

Buying Reputation on eBay¹

By Federico Dini² and Giancarlo Spagnolo³

First draft: January 2006

This version: September 2007

QUADERNO CONSIP VIII - October 2007

Available for download from www.consip.it and www.gianca.org

Abstract: Electronic feedback mechanisms allow trading partners to express a public rating on each other's performance after they transact. eBay argues its feedback mechanism is key in fostering honest behaviour among anonymous trading partners, and analogous mechanisms are being implemented in many (private and public) electronic trading platforms around the world. Reliability of information produced by these systems, however, is both crucial and hard to guarantee. In this paper we highlight a statistical puzzle in terms of feedback behaviour on eBay's and discuss the forms of feedback manipulations discussed by literature as possible sources of the puzzle. We then find that positive feedback can easily be 'bought' on eBay against cash, and argue that this 'feedback trading' on a 'market for positive feedback' could be a source of the identified puzzle. We then report results from field experiments we conducted to identify the minimum monetary price at which unwarranted positive feedback can be bought, and derive potential implications of this "market for feedbacks" on the other markets present on eBay.

Keywords: eBay, electronic platforms, feedback mechanisms, feedback manipulation, online reputation, seller reputation, scammers.

¹ We are grateful to Chris Dellarocas, Tobi Klein, Chris Lambert, Konrad Stahl, Lucio Picci, our colleagues at Consip Research Unit, graduate students in a lively reputation seminar at the University of Mannheim in 2004, and seminar audiences at the Stockholm School of Economics, Consip and the II International Public Procurement Conference in Rome (September 2006) for helpful comments or stimulating discussions. All remaining errors are our own.

² Research Unit, Consip S.p.A. (the Italian Public Procurement Agency). e-Mail: federico.dini@tesoro.it. URL: <http://www.consip.it/on-line/Home/Studieprogetti/UfficioStudi.html>

³ Consip Research Unit, University of Rome 'Tor Vergata', SITE-Stockholm School of Economics and CEPR. E-Mail: giancarlo.spagnolo@uniroma2.it. URL: www.gianca.org

1. Introduction

In many online marketplaces, including eBay, Amazon, and Yahoo, sellers' opportunism is a major concern for buyers, as legal enforcement is absent or too costly. Moreover, transactions with a specific seller is typically occasional for a specific buyer, they are not embedded in a repeated bilateral relationship so that expected gains from future cooperation cannot mitigate the parties' temptation to behave opportunistically.

Additional concerns arise from the relative anonymity and the geographical dispersion of trading partners that make transaction disputes much more complex and costly than in traditional "face-to-face" frameworks. Trading in online platforms can be particularly risky for buyers because shipment occurs only after payment, so the seller is tempted to cheat and the buyer can check the quality the item "too late", when there is little to do in case of dissatisfaction.

Feedback (or reputation) mechanisms, by allowing partners to rate each other for their performance in the transaction and posting ratings near each party homepage on the platform may mitigate informational asymmetries and opportunism, fostering trust and trade. The idea is to stimulate traditional reputational forces, crucial in any market, by using cheap electronic instruments to collect, aggregate, and disseminate information about individual traders' past performance across the community of potential future trading partners.

The essence of any reputation system, however, is the reliability of the information it produces. Ratings (feedback) and the reputation indicators they generate (feedback scores) will improve efficiency of transactions only if they contain at least some "truthful" information. Unfortunately, electronic transmission of information involves

anonymous individuals (essentially strangers) who cannot assess the reliability of information/feedback received by trusting the source of the information (Resnick and Zeckhauser, 2002).

Unreliability of information also stems from the “soft” (and thus unverifiable) nature of the feedback content, which makes it practically impossible to identify and punish untruthful reports. “Hard” information, like measures of performance that can be verified by third party (e.g. timing of payment through credit card, timing of delivery registered in the delivery bill), may be contracted upon and misreporting of this information can be demonstrated in a court and sanctioned. On the contrary, any measure of buyer's satisfaction is soft, since it relies on a subjective judgement, and therefore is more subject to falsification.

Indeed, as online reputation systems and analogous electronic carriers of soft word-of-mouth information (e.g. [CNET](#), [Epinions](#), [Citysearch](#), [Moviefone](#)) become of common use, influencing substantively consumers' behaviour, the incentives to manipulate soft information become ever stronger. There is quite some evidence already that dishonest electronic traders (in this case also known as “scammers”) attempt to acquire quick and undeserved reputations, or even to spoil the reputation of competitors, facilitated by the low cost of submitting online feedbacks and by the anonymity of trading partners.⁴

In this paper we start from observing a puzzling asymmetry in terms of feedback behaviour on eBay, as recorded by Resnick and Zeckhauser (2002). Most positive feedbacks (or “positives”, henceforth), the large majority of feedbacks given on eBay, are reciprocated, i.e. are met by a corresponding positive feedback from the trading partners. Negative feedbacks, instead, are very often not retaliated, even

⁴ For example, Cabral and Hortacsu (2006) show that sellers often start accumulating feedbacks by acting as buyers of lower value items; Harmon (2004) reports that when in February 2004, because of an error, Amazon.com's Canadian site revealed the identities of some book reviewers, it turned out that many rating were written by the books' publishers, authors, and competitors. See Mayzlin (2006) and Dellarocas (2006) for more on this.

though matching a negative with another negative feedback may deter future partners from rating negatively and opens the possibility to mutually withdraw both ratings.⁵

We then briefly review the various forms of online feedback manipulations already pointed out in the literature to see if these can be the source of the puzzling asymmetry: if you can acquire positive feedbacks through manipulations, then it can be that the manipulation method brings about both the very high percentage of positive feedbacks and the relatively much higher rate of reciprocation of positives we found puzzling.

Continuing our inquiry in the field (i.e. searching eBay) we found out that it is has become rather easy to ‘buy’ positive feedbacks in fake transactions, and that this typically lead to both side of the transaction receiving a positive. Because in this ‘market for positive feedback’ (described more in detail below) both parties in a transaction typically obtain a positive feedback, this form of manipulation tend to inflate both the amount of positive feedback and the frequency in which a positive feedback is reciprocated.

To improve our understanding of this ‘market for feedbacks’ and to quantify the minimum price at which feedback can be bought we conducted field experiments trying to obtain positive feedback at an ever lower price. As we will describe in detail, we found that positive feedbacks could be ‘bought’ for as little as **0,12 USD**.

We conclude by thoroughly discussing the likely implications of this “market for feedbacks” on the functioning of other markets present on eBay, in the light of most recent research on the effects of trading names and reputations and of posting fake recommendations in online opinion forums.

⁵ Both parties must agree to withdraw their ratings on eBay; see Klein et al. (2007) for details of eBay’s feedback mechanism and the recent changes it undertook.

It is worth noting that ratings' manipulation can potentially hit any market or institution producing "soft" evaluations. In most supervision, rating, auditing and certification relationships, the agent (the supervised, ratee, auditee, etc.) may collude with the principal (supervisor, rater, customer, auditor) to get a better assessment than deserved. Feedback purchase in eBay is little different from a collusive agreement between a rating agency and a rated firm to misrepresent the firm behaviour, hence the discussion in the paper has a more general value than one could think.⁶

The rest of the paper is organized as follows. Section 2 discusses a puzzling asymmetry in feedback behaviour that seems to have gone unnoticed in the literature. Section 3 reviews the forms of feedback manipulation highlighted by the existing literature and the possibility that they are at the origin of the puzzle, and note the emergence of 'markets for feedbacks'. In section 4 we present the results of our field experiments, which provide evidence on how easy and cheap it is for any (honest or dishonest) user to collect artificial feedback on eBay. In section 5 we discuss likely implications of such a 'market for feedback'. Section 6 draws the conclusions.

2. Testing a puzzling asymmetry in the behaviour of feedback providers

Friedman and Resnick (2002) show that correlation in feedback provided in eBay by traders is high and positive. However, they appear not to note that in their data set behaviour is asymmetric in terms of reciprocation/retaliation: while reciprocation is strong for positive feedbacks, retaliation is quite weak for negatives.

⁶ See Baliga (1999), Dini and Spagnolo (2004), Lippert and Spagnolo (2006), and Peyrache and Quesada (2005) for recent contributions on the topic of (soft) information manipulations or transmission. Bar-Isaac and Tadelis (2007) offer an easily accessible general survey of the reputation literature in economics; for a more detailed and technical treatment see Mailath and Samuelson (2006).

We now investigate whether this puzzling asymmetry emerging from the data collected by Friedman and Resnick (2002) is statistically significant. For this purpose we run 4 sets of tests to verify statistical significance of the:

- Set 1: single reciprocation coefficients;
- Set 2: unconditional reciprocation difference (of positives - negatives);
- Set 3: conditional reciprocation difference (positives - negatives) with respect to both parties providing the feedback;
- Set 4: conditional reciprocation difference with respect to seller or buyer providing the feedback.

The first set of tests aims at validating the significance of single coefficients. The second and the third test are the key tests, validating the significance of the conditional (to providing feedback) and unconditional reciprocation difference. The last set of tests accounts for seller vs. buyer providing the feedback.

Data and statistical method

We exploit the data-set used in Resnick and Zeckhauser (2002), which consist in 36,238 transactions from eBay in the period February 1 - June 30, 1999.

The nature of the data suggest the use of non-parametric statistics. Reciprocation coefficients can be straightforwardly interpreted as correlations, as they provide the probability to get a 1,0,-1-feedback given that the opponent provided a feedback of the same sign.

Choosing the right statistical test is quite easy. In general parametric tests are recommended when the distribution of the variable in the population is known and in particular when it is Gaussian. Non-normality strongly suggests non-parametric statistics. Since we have no information about the population and the sample distribution of reciprocation coefficients, then we should restrict our attention to non-

parametric statistics which abstract from the sample distribution of the variable of interest.⁷ In our case non-parametric approaches are supported by three further elements. First, data are categorical, in the sense that they are just a pure score. When the outcome variable is a categorical, non-parametric tests are recommended. Second, since “positives” are extremely dominant over “negatives” – especially when both parties provide the feedback – our data are likely to be drawn from a non-Gaussian distribution. When normality is violated, the use of parametric methods is likely to yields biased results. Third, we have a large sample.⁸ Large data sets face no particular problems with respect to the choice of parametric vs. non-parametric tests: by the Central Limit Theorem non-parametric are as powerful as parametric ones in large sample, regardless of the variable distribution.

Specification of tests

To verify the statistical significance of single coefficients (set 1) we use the following t-student test:

$$t = \frac{P}{\sqrt{[(1-p^2)/(N-2)]}},$$

where p is the correlation and N the number of observation. The test rejects the null hypotheses $H_0: p=0$ against the alternative hypotheses $H_0: p \neq 0$ if the computed value is greater than 1.96 at 5% level of confidence.

This test is usually adopted to check the validity of the non-parametric Spearman correlation coefficients. The proportions we compute are interpretable as correlation coefficients, namely the proportion of time (or probability) that a certain feedback is reciprocated by two trading partners. The P test appears appropriate to check whether

⁷ More precisely, non-parametric methods do not rely on the estimation of parameters (such as the mean or the standard deviation) describing the distribution of the variable of interest in the population. These methods are also, and more appropriately, called parameter-free methods or distribution-free methods.

⁸ In these kinds of statistics large samples contains more than 50 observations.

single reciprocation coefficients are statistically significant. In order to verify whether the difference between positive and negative reciprocation coefficients (sets 2-4) is statistically significant we use the following proportion test:

$$z = \frac{p_1 - p_2}{\sqrt{p(1-p)\left(\frac{1}{N_1} + \frac{1}{N_2}\right)}}$$

where, z follows a standard Gaussian distribution, and p_1 and p_2 are the proportions of two different samples to be compared with the overall proportion p , computed as the two samples were pooled into a larger one. The test z is distributed as a standard Gaussian variable. This test is specified by the following set of hypotheses: $H_0: p_1 = p_2$ against $H_1: p_1 > p_2$.

In our case we set the following hypotheses: $H_0: p_- = p_+$ against $H_1: p_- > p_+$. The choice of H_1 , which determines a one-tail test, is motivated by the fact that correlation between positive feedbacks is much larger than the one of problematic (negative plus neutral) feedbacks. The test rejects H_0 at 1% and 5% level of confidence, if computed values are greater than $z=2.33$ and $z=1.64$, respectively.

Results

The tests confirm that the puzzle is statistically significant (see Table 2).

These are main results:

Single coefficients. Reciprocation coefficients are all significant. Computed t values range 18-103, considerably larger than the threshold value 1.96 at 5% level of confidence.

Differences. The difference between positive and negative coefficients are significant as well. Unconditioned differences are statistically significant, ranging 16.8-24.9. The

value 16.6 refers to the difference between 0.814-0.31, while 24.9 refers to the difference between 0.701-0.301. Conditioning to both parties providing the feedback increases the significance to the range 63.2-34.2.

3. Manipulation of online feedback mechanisms

The strong correlation between positive feedbacks may be due to generalized good performance that pushes traders to leave a high frequency a positive feedback. But this correlation may also be due to simple reciprocation, or to other manipulations of the type we will soon illustrate. Analogously, the weak correlation for negatives may hide the fear of retaliation, or the lack of manipulations that induce asymmetry and negative feedbacks (who should want them?).

The issue of manipulations of online feedback mechanisms has been only recently analyzed by the economic literature.⁹ Let us briefly summarize some of the most common forms of manipulations: i) *fake messages* in online discussion forums: firms invest in false messaging to increase the reputation of their products at the expense of the one of competitors ii) *feedback reciprocity*: users post a reciprocal positive feedback at the end of the transaction to merely boost reputation. iii) *unfair ratings*: users post positive or negative ratings to promote partners or handicapping competitors. iv) *identity changes*: people re-enter the marketplace with a new profile after experiencing negative feedback that compromises reputation.

3.1 False messaging in promotional chats

Recommendation web sites, chat rooms and customer review sections allow consumers to overcome geographic boundaries and to communicate for mutual

⁹ See e.g. Klein et al. (2006) and Dellarocas and Wood (2006). Related work by Mayzlin (2006) and Dellarocas (2006) discusses problems of manipulation on opinion forums.

interests. In recent years, online forums have been growing in popularity and have become an important component of portals, e.g. Yahoo! Movies, traditional news sites, as BBC Talking Point, consumer report sites, as CNET, Epinions, Citysearch, Moviefone, and online retailer web sites, e.g. Amazon, and Circuit City (see e.g. Dellarocas, 2006). These online discussion forums, where people exchange impressions and information about goods and services, are today a key resource for both consumers and firms. The former can learn quickly about new goods/services available on the market and benefit from previous consumers' recommendations. The latter, can easily and effectively promote their products, reaping the benefits from effective interactivity among consumers and less regulation of online promotions with respect to traditional advertising (Mayzlin, 2006). However, exploiting the anonymity enjoyed by participants of these online communities firms can easily disguise their promotion as consumer recommendations: firms can infiltrate in consumers' discussions to supply promotional messages and influence the judgments of their products. Since agents' identities are rather “nebulous” in the virtual world, it can be quite difficult, especially for first-time consumers, to distinguish between truthful recommendations and strategic false messages.

3.2 Unfair ratings

Unfair ratings aims at inflating the reputation of partners and/or destroying the one of competitors. This phenomenon may arise in any online marketplace in which users (as buyers or previous consumers) can rate other users (sellers or firms promoting their products). This is the case of Amazon, eBay, Yahoo, but also consumers report sites as Epinions and Citysearch mentioned in the previous section. For instance, any user (buyer or seller) can post a voluntary rating to a book on sell on an Amazon e-shelf. It is not necessary to buy a book on Amazon to leave a feedback on that book. A group of badly-intentioned users may coordinate to boost the reputation of one book or to compromise the one of others, facilitated by the possibility to post the

rating for free. Another example is eBay. On eBay, one seller may in principle enter the community with a secondary false account, bid for an item of a competing (partner) seller and post him a negative (positive) unfair feedback. He can then vanish without paying for the item and continue to sell his products with his primary account.¹⁰

3.3 Identity changes

Manipulations of online feedback mechanisms may take the form of identity change. One key element of reputation on the Internet that does not typically arise in non-virtual settings is the possibility for e-traders to easily change their identity. Identity change only requires few minimal keyboard operations. One seller faces with many alternatives: i) interacting fully anonymously (by changing identity every time), ii) maintaining the same identity for a given period, or iii) for all his trading life. These options make the transmission of reputation information a strategic variable, controlled by each player. This introduces opportunities to misbehave without incurring in reputational consequences. For instance, in eBay one seller may cheat buyers in first transactions, and in case a negative feedback he can easily “clean” his reputation profile by re-entering the community with a new identity. Some solutions to the problem are discussed by the recent literature. Among them: i) straightforwardly, eliminating anonymity or avoiding name changes; ii) requiring entry fees; iii) assigning an “introductory” low reputation to new market entrants.¹¹

¹⁰ Dellarocas (2000, 2003) defines unfair high ratings (also known as “ballot stuffing”) as when a seller colludes with a group of buyers in order to be given unfairly high ratings which in increase selling prices. Unfairly low ratings (also known as “bad-mouthing”) is when one or more sellers collude with buyers in order to “bad-mouth” other competing sellers. He argues that the effect of such unfair ratings is the dispersion of ratings for a given seller. In particular, since buyers may have different perceptions of the same seller’s performance, it will be rather difficult or even impossible to distinguish (fair) ratings due to such differences from (unfair) ratings.

¹¹ See, for instance, Dellarocas (2005), Cheng and Friedman (2005), Friedman and Resnick (2001) and Resnick and Zeckhauser (2002). One interesting solution are rewards/fees-based systems (Avery et. al. and 1999, Miller et. al. 2005, Papaioannou and Stamoulis 2005, Jurca and Faltings 2004 and Dellarocas 2005, 2006). The idea is that feedback referred to one seller tend to be correlated (good sellers will get on average positive evaluations, bad sellers will instead get negative ones) and any feedback that is much different from the others should be considered unfair and thus subject to penalty

3.4. Reciprocity

Reciprocity occurs when trading partners implicitly or explicitly agree to post a reciprocal positive feedback at the end of the transaction although none of them performed particularly well. It arises in two-way mechanisms (as the one of eBay) where transacting parties are allowed to rate one another at the end of the exchange. In theory, bilateral rating systems are used to induce both parties to fulfil their obligations. In practice, however, they may produce adverse effects on the informational content of posted feedback for the effect of reciprocation (and retaliation) arising when users can respond to the received feedback. The empirical research highlights that reciprocity is a relevant concern. In eBay it is common practice for traders to negotiate welfare improving actions before leaving negative feedback. This might also induce traders to report only extremely bad performance.

Resnick and Zeckhauser (2002) confirm that the two-way nature of the eBay's mechanism provides additional inducements of positive feedback. They prospect that leaving a positive feedback is a dominant strategy in two-way mechanisms. Dellarocas, Fun and Wood (2004) argue that the expectation of feedback reciprocation increases the participation to reputation systems. Knowing that they will be reciprocated with high probability, people are encouraged to leave a positive feedback. Chwleos and Dhar (2005) confirm the results of Dellarocas, Fan and Wood (2004) when comparing the Amazon and the eBay's feedback mechanisms. They attribute to the nature of the mechanism (one-way of Amazon vs. two-way of eBay) the differences in the behaviour of feedback providers in the two systems. In particular, the eBay's mechanisms users post and reciprocate positive feedback with significant higher frequency with respect to Amazon. These results clearly suggest that turning to a one-way feedback mechanisms would greatly reduce reciprocation. This solution is supported by Dellarocas, Dini and Spagnolo (2006) who address the issue focusing on public e-procurement platforms.

All these papers suggest that reciprocation is a concern for eBay-style feedback mechanisms. To counter this problem, Klein et al. (2006, 2007) suggest to introduce a blind period with fixed deadline, after which feedbacks cannot be left or changed anymore (no more ‘mutual withdrawal’ possibilities), and during which whether and what feedback was left by a partner cannot be observed.

3.5 Scamming the eBay reputation system

The problem of false reputation seems to hit particularly the eBay's feedback system. But how do scammers practically manipulate eBay's feedback system? We discuss below some of these ways.

Feedback auction. Scammers can transparently auction the feedback, saying “I sell positive feedback”. Then the winning buyer gives the seller positive feedback in exchange of a positive feedback from him. A more covert way to auction the feedback is listing the words “positive feedback” or “feedback exchange” in the title of the auctioned item. In so doing the seller can disguise a feedback exchange with a real transaction. Item titles may also be of the type “free stickers for leaving positive feedback”, “new recipe gets positive feedback from all”, or “the most positive way to buy pre-written feedback”. Titles always contain the words “positive feedback” to signal other scammers that positive feedback can be easily obtained in that auction. The buyer then places his bid for the item, and both traders get positive feedback without completing the transaction.¹²

Buying and re-selling feedback. Sometimes scammers purchase and re-auction low-value things like recipes, e-books, wholesale lists, free information and information

¹² This means that there is no payment from the buyer to the seller, although this pays the fixed fees for listing of the item.

booklets. For instance, you can buy on eBay an e-book titled “get 100% “positive feedback quick” for less than 1\$. Any user can buy and resell the book as many times as he desires to collect positives.

Multiple accounts. This is a very simple way for dishonest traders to obtain positives quickly. Basically, they setup multiple accounts on eBay and sell multiple items. Then they purchase the items and leave themselves positive feedback. There is evidence that in order to appear reliable and honest sellers, scammers even create more than 100 accounts.

Feedback Theft. This is when somebody takes a feedback that does belong to someone else. In other words some users can take over the control of another user’s personal eBay account. The scammer then exploits the stolen eBay seller’s good reputation to trade with other sellers.

4. eBay’s “market for feedbacks”: simple field experiments

The list above in itself does not provide a clear confirmation of our conjecture that manipulations could be at the origin of at least one of the determinants of the puzzling asymmetry highlighted in section 2. However, checking the various scamming techniques on eBay led us to discover a real ‘market for positive feedback’, where some agents openly state they ‘sell positive feedbacks’, or post fake or symbolic objects for sale - like an electronic one-page book with the most standard recipe for chocolate cake - at a similarly symbolic price, just to be able to exchange a positive feedback at the minimum cost (these offers can easily be found searching for the keyword ‘positive feedback’).

This section describes and report results from some experiments we made on that market in the period October-December 2005 to test how viable these ways to collect positive artificial feedback are.¹³ In particular, the aim of the experiments is to estimate the minimum price of a feedback, and thus the potential value of “marginal reputation”. All experiments were conducted in the fall-winter 2005 through an account created on eBay with nickname “**Convettore100**”.

Buying and re-selling

On October 16th 2005 we bought on eBay the e-book “How to earn up to 100 Feedback” for a price of €0,80. The very short book (just 2,5 pages!) tells why users should get feedback quickly and describes some methods to collect many feedback quickly and at a low cost. Some hours later we paid the book and the seller immediately gave us a positive feedback, leaving the comment “Quick payment, Thanks”. Two days later we posted our positive feedback to the seller leaving the comment “Fast shipment, thanks”. Since the book was directly e-mailed to us, we were charged no additional shipping costs. Therefore, the cost for collecting this positive feedback was just the price of the book: €0,80.

On October 30th we re-auctioned off the book for the price of €¹⁴ but no bid has been placed for it. However, on December 15th an user contacted us proposing to exchange a feedback without completing the transaction for the book.¹⁵ Since we received the message after the closing time of the auction, we had to re-auction the item to allow the user to place his bid. Once the bid was placed, we immediately sent our positive feedback to him with the comment “All perfect, reliable eBay”. He return the positive feedback on December 30th the with the comment “+ ok +” (see Figure 1). The cost of collecting the feedback by re-selling was €0,6. This cost is

¹³ We do not consider feedback theft because illegal. Reciprocation is essentially similar to the other legal methods, but it may occur after a non-artificial transaction is completed by the parties. Then, it is in general much more costly than any other method.

¹⁴ The reserve price for such items is usually no greater than 2 Euro.

¹⁵ On eBay transactions can be made fully artificially. Traders may agree to exchange the feedback without performing a physical exchange of money for the item. However, as mentioned before, the transaction is not completely “neutral” for the seller since this pays eBay a fixed fee proportional to the awarding value of the item.

computed as follows: €0,1 is for listing the item, €0,2 for giving the item particular visibility on the market. Since we auctioned the item twice the total fixed cost was €0,6.¹⁶

On average the buy and re-sell method determines a cost of $(0,6+0,8)/2=€0,7$ per positive feedback.

Multiple accounts

On December 12th we auctioned-off again the object bought on October 16th. We then placed our bid with another false account. However, the bid was declared invalid by eBay.¹⁷ Therefore, that account became unavailable to continue our experiment. If this system worked out, we would have incurred in a cost of €1,59 to collect one positive (the cost is computed as follows: €0,1 for listing the item, €0,2 for giving more visibility to the item and €1,29 for credit card fees). But eBay was careful enough to avoid this.

Auctioning the feedback

On December 26th we collected another positive feedback by auctioning the feedback. Once the positive feedback had been listed we contacted the same user with whom we exchanged the feedback on December 15th and proposed him an additional transaction. The user accepted. After he placed his bid we immediately closed the auction and awarded the object (this was on December 28th). Once we closed the auction we sent our positive feedback to him with the comment “Fast and reliable”. He returned use the positive on December 30th the with the comment “Great seller” (see again Table 1). Auctioning the feedback costs us €0,12. The cost is composed as follows: €0,1 is the cost for listing the item, €0,02 is the transaction fee.

¹⁶ €0,2 could have been saved by renouncing to give the item particular visibility to the market.

¹⁷ In the message we received from eBay there was written that personal data inserted were not matching with the user's profile.

Summary of the results

The results of our experiments are summarized in table 1. On the basis of our experiments, the cheapest way to collect artificial reputation on eBay is *buying feedback*: the cost is €0,12.¹⁸

Method	Auction Data		Cost Data		
	Reserve price	Awarding price	Fixed fees	Transaction fees	Actual cost
Auction	0,1	0,51	0,10	0,02	0,12
Buy and resell	<i>Buy</i>	-	0,00	-	0,80
	<i>(Re)sell</i>	0,01	0,01	0,30	-

Structure Fee on eBay: fixed fee for listing items with a reserve price of less than 1.99 Euro are 0.1 Euro. An additional fee of 0.2 Euro is charged if the seller wants to insert more details on the object listed. Transaction fees are 4,5% up to 50 Euro.

5. Discussion: the likely effects of the ‘market for feedback’

The previous simple field experiments shows how easy and cheap it is to manipulate the feedback record buying positives on eBay’s ‘market for positive feedbacks’. Because transactions on such market imply reciprocated positive feedbacks, that market could be one determinant of the puzzling large number of reciprocated positives discussed in Section 2. In this section we discuss the likely consequences of such fake reputations on other, ‘normal’ markets active on eBay.

In principle, when reputation indicators are driven by information manipulation rather than “good behaviour” they may poorly predict future performance and may then distort the decisions of unaware people using the reputation information (for instance when deciding whether and which amount to bid in an eBay auction). However, understanding the impact of such manipulations is not so easy, as we need

¹⁸ All purchases and sells we made were performed by credit card. Cost estimate do no include credit card fees (which is usually about €1). However, the per-transaction fee is lower the larger the number of transactions. Sellers willing to acquire a reputation by purchasing or (re-)selling feedback are expected to trade much in the market so credit card fees would become negligible. Notice also that (re-)selling the feedback costs more than buying it, since any (re-)sell implies a fixed listing fee.

to account for a number of key factors. In particular, we must try to answer the following questions: 1) Can buyers distinguish at low cost between fair and bought feedback? 2) How can vendors build up a high feedback score? 3) What is the cost of high building it? 4) What types of sellers (honest or dishonest) is more interested in acquiring a high feedback score?

Can buyers distinguish between fair and unfair report?

Is it possible for buyers to distinguish between fair and unfair reports received by vendors? If this was possible at low cost buyers could do a first screening among vendors. Although the buyer does not know how the vendor will behave after achieving a high feedback score, he may reduce the risk of unsuccessful transaction renouncing to bid upon observing that the vendor's feedback score is determined by manipulations. On eBay, the vendor's feedback history provides detailed information about all the items he sold, as the awarding value, shipping costs, etc., so a buyer can potentially check if the vendor's high feedback score is driven by feedback trade rather than good behaviour. However, such detailed information are made available to the public only for 90 days, so that buyers can screen vendors only for a limited period.

How can vendors build up high feedback scores?

We discussed three main ways by which vendors can build up high feedback score. One is, of course, effort and good performance, helped by reciprocation. Another is buying unwarranted positive feedbacks on the 'market for feedbacks' that we discovered. The last method is "buy first and sell later", identified by Cabral and

Hortacsu (2006) as a rather popular strategy.¹⁹ The point worth noticing is that in either case, improving is costly. In the first way, the vendor exerts costly effort, i.e., he delivers the promised good “renouncing” to deliver another inferior, but less costly good. In the second method, he does exert no effort, but he pays some cash and loses time (although the unitary cost of feedback is low, the vendor needs to invest a minimal amount of money and time to collect the number of positives sufficient to generate some “reputation effect”). In the third method, the vendor fairly spends money to purchase low value items.

At first glance we could be tempted to conjecture that the ability to purchase feedback on eBay, by making unreliable the information contained by feedback scores, inevitably destroys the vale of a feedback system. Such a conjecture, however, would be premature and most likely wrong. The reason is that as long as reputation is costly to obtain, the fear of losing it could in principle discipline future behaviour independent of how it was obtained. This is what makes the exact price/cost of buying reputation interesting. Note that in *all* the three circumstances above the vendor does incur a *cost* for building up and/or improving his score. This is object of further discussion later when we address the issue of who wants high reputation.

What if the cost of positive feedback was zero?

Our experiments show that vendors incur a positive cost to purchase positive feedback. It is useful to think first – as a benchmark – to the consequences of manipulations if they led the cost of collecting unwarranted positive feedback to zero? Given that any agent – good or bad – can always obtain a positive feedback score, in such situation positive feedback have no value unless some other agents are

¹⁹ Since it is much easy to collect positive feedback after a purchase, some vendors purchase at the beginning and then switch to selling after acquiring a sufficiently high score.

naïve or poorly informed. However, negative (or neutral) feedback would still have value. Indeed, with costless positive feedback acquisition only negative feedback would have informational content.

This is an extreme result, but it clearly extends to less extreme situations: negative feedbacks tend to be much more informative than positive as long as positive feedbacks are obtainable at low cost, e.g. through reciprocation, and this is probably the (good) reason why eBay publishes the number of negative feedbacks separately. Note that several online feedback mechanisms allow ratings in a positive scale and define the feedback score as the average of these ratings. Examples of this kind are Amazon or Yahoo!, in which users rate from 1-5 or 1-10 respectively, and the score is given by the average of ratings. These metrics create two problems. First, they do not allow for strictly negative feedback, as eBay instead correctly does, therefore preventing bad transactions to be effectively punished. Most importantly, with average-based measures negative feedbacks are “blended with” positive, so the impact of a single negative feedback is much more limited.

Who wants a high reputation?

One important point is that vendors may want to “buy” positive feedback to either foster profitable but honest future business, or to undertake profitable but dishonest business with future buyers. Which of these two situations is more likely to occur? This may depend on “who” wants high reputations, that is, whether good/honest or bad/dishonest vendor are more willing to invest in acquiring a costly reputation.

Mailath and Samuelson (2001) theoretically address the question of what type of firm (good/competent or bad/inept) is willing to buy what ‘type’ of reputation

(high, medium, or low) when reputation for quality is on sale. They conclude that high reputation attracts both good and bad firms – as the former can easily maintain high reputation whereas the latter can immediately exploit it. They also find that intermediate (or average) reputation only attracts good firms because they can boost it by exerting effort, whereas for bad firms that is not valuable. The level of reputation bought solely by good firms is the intermediate one. The result appears strictly related to our question, even though its interpretation is not easy: what is an ‘intermediate’ reputation on eBay, and how do we define ‘buying an intermediate reputation’ when agents can only buy one feedback after the other? One possible solution is perhaps to interpret the result in terms of *when does an agent stop buying feedback and starts selling*. If he or she turns relatively early to normal selling, it should be more likely that he is a good quality seller that will behave well in order to further increase his feedback score and become an established high quality seller. If, instead, the trader bought a relatively large number of positive feedback, then he or she will be more likely a bad quality seller that plans to milk the acquired reputation by cheating on future buyers.

Tadelis (1999) investigates a similar problem in a theoretical model on the market for “brand names”, where firms differ in terms of intrinsic quality (so that adverse selection is the problem, while there is no room for opportunism moral hazard), and brand names are little else than reputation indicators. He finds that as long as name sales are unobservable to buyers, there are no circumstances in which only one type of firms (good or bad) buys only one type of brand name (good or bad).²⁰ In particular, good firms value good names more than bad firms because it is easier for them to maintain the name. However, it is more difficult for bad firms to build a name, thus they find names more valuable than good firms at the beginning. Because of these conflicting forces, competition for brand names (for good reputations) between good quality and bad quality firms does not bring to a full separation

²⁰ The result appears robust to changes in the modeling environment; for example, in Tadelis (2002) an analogous ‘no separation’ result is obtained when both moral hazard and adverse selection are present. See also Tadelis (2003).

between good and bad firms. If it would, there would be no way for buyers to update in a negative direction a high reputation, but this would make brand names much more valuable for bad firms planning to ‘milk’ the high reputation by providing poor quality at high prices.

Tadelis’ results are suggestive, they appear to imply that as long as agents’ or firms’ intrinsic but unknown quality is the important issue, so that adverse selection is the main problem, a market for reputation constrains the ability of a reputational forces to separate those types. On the other hand, in situations where agents’ choice whether to behave opportunistically is the crucial problem, a market for reputation may in fact be less problematic (then the reputation is just a valuable bond that increases the cost of opportunism, independent how it is gained, as in Klein and Leffler 1981), and may even help discipline opportunism when a firm plans to exit a market, as shown in Tadelis (2002) and Bar-Isaac (2006).

These results, however, cannot be directly applied to our context because in eBay’s market for positive feedbacks there appears to be no role for competition for positive feedbacks that could drive up the price of feedbacks, as it happens in Tadelis model for the price of brand names: anyone can buy feedback at very low price in eBay’s ‘market for positive feedback’ – acting either as seller or as buyer, depending where the short side of such market is – as positive feedbacks can be produced at extremely low cost in unlimited amounts. Anyone appear able to start selling and buying positive feedback in eBay’s ‘market for positive feedback’ at the constant minimum price linked to transaction costs plus the time loss of participating in that market.

The results in Mayzlin (2006) and Dellarocas (2006) about the informativeness of online forums in the presence of false messages may also shed some light on our question. Mayzlin (2006) is the first to look at the likely effects of false messages posted on electronic opinion forums. In particular, she investigates whether recommendations remain credible when firms can post fake recommendation for their

products, and to which type of good, superior or inferior, they exert more promotional effort. For this purpose, she builds up a theoretical model where two competing firms hold private information on the quality of their products. The firms send recommendations to influence consumers' purchasing decision. Previous consumers who have experience about products also leave online recommendations. The online discussions then result in a mixture of truthful recommendations from consumers as well as promotional activities from interested firms, with the consumer being unable to distinguish biased from unbiased recommendations. Mayzlin's most important finding is that, if the ratio profits/cost of manipulations is sufficiently large, there exists an equilibrium in which both firms manipulate (i.e. send promotional/fake messages), but overall, the online forum remains informative. Consumers correctly follow the online advices, benefiting from the interactivity of discussions, typically arising in online forums and absent in standard advertising, that allows them to better learn the quality of products. However, consumers face a higher positive probability of making a wrong choice with respect to the scenario with no fake promotions.

Dellarocas (2006) extends Mayzlin (2006) finding the counterintuitive result that manipulations may either increase or decrease the informativeness of the forum. When the number of users posting honest recommendations is small enough, then manipulations become beneficial to the forum. The basic intuition is that since promotions are costly, then they are more intensively undertaken by good firms. Hence, good recommendations (i.e. false messages regarding to good products) will be larger and more frequent than bad recommendations (i.e. false messages regarding poor quality products), preserving the overall ability of customers to distinguish among good bad products. Firms will invest in promotions as long as the number of honest customers in the forum is low; in this case promotions will dominate consumers' true recommendations and affect their choices. If the number of

customers' recommendations was very large, firms would no longer be effective in driving consumption choices whatever the effort put in sending messages.

Will the 'market for feedback' make the feedback system valueless?

Cabral and Hortacsu (2006) note that the existing literature does not have a clear prediction on who is more likely to buy good reputation. As discussed above, Tadelis (1999) shows that any reputation can be bought by any type of firms. Similar results are also obtained by Samuelson and Mailath (2001), although they identify one level of average-reputation that induces only one type of firm (good) to buy it. However, given these results together with those of Mayzlin (2006) and Dellarocas (2006), we conjecture that the eBay's feedback system is likely to maintain substantial value as long as buying a positive feedback is costly (price + transaction costs do not shrink to zero).

6. Concluding remarks

Feedback mechanisms are used by many online marketplaces in order to stimulate traders' good performance and to increase transaction efficiency. As these mechanisms tend to be of common use, the incentives for e-traders to manipulate them become stronger, thus undermining the reliability of reputation information they produce. Feedback manipulations rise today one important concern for platforms using reputation mechanisms. In this paper we highlight a puzzling asymmetry in feedback behaviour on eBay, and identify a form of manipulation – giving rise to a real 'market for feedbacks' - that could be at the origin of such asymmetry. We reported results from field experiments showing that positive feedback could be bought for 0.12 cents, including transaction costs, then we discuss which could be

the consequences of such feedback trade for the role of eBay's feedback system in 'normal' eBay markets. This paper suggests that, apart from excluding who trades feedback from the platform, an increasing fixed fee on each transaction seems to be the only possible avenue eBay has to reduce the negative impact of the market for feedback on confidence in the feedback system.

References

Avery, C., P. Resnick and R. Zeckhauser, 1999. "The Market for Evaluations", *American Economic Review*, 89, pp. 564-584.

Baliga, S., 1999. "Monitoring and Collusion with Soft Information", *Journal of Law, Economics & Organization*, 15, pp. 434-440.

Bar-Isaac, H., 2006. "Something to prove: Reputation in Teams", forthcoming in the *Rand Journal of Economics*.

Bar-Isaac, H. and Tadelis, 2007. "Seller Reputation", manuscript, Berkeley and New York University.

Cabral L. and A. Hortacsu, 2006. "The Dynamics of Seller Reputation", manuscript, New York University

Cheng, A. and E. Friedman, 2005. "Sybilproof Reputation Mechanisms", in *Proceedings of the 3rd Workshop on the Economics of Peer-2-Peer Systems (P2PECON)*, Philadelphia, PA.

Chwelos P. and T. Dhar, 2005. "Caveat Emptor: Differences in Online Reputation Mechanisms", manuscript, University of British Columbia.

Dellarocas C., 2006. "Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms", *Management Science* Vol. 52, No. 10, 1577–1593.

Dellarocas, C. 2005. “Reputation Mechanism Design in Online Trading Environments with Pure Moral Hazard”, *Information Systems Research*, 16 (2), June.

Dellarocas, C. 2003. “The Digitization of Word-of-Mouth: Promise and Challenges of Online Feedback Mechanisms”, *Management Science*, 49 (10), 1407- 1424.

Dellarocas C., 2000. “Immunizing online reputation reporting systems against unfair ratings and discriminatory behaviour”, in *Proceedings of the 2nd ACM Conference on Electronic Commerce*. Association for Computing Machinery, Minneapolis, MN, 150-157.

Dellarocas C., F. Dini and G. Spagnolo, 2006. “Designing Reputation Mechanisms”, in N. Dimitri, G. Piga and G. Spagnolo (Eds.), *Handbook of Procurement*, Cambridge University Press, Cambridge, N.Y..

Dellarocas, C. and C.A. Wood. 2006. “The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias”, forthcoming in *Management Science*.

Dellarocas, C., M. Fan and C.A. Wood. 2004. “Self-Interest, Reciprocity, and Participation in Online Reputation Systems”, MIT Sloan School of Management, Working Paper No. 205.

Dini, F. and G. Spagnolo. 2004. “Reputation Mechanisms and Electronic Markets: Economic Issues and Proposals for Public Procurement”, in *Challenges in Public Procurement: an International Perspective*, ed. by K. V. Thai, A. Araujo, R. Y. Carter, G. Callender, D. Drabkin, R. Grimm, K. R. Ejlskov Jensen, R. E. Lloyd, C. P. McCue, and J. Telgen. Academic Press.

Friedman, E. and P. Resnick. 2001. The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy*. 10 (1), 173-199.

Harmon, A. 2004. "Amazon glitch unmask war of reviewers". *New York Times* (February 14).

Jurca, R. and B. Faltings. 2004. "Eliciting Honest Feedback without Independent Verification Authorities", in *Proceedings of the Sixth International Workshop on Agent Mediated Electronic Commerce, (AMEC VI 2004)*, New York, USA, July 19 2004.

Kandori, M. 1992. "Social Norms and Community Enforcement", *Review of Economic Studies*, 80, 59-63.

Klein, T., Lambertz, C., Spagnolo, G. and Stahl, K. 2006. „Last Minute Feedback“, *CEPR Discussion Paper No 5693*.

Klein, T., Lambertz, C., Spagnolo, G. and Stahl, K. 2007. „The Actual Structure of eBay's Feedback Mechanism and Early Evidence on the Effects of Recent Changes“, *SFB/TR 15 Discussion Paper No. 220*, Universities of Mannheim, Bonn, Berlin and Munich, available from <http://ideas.repec.org/p/trf/wpaper/220.html>.

Klein, B. and K. B. Leffler. 1981. "The Role of Market Forces in Assuring Contractual Performance", *Journal of Political Economy*, 89, 615-641.

Lippert, S. and G. Spagnolo, 2006. „Networks of Relations, Word-of-Mouth Communication, and Social Capital“, manuscript, Massey University and University of Rome 'Tor Vergata'.

Mailath G.J and L. Samuelson. 2001. "Who Wants a Good Reputation?" *Review of Economic Studies*, 68, 415-441.

Mailath G.J and L. Samuelson. 2006. *Repeated Games and Reputations: Long-run Relationships*, Oxford University Press."

Mayzlin, D. 2006. "Promotional Chat on the Internet". *Marketing Science*, 25(2) 157–165.

Milgrom, P., D. North, and B. Weingast. 1990. "The role of institutions in the revival of trade: the law merchant, private judges, and Champagne fairs". *Economics and Politics* 2, 1023.

Miller, N., Resnick P. and R. Zeckhauser. 2005. "Eliciting Honest Feedback in Electronic Markets". Forthcoming in *Management Science*.

Papaioannou, T. and G. D., Stamoulis, 2005. "An Incentives Mechanism Promoting Truthful Feedback in Peer-to-Peer Systems". In *Proceedings of IEEE/ACM CCGRID 2005 (Workshop on Global P2P Computing)*.

Peyrache E. and L. Quesada. 2005. "Credibility of Auditors and Incentives to Collude", Working paper, Groupe HEC and Universidad Torcuato Di Tella.

Resnick, P. and R. Zeckhauser. 2002. "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System", in Michael R. Baye, ed. *The Economics of the Internet and E-Commerce* (Advances in Applied Microeconomics, Vol. 11). JAI Press.

Tadelis S., 2003. "Firm Reputations with Hidden information," *Economic Theory*, 21(2-

Tadelis S., 2002. “The Market for Reputations as an Incentive Mechanism”, *Journal of Political Economy*, 92(2), 854-882.

Tadelis S., 1999. “What’s in a Name? Reputation as a Tradeable Asset”, *The American Economic Review*, Vol. 89(3), 548-563.
3), 635-651.