.016
Revoked Feedback Percentage			44.695***	0.000
# of observations	5946		5946	
Pseudo R ²	0.0116		0.0443	

⁺p<0.10 *p<0.05, **p<0.01, ***p<0.001

However, as we have discussed, potential financial loss under the new fee policy may also lead sellers to strike. To control for this potential explanation and other confounding factors, we employ a logit regression model to predict the propensity for strike. The results in Table 3 also suggest that sellers who strategically revoked negative feedbacks were more inclined to strike.

¹ Revoked feedbacks might also be caused by a seller's correcting his/her genuine mistake. We randomly chose 100 revoked feedbacks and found that the cases of genuine mistakes are pretty rare.

4.2 Changes in Seller Behavior: The Ugly Seller

If the strikers are indeed “bad sellers” (i.e., a case of adverse selection), we should expect to see the new reputation system more effectively reveal these “bad sellers” by their much higher negative feedback percentage. On the other hand, if sellers were indeed “ugly sellers” (i.e., a case of moral hazard) they should be able to improve their quality and improve their reputation under the new reputation system. As shown in Table 4, there is no difference in negative feedback percentage between strikers and other two types of sellers after the ban on revoking. Also, the absolute increase or the percentage increase in negative feedback percentage for strikers is much smaller than the increase for general sellers and forum sellers. This supports the moral hazard assumption. In other words, sellers (i.e. the strikers in our sample) are the “ugly sellers” - inherently good sellers who strategically choose to behave dishonestly.

Table 4. Post-Change Reputation Profile Comparison: Strikers vs. Non-Strikers

	Score	Positive	Negative	Neutral	eBay-withdrawn
Strikers	204.25	98.87%	0.90%	0.17%	0.06%
General Sellers (Forum Sellers)	402.03 (283.24)	98.88% (99.18%)	0.69% (0.59%)	0.34% (0.15%)	0.09% (0.08%)
T-test	-3.32*** (-1.68)	-0.05 (-1.26)	1.04 (1.27)	-3.56*** (0.25)	-0.67 (-0.67)

⁺p<0.10 *p<0.05, **p<0.01, ***p<0.001

4.3 Robustness check

Table 5. Reputation Profile Comparison: Revokers vs. Non-Revokers

	Pre-Change					Post-Change			
	Positive	Negative	Neutral	Revoked	eBay-withdrawn	Positive	Negative	Neutral	eBay-withdrawn
Revokers	97.08%	0.53%	1.02%	1.27%	0.10%	97.35%	1.87%	0.68%	0.10%
Non-Revokers	99.26%	0.26%	0.42%	0.00%	0.06%	99.15%	0.50%	0.29%	0.09%
T-test	-16.08***	4.74***	12.03***	18.57***	0.78	-7.98***	6.52***	6.38***	0.78

⁺p<0.10 *p<0.05, **p<0.01, ***p<0.001

We further combined all three categories of sellers and partitioned the combined sample into two groups, based on their revoking behavior. We defined “revokers” as sellers who initially received more than 4 negative feedbacks and then revoked at least 40% of these negative feedbacks. “Non-revokers” are sellers who never revoked negative feedbacks. This resulted in 249 revokers

and 2786 non-revokers. As shown in Table 5, the increase in negative feedback percentage for revokers is only 0.07% (from 1.27%+0.53%=1.80% to 1.87% whereas it is 0.24% (from 0.26% to 0.50%) for non-strikers. This implies that compared with non-revokers, revokers have changed their behavior in a positive way to mitigate the increase in negative feedbacks caused by the reputation system change. This further supports the moral hazard assumption about seller behavior. As an additional robustness check, we restricted or relaxed the definition criteria for revokers. We find that the moral hazard assumption is still supported. Defining revokers as the top 5% or 10% sellers in terms of their revoked feedback percentage produced similar results.

5. Conclusions

In this paper, we study the strategic gaming behavior resulting from the ability of sellers to revoke their feedback on eBay. We find evidence that certain sellers strategically utilize revoking to “improve” their reputation. Compared with general sellers, strikers had a much higher revoked feedback percentage. We find that the changes to the reputation mechanism instituted by eBay, has had a significant influence on these sellers’ behavior. We find evidence that they put in more effort into their transactions and receive the same level of negative ratings as other sellers. These findings contribute to the literature of gaming behavior in the online markets and online reputation system design.

References

- Bolton, G., Greiner, B., and Ockenfels, A. “Engineering Trust: Strategic Behavior and the Production of Reputation Information,” *Working paper*, Harvard Business School, 2007.
- Cabral, L.M.B., and Hortacsu, A. “The Dynamics of Seller Reputation: Theory and Evidence from eBay”, *CEPR Discussion Paper No. 4345*, April 2004.
- Dellarocas, C. “Reputation Mechanism Design in Online Trading Environments with Pure Moral Hazard,” *Information Systems Research* (16:2), 2005, pp.209-230.
- Dellarocas, C., and Wood, C.A. “The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias,” *Management Science* (54:3), 2008, pp. 460-476.
- Dewan, S., and Hsu, V. “Adverse Selection in Reputations-Based Electronic Markets: Evidence from Online Stamp Auctions,” *Journal of Industrial Economics* (52:4), 2004, pp. 497-516.
- Masclot, D., and Pénard, T. “Is the eBay Feedback System Really Efficient? An Experimental Study,” *Working Paper*, University of Caen, 2008.
- Reichling, F. “Effects of Reputation Mechanisms on Fraud Prevention in eBay Auction,” *Working Paper*, Stanford University, 2004.
- Resnick, P., Zeckhauser, R., Swanson, J., and Lockwood, K. “The Value of Reputation on eBay: A Controlled Experiment,” *Experimental Economics* (9), 2006, pp. 79-101.
- Resnick, P., and Zeckhauser, R. “Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay’s Reputation System,” *The Economics of the Internet and E-Commerce*, Elsevier Science, 2002, pp. 127-157.
- USA Today. “*Online auction listings down 13% in boycott of eBay*,” February 24, 2008 (available online at: http://www.usatoday.com/tech/techinvestor/corporateneews/2008-02-24-ebay-strike_N.htm)
- Vasalou, A., Hopfensitz, A., and Pitt, J.V. “In Praise of Forgiveness: Ways for Repairing Trust Breakdowns in One-Off Online Interactions,” *International Journal of Human-Computer Studies* (66), 2008, pp. 466-480.