Tilburg University

Adverse Selection and Moral Hazard in Anonymous Markets

Klein, T.J.; Lambertz, C.; Stahl, K.

Publication date:
2013

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Klein, T. J., Lambertz, C., & Stahl, K. (2013). *Adverse Selection and Moral Hazard in Anonymous Markets*. (CentER Discussion Paper; Vol. 2013-032). Econometrics.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

No. 2013-032

**ADVERSE SELECTION AND MORAL HAZARD
IN ANONYMOUS MARKETS**

By

Tobias J. Klein, Christian Lambertz,
Konrad O. Stahl

June, 2013

ISSN 0924-7815
ISSN 2213-9532

Adverse Selection and Moral Hazard in Anonymous Markets*

Tobias J. Klein

Christian Lambertz

Konrad O. Stahl[†]

June 2013

Abstract

We study the effects of improvements in eBay’s rating mechanism on seller exit and continuing sellers’ behavior. Following a large sample of sellers over time, we exploit the fact that the rating mechanism was changed to reduce strategic bias in buyer rating. That improvement did not lead to increased exit of poorly rated sellers. Yet, buyer valuation of the staying sellers—especially the poorly rated ones—improved significantly. By our preferred interpretation, the latter effect results from increased seller effort; also, when sellers have the choice between exiting (a reduction in adverse selection) and improved behavior (a reduction in moral hazard), then they prefer the latter because of lower cost.

JEL classification: D83, L15.

Keywords: Anonymous markets, adverse selection, moral hazard, reputation building mechanisms.

*We would like to thank the participants of the 2013 ICT Workshop in Evora and of the 2nd MaCCI day in Mannheim, and in particular Benjamin Engelstätter, Volker Nocke, Martin Peitz, Yossi Spiegel and Stefan Weiergräber for helpful comments. Financial support from the Deutsche Forschungsgemeinschaft through SFB TR-15 and Microsoft through a grant to TILEC is gratefully acknowledged. The latter was provided in accordance with the [KNAW Declaration of Scientific Independence](#). The views expressed here are not necessarily the ones of the Microsoft Corporation.

[†]T. J. Klein: CentER, TILEC, Tilburg University; C. Lambertz: University of Mannheim; K. O. Stahl: University of Mannheim, C.E.P.R., CESifo, and ZEW. E-Mail for correspondence: kos@econ.uni-mannheim.de.

1 Introduction

Informational asymmetries abound in anonymous markets such as those opening in the internet almost every day. Before trading takes place, the typical buyer often cannot be sure of the quality of the good to be exchanged. Even if she can be sure, for instance when purchasing a certified commodity, she does not know the seller's effort spent on its delivery. In turn, when the good is traded anonymously, the typical seller does most often not know the buyer's payment morale.

Without remedies, these informational asymmetries invite adverse selection and moral hazard. *Adverse selection* may arise along Akerlof's (1970) classical argument. Sellers offering costly high quality products leave—or may not even enter—the market, as long as quality is unobservable to the buyers, and thus their willingness to pay does not exceed average quality. For complementary reasons, sellers of low quality products or low quality service tend to self-select into such a market, because they can cheat on buyers by claiming to offer high quality products. *Moral hazard* may arise because effort on both sides of the market is endogenous: sellers may package goods badly, or delay, or default on delivery. Likewise, buyers may delay, or default on payments.

Whereas the consequences of adverse selection and moral hazard are well understood conceptually, we still lack empirical tests on the direction of the effects as predicted by theory, and evidence on the magnitude of these effects. Anonymous markets such as those organized in the internet provide a useful environment for collecting evidence and conducting tests. Faced with adverse selection and moral hazard in these anonymous markets, the market organizers designed remedies early on. In particular they constructed mechanisms under which buyers and sellers mutually evaluate their performance; and documented them, so that agents on both sides of the market could build reputation capital.

These reporting mechanisms were adjusted over time in reaction to opportunistic reporting behavior on one, or both sides of the market, that was unexpected by their designers. These changes in the reporting mechanism are typically also unexpected by market participants, and thus can be perceived as natural experiments. We collected reporting data before, and after such changes to address a question we feel to be central in evaluating the performance of anonymous markets: *What are the qualitative and quantitative effects of an unexpected improvement*

in the typical market participant's reporting possibilities about her counterpart's performance?

It would be optimal to study the effects on the basis of the agents' factual behavior in the individual transactions. Yet that behavior is not documented by the intermediaries, but rather its evaluation by the counterparty, which needs interpretation. At the same time, it provides more information than is typically available in traditional offline markets, where data on seller behavior other than assortment and prices are typically not available.

eBay's classic reputation mechanism allows buyers and sellers to mutually evaluate their performance in just completed transactions. In May 2007, eBay added a new, second rating system, called Detailed Seller Ratings (DSR), that allows buyers to rate seller performance in detail—but not *vice versa*. The DSRs are reported as moving averages over the last 12 months, so unlike under the classic feedback scheme, the buyer's individual evaluation can no longer be identified by the seller under the DSR scheme. One year later, in May 2008, eBay also changed the symmetry between buyer and seller rating in its classic feedback scheme, by forbidding negative seller rating and with it, removing buyer fear of seller retaliation to a bad rating by the buyer, that was likely to have had an influence on buyer ratings before the change was enacted.

Lacking documentation about the agents' behavior in the transaction, which is not recorded by eBay, we exploit these changes in eBay's feedback mechanism as follows. Since under the DSR scheme, the buyer's evaluation of a particular transaction is no longer identifiable by the seller, the DSR can now be considered an unbiased reflection of buyers' satisfaction of that seller. We use this unbiased measure to identify how buyer satisfaction was affected by the change in eBay's classic feedback mechanism one year later.

We show that this second change leads to a significant increase in buyer satisfaction with the incumbent sellers. The effect is substantively and significantly stronger for sellers that had performed poorly before the change. In fact, rather than exiting from the market—which would have been the alternative reaction to the change in the reporting mechanism—, sellers improve on all DSR dimensions.

Towards our preferred interpretation of these results, we develop a toy model from which we predict effects on seller adverse selection and moral hazard. Within this toy model, removing buyer fear of adverse retaliation by the seller incentivizes the buyer to report truthfully rather than opportunistically, and in particular to report bad experiences. That, in turn, ameliorates *adverse selection*, by leading bad seller types to exit the market; or ameliorates *moral hazard*,

by the remaining sellers engaging in more effort towards improving on buyer satisfaction.

We assess the robustness of our empirical results against alternative interpretations and conclude that our interpretation fits best, that is: *the reduction of the informational asymmetry results in a reduction of seller moral hazard, rather than a reduction of adverse selection.*

Towards detailing procedure and results, we proceed as follows. In the next Section 2, we report on the literature pertinent to what we do. In Section 3, we describe the eBay Feedback Mechanism, and in particular the two changes we focus on. Section 4 contains the description of our data. In Section 5, we present our central results. In the ensuing Section 6 we develop our toy model, from which we derive our preferred explanation and interpretation of the results. In Section 7 we assess the robustness of our interpretation and provide additional supportive results. In particular, we show that the effect of eBay’s introduction of a new way of sorting listings had, if at all, then only a small short run effect on seller behavior. We conclude with Section 8.

2 Literature

Adverse selection and moral hazard are related to asymmetric information between two contracting parties.¹ Take a car insurance company and an individual insurance taker. Adverse selection arises if the individual self selects into buying high coverage because she knows to be an unsafe driver. Moral hazard is present if she pursues lesser accident-preventing effort in the face of the better coverage, and is therefore more likely to have an accident.

It is difficult to disentangle adverse selection and moral hazard by just observing the incidence of accidents conditional on coverage. Therefore, authors have first focused on showing that asymmetric information affects economic outcomes. In general, this is done by relating individual choices to *ex post* outcomes (see Chiappori, 2000, for an early review). For example, Chiappori and Salanié (2000) use data on contracts and accidents in the French market for automobile insurance to test whether insurance contracts with more comprehensive coverage are chosen by individuals who then have higher claim probabilities. If this is the case, then this can either be explained by moral hazard, or adverse selection, or both, without further discrimination.²

¹Sometimes, a distinction is made between *ex ante* moral hazard and *ex post* moral hazard. The terminology adverse selection refers to the former in this paper, and the terminology moral hazard to the latter.

²See also Finkelstein and Poterba (2004) for a similar approach in the context of annuitization and mortality and Fang, Keane, and Silverman (2006) in the context of health insurance.

From the theorist’s point of view, the inability to disentangle adverse selection and moral hazard effects does not come as surprise: the analyst typically cannot observe self selection *ex ante* by type because the type is largely private information. In addition, with an endogenous change of effort, that type can modify the outcome. In the interpretation of our results preferred by us, we can strictly separate the two effects *ex post*, by observing poorly performing sellers’ exit from the market as reflecting a reduction in adverse selection; and ongoing sellers’ effort taking, as reflecting a reduction in moral hazard.

At any rate, following up on [Chiappori and Salanié \(2000\)](#), [Abbring, Chiappori, and Pinquet \(2003\)](#) and [Abbring, Heckman, Chiappori, and Pinquet \(2003\)](#) argue that dynamic insurance data allow researchers to isolate moral hazard effects, by looking at insurance contracts in which the financial loss associated with a second claim in a year is bigger, so that exercising moral hazard becomes more costly, and therefore the incentive to do so decreases. One can isolate moral hazard effects in this context because one naturally follows an individual over time, and therefore the factors influencing adverse selection stay the same, while incentives to exert moral hazard change. In the context of deductibles in health insurance [Aron-Dine, Einav, Finkelstein, and Cullen \(2012\)](#) follow-up on this idea and investigate whether individuals exhibit forward looking behavior, and reject the hypothesis of myopic behavior.

Focusing on adverse selection, [Einav, Finkelstein, Ryan, Schrimpf, and Cullen \(2011\)](#) show that some individuals select insurance coverage in part based on their anticipated behavioral response to the insurance contract, and term it “selection on moral hazard”. For this, they exploit variation in the health insurance options, choices and subsequent medical utilization across different groups of workers at different points in time. Also [Bajari, Hong, and Khwaja \(2006\)](#) study individual selection of insurance contracts. They provide, as we do in our very different context, evidence of moral hazard, but not of adverse selection. Their result is based on a structural model of demand for health insurance, in which, in order to isolate selectivity *ex ante* and lacking exogenous variation, they need to control in an elaborate way for individual risk and risk preference.

We instead develop our results from a natural experiment, involving, in our interpretation, self selection and adjustment of moral hazard *ex post*. We follow sellers over time, which allows us to control for unobserved differences across sellers by means of fixed effects when studying moral hazard. We then study whether an improvement of the mechanism led to increased exit

from the market on the one hand, and a reduction of moral hazard on the other hand.

This paper is more generally related to [Avery, Resnick, and Zeckhauser's \(1999\)](#) somewhat sweeping general hypothesis that the internet has greatly reduced the cost of distributing information and that there is an efficient provision of evaluations by users. See also [Dranove and Jin \(2010\)](#), [Bajari and Hortascu \(2004\)](#) and [Cabral \(2012\)](#) for reviews of the theoretical and empirical literature on quality disclosure on the internet.

For eBay, the general finding is that better ratings benefit sellers by an increase in the probability to sell a product, and in its selling price. See, e.g., [Melnik and Alm \(2002\)](#), [Lucking-Reiley and Reeves \(2007\)](#) and [Jin and Kato \(2008\)](#) for evidence using field data, and [Resnick, Zeckhauser, Swanson, and Lockwood \(2006\)](#) for experimental evidence.³

These results show that ratings on eBay convey information, but it is unclear how much. The reason is that, due to the design of the reputation mechanism, ratings were biased before the implementation of DSR, and the removal of symmetric classic feedback. [Resnick and Zeckhauser \(2002\)](#) provide reduced-form evidence that points towards underreporting of negative experiences and [Klein, Lambertz, Spagnolo, and Stahl \(2006\)](#) complement this by showing that the probability to leave a negative rating increases substantially towards the end of the period in which feedback can be left.

[Klein, Lambertz, Spagnolo, and Stahl \(2009\)](#) provide detailed information on the actual structure of the feedback mechanism and provide first descriptive evidence on the newly introduced DSRs. [Bolton, Greiner, and Ockenfels \(forthcoming\)](#) also provide such evidence and complement it with an experimental study. Focusing on why classic ratings are left at all, [Dellarocas and Wood \(2008\)](#) estimate a model of rating behavior, assuming that ratings, once given, are truthful, and estimate the true underlying distribution of satisfaction. This can be seen as controlling for the selection bias that comes from traders being much more likely to leave a rating when satisfied.

[Cabral and Hortascu \(2010\)](#) provide indirect evidence for the presence of moral hazard on eBay. They find that, when a seller first receives negative feedback, his sales rate drops and he is more likely to receive negative feedback and to exit. Moreover, they find that just before exiting, sellers receive more negative feedback than their lifetime average. With our paper we

³In a different context, [Anderson and Magruder \(2012\)](#) relate online ratings of restaurants to restaurant reservation availability and find that an extra half-star on the popular platform [Yelp.com](#) causes restaurants to sell out 19 percentage points more frequently.

complement the aforementioned studies by providing direct evidence on one of the most policy-relevant questions, namely the relationship between the design of the feedback mechanism and the presence of moral hazard or adverse selection.

3 eBay’s Feedback Mechanism

eBay’s feedback mechanism by which sellers and buyers could evaluate the performance of their trading partners was introduced in February 1996, just a few months after the first auction had taken place on its website.⁴ In its earliest form, the system allowed any eBay user to leave feedback on the performance of any other user, i.e., a “positive,” “neutral,” or “negative” rating accompanied by a textual comment. This feedback was immediately observable on his or her “Feedback Profile” page, together with all ratings and comments that a user had ever received by other users.

In February 2000, four years after its institution, the mechanism was changed to transaction-specific feedback. Since then, all new ratings must relate to a particular transaction, i.e. only the seller and the buyer in a particular transaction can rate each other regarding their performance in that transaction.

From early on, the feedback mechanism has led to conflicts and heated discussions about unfairly biased reporting. As a consequence, eBay repeatedly modified the system. The two major changes we focus on here were made in May 2007 and May 2008, respectively. In May 2007, eBay introduced a new form of unilateral rating by buyers: Detailed Seller Ratings (DSR). In addition to the original bilateral rating available heretofore, buyers could now separately rate the transaction items *accuracy of the item description*, *communication*, *shipping speed*, and *shipping charges* with one to five stars. These detailed ratings are left anonymously and only unilaterally by the buyer; they are anonymized by being published in summarized form only, so that the individual rating cannot be identified by the seller.

This change tackles what was felt to be a substantial flaw in eBay’s original bilateral feedback mechanism, namely the buyer’s fear of retaliation when leaving a negative rating *before* the seller—a problem well known to many eBay users and well discussed among scholars for some

⁴An early description of the basic mechanism and an analysis of rating behavior are given in [Resnick and Zeckhauser \(2002\)](#).

Figure 1: Changes to the feedback mechanism

time.⁵ An important detail is that DSRs can only be left when a classic rating is left. The two ratings need not be consistent, however. That is, for the very same transaction, a buyer could leave a positive classic rating identifiable by the seller—and a negative, truthful set of DSRs not identifiable by him.

In May 2008, the classic bilateral feedback mechanism was transformed to effectively a unilateral one as well: sellers could only leave positive ratings on buyers—or none at all. Thereby, eBay removed the possibility that the seller would strategically postpone his rating, in order to implicitly or explicitly threaten the buyer with retaliation to a negative rating.⁶ The two changes are summarized in Figure 1.⁷

Since buyers can leave DSR without threat of retaliatory feedback by the seller, we take it that a buyer’s DSR constitutes a possibly subjective, but strategically unbiased evaluation of seller performance.⁸ We use this as the basis for investigating how individual seller performance reacts to the May 2008 change, when all ratings were effectively made unilateral. As the May 2008 change enacted by eBay addresses the classic feedback system that was actually better established amongst users than DSR, we do expect a significant impact of that second change, even though the May 2007 change to DSR already had made room for candid ratings that could

⁵Klein, Lambertz, Spagnolo, and Stahl (2006) gave an early account of this.

⁶In fact, eBay stated the reasons for this step in a public announcement in January 2008: *Today, the biggest issue with the system is that buyers are more afraid than ever to leave honest, accurate feedback because of the threat of retaliation. In fact, when buyers have a bad experience on eBay, the final straw for many of them is getting a negative feedback, especially of a retaliatory nature...Now, we realize that feedback has been a two-way street, but our data shows a disturbing trend, which is that sellers leave retaliatory feedback eight times more frequently than buyers do... and this figure is up dramatically from only a few years ago. So we have to put a stop to this and put trust back into the system...here’s the biggest change, starting in May: Sellers may only leave positive feedback for buyers (at the seller’s option).* (Taken from <http://announcements.ebay.com/2008/01/a-message-from-bill-cobb-new-pricing-and-other-news/>, last accessed in June 2013.)

⁷Additional changes aiming at alleviating sellers’ concerns about strategic feedback extortion by the buyers but not of interest in this study were at several points in time. For instance, buyers could threaten to leave a negative rating if not given a discount, without any fear of retaliation by the seller. To remedy this, eBay abandoned any options to mutually withdraw feedback. Another change was the introduction of a new way of sorting auction listings. We discuss the latter change in more detail in the form of a robustness check in Section 7.3.

⁸As eBay does not follow up the transactions process, there was no unbiased information about eBay transactions available at all before the DSR mechanism was implemented.

have affected seller behavior.

We concentrate on two effects expected with this change: first, that poorly performing sellers would leave the market place, resulting in a *reduction in adverse selection*; and second, that the continuing sellers would be disciplined and wage more effort in satisfying buyers, resulting in a *reduction in moral hazard*.

4 Data

Our data are monthly information on feedback received for about 15,000 eBay users over a period of three years, between July 2006 and July 2009. The data were collected from eBay’s U.S. website using automated download routines and scripts to parse the retrieved web pages for the details in focus. In May 2007, we drew a random sample of, respectively, 3,000 users who offered an item in one of five different categories. The categories were (1) Laptops & Notebooks, (2) Apple iPods & Other MP3 Players, (3) Model Railroads & Trains, (4) Trading Cards, and (5) Food & Wine.⁹ We chose these categories because they were popular enough to provide us with a large list of active sellers, and because they appeared reasonably heterogeneous to us not only across categories, but also within categories as none of them was dominated by the listings of a few sellers. From June 1st, 2007 onwards we downloaded these users’ “Feedback Profile” pages on 18 occasions, always on the first day of the month. The last data collection took place on July 1st, 2009. The information dating back from May 2007 to July 2006 was inferred from the data drawn in June, 2007.¹⁰

Towards capturing changes in adverse selection, we needed to specify *seller exits* in the observation period. We identified as the date of exit the date after which a user did not receive any new DSRs during our observation window. This is a proxy, as it may also apply to users not completing any transaction for a period of time beyond our observation window, but being active thereafter. However, when we follow individuals who become inactive in the very beginning of our sampling period, then we see that it is very unlikely that such users will become active again

⁹See Table 5 in Appendix A for the exact categories.

¹⁰See Figure 8 in Appendix A for a graphical representation of the times at which we collected data. This data collection design implied that we were unable to recover seller entry in the relevant time interval.

within a couple of months.^{11,12}

Out of the 15,000 user names we drew in May 2007, we were able to download feedback profiles for 14,937 unique users in our first data collection on June 1st, 2007.¹³ One year later, we could still download data for 14,684 users, and two years later for 14,360 users.¹⁴

Table 1 gives summary statistics. As described above, the first data collection took place on June 1, 2007. On that day, the average user in our sample was active on eBay for almost four years. Proxying user experience by the length of time a user has registered, the most experienced user in our sample had registered with eBay more than eleven years before we collected our first data, and the least experienced user just a few days before our observation window opened. About 2,000 of our users had registered their accounts before the turn of the millennium, and about 3,000 users only within two years before the May 2008 changes.

On eBay, the feedback score is given by the number of distinct users who have left more positive classic ratings than negative ones, minus the number of users who have left more negative ratings than positive ones. At the time the observation window opened, the mean feedback score of our users was 564, the median score was 88, and 769 users had a feedback score of zero. The average share of positive feedback users have ever received was 99.09 per cent, which corresponds well to findings in other studies. The median number of feedbacks received during the year before that was 43. In the following year, users received roughly as many classic ratings as in the year before, and also the percentage positive ratings was very similar. On June 1, 2008, statistics for the DSRs are available for the 4,429 users who received more than 10 DSRs. The reason is that otherwise, anonymity of the reporting agent would not be guaranteed, as a seller could infer the

¹¹We are more likely to mis-classify sellers as inactive towards the very end of our observational period. Yet this is unlikely to affect our results, because we have collected data for more than a year after the feedback change, and we mostly use information around the May 2008 change to the system. See also footnote 21.

¹²We also re-did our analysis using a different definition of exit, namely the time after which a user did not receive any classic ratings anymore. This criterion captures the activity of users when active as a buyer or a seller, as classic ratings can be received when acting in either role. If users are equally likely to stop being active as a buyer before and after the change to the classic feedback mechanism then finding an increase in the probability that became inactive according to this more strict criterion also indicates that adverse selection was affected by the feedback change. The empirical patterns are very similar for both definitions of exit and therefore we do not report the results below.

¹³There were download errors for 11 users and we decided to drop three users from our panel for which eBay apparently reported wrong statistics. Moreover, there are 48 users in our sample who had listings in two categories (and therefore are not unique), and two users who had listings in three of our five categories. We dropped the duplicate observations.

¹⁴We devoted a large amount of effort to following users when they changed their user names. Still, we were not successful in doing so for all users. One reason for this is that after an account is closed, the associated feedback profile is no longer available, and therefore, we could not find it. Over time, it became more difficult to download data, as eBay put measures in place that prevented us from downloading enough pages within a short amount of time.

Table 1: Summary statistics

			percentile				
	obs.	mean	5	25	50	75	95
June 1, 2007							
duration membership in years	14,937	3.83	0.09	1.33	3.54	6.12	8.46
feedback score	14,937	563.66	0.00	18.00	88.00	339.00	2099.00
percentage positive classic ratings	14,189 ^a	99.09	97.10	99.70	100.00	100.00	100.00
member is PowerSeller	14,937	0.07	-	-	-	-	-
number classic ratings previous 12 months	14,937	273.10	0.00	10.00	43.00	161.00	975.00
percentage positive classic ratings previous 12 months	13,943 ^b	98.95	96.49	100.00	100.00	100.00	100.00
June 1, 2008							
number classic ratings previous 12 months	14,684	282.16	0.00	10.00	45.00	164.00	1042.00
percentage positive classic ratings previous 12 months	13,812 ^c	97.95	93.10	99.53	100.00	100.00	100.00
number DSR previous 12 months	4,429 ^d	378.78	12.00	28.00	78.25	265.50	1378.25
DSR score	4,429	4.71	4.35	4.65	4.75	4.82	4.90
number DRS relative to number classic feedbacks	4,429	0.42	0.10	0.27	0.44	0.59	0.70
June 1, 2009							
number classic ratings previous 12 months	14,360	200.45	0.00	2.00	20.00	97.00	761.00
percentage positive classic ratings previous 12 months	11,524 ^e	99.47	98.18	100.00	100.00	100.00	100.00
number DSR previous 12 months	3,272 ^f	376.41	12.00	26.38	72.00	255.75	1378.00
DSR score	3,272	4.78	4.53	4.73	4.82	4.88	4.95
number DRS relative to number classic feedbacks	3,272	0.46	0.11	0.29	0.48	0.63	0.74

Notes: The table shows summary statistics for our sample of sellers. There are three panels, each one for a different point in time for which we report summary statistics. These points in time are the first point in time at which we collected data, as well as one and two years after that. DSRs were introduced in May 2007, so the first point in time is the beginning of the first month after this. The change in the classic feedback mechanism whose effect we analyze occurred in May 2008, i.e. in the month prior to the second point in time for which we report summary statistics. The third point in time is one year after that. The feedback score is the number of users who have mostly left positive feedback in the classic system, minus the number of users who have mostly left negative feedback. The PowerSeller status is awarded by eBay if a seller has a particularly high transaction volume and generally a good track record. DSR is the average DSR, per user, across the four rating dimensions. ^aCalculated for those 14,189 users whose feedback score is positive. ^bCalculated for those 13,943 users who received classic feedbacks in the previous 12 months. ^cCalculated for those 13,812 users who received classic feedbacks in the previous 12 months. This is also inferred from other points in time, hence the number of observations is higher than it is for other statistics in the same panel. ^dCalculated for those 4,429 users who received enough DSRs so that the score was displayed. The statistics in the following two rows are calculated for the same users. ^eCalculated for those 11,524 users who received classic feedbacks in the previous 12 months. ^fCalculated for those 3,272 users who received enough DSRs so that the score was displayed. The statistics in the following two rows are calculated for the same users.

rating from the change in the DSR. DSR scores are available for about 15% of the users one month after their introduction in May 2007, and for about 30% of users one year later. The DSR score we report on here is the average reported score across the four rating dimensions. Yet another year later, the picture looks again similar, except for the number of classic ratings received, which has decreased.

At this point, it is useful to recall the objective of our analysis: it is to study sellers’ reactions to the May 2008 system change, on the basis of unbiased ratings by their buyers effective with the introduction of DSR one year before. Users may sometimes act as sellers, and sometimes as buyers. With our sampling rule, we ensure, however, that they were sellers in one of the five specified categories in May 2007. Moreover, DSRs can only be received by users when they act as sellers. Hence, the average DSR score will reflect only how a user behaved in that very role.¹⁵ Still, it is also important to keep in mind that we will not be able to observe the reaction of sellers who receive less than 10 DSRs per year.¹⁶ However, looked at it in a different way, we capture behavior that is associated with most of the transactions on eBay, as those sellers who receive less than 10 DSRs per year are not involved in most of the sales on eBay.

5 Results

5.1 Incumbent sellers’ reactions

The May 2008 change provided additional means to buyers to voice negative experiences without fear of negative seller reaction. This should have incentivized continuing sellers to prevent negative buyer ratings by significantly reducing shirking, i.e. not describing and pricing goods as of higher than the true quality; and increasing their effort in the other dimensions towards satisfying buyers. Therefore, we expect a significant increase in buyer satisfaction, as measured

¹⁵One may still wonder how often the users in our sample acted as buyers. On June 20th, 2008, eBay reveals in a statement that buyers leave DSR 76% of the time when leaving “classic” feedback. In our data collection just before this statement, the mean overall “DSR to classic” ratio of users for whom a DSR is displayed is about 43%. The difference between those 76%, where users acted as sellers, and the 43%, where they acted as buyers or sellers comes about because they may also have acted as buyers. Looked at it in a different way, the 43% in our sample is a lower bound on the probability that a user has acted as a seller in a given transaction, because DSRs can only be left when a classic rating is left at the same time. It is a rather conservative lower bound because it assumes that a user receives a DSR every time he acts as a seller and it takes a bit more effort for the buyer to leave a DSR, as compared to leaving only a classic rating. In any case, DSR scores reflect how users have behaved when acting in their role as sellers.

¹⁶Throughout, we will control for seller fixed effects. This is important because sellers for whom DSRs are available may be different from those for whom DSRs are not available; and because sellers who exit at some point may be different from those who will not exit. See also the discussion below.

Figure 2: Evolution of Detailed Seller Ratings

Notes: The figure shows how DSRs changed over time. The vertical line denotes the time of the May 2008 change to the classic feedback mechanism. The dots are averages across users for whom DSRs are displayed and the error bars depict corresponding 95% confidence intervals. The circles are linearly interpolated values for the periods in which we did not collect data. Before averaging DSRs across users we calculate the average DSR per user, across the four categories. The horizontal dashed lines visualize that the dots are averages over the 12 months prior to the point in time at which the DSRs are displayed.

by the DSR scores.

This expectation is verified in Figure 2 in which we summarize the evolution of DSRs in the relevant two 12-month intervals before and after the May 2008 change. Recall that at any point in time, DSR indices are published in four categories, for every seller that has received more than 10 DSRs up to that point, with ratings aggregated over the respective preceding 12 months. In the figure, each dot reflects the overall average across sellers and categories. The patterns by category resemble one another closely and are reported in detail in Figure 10 in Appendix B.

When interpreting Figure 2 it is important to once again keep in mind that DSR scores show the average of all DSR ratings given in the previous 12 months. Therefore, if all ratings received after the change were higher, one would observe unchanged ratings reflected in a flat curve before the change, then an increasing function in the 12 months after the change, and again a flat curve (at a higher level) after that. The full effect of the change equals the difference

between the DSR rating one year after the change and the DSR rating right before the change, and is approximately proportional to the slope of the line after the change.

The horizontal lines in Figure 2 depict the relevant average scores that were accumulated at the beginning of May 2008 and June 2009, respectively.¹⁷ The difference in the horizontal lines equals the difference between buyer ratings one year after the change, i.e. after the change in buyer ratings is completely reflected in the 12 months moving average, relative to the score just before the change. The figure clearly shows that the DSRs have increased after the May 2008 change.¹⁸

We performed regressions to quantify the effect controlling for fixed effects. Towards their specification, denote by DSR_{it} the average score across the four DSR rating dimensions reported for seller i in period t . Recall that our data is always drawn on the first day of the month, and that DSR_{it} is the average of all ratings seller i has received over the previous 12 months. Let w_{it}^τ be the implicit weight that is put, in the construction of the index, on $d sr_{i\tau}$, the average of all ratings given in month τ . Clearly, this weight is zero for $\tau < t - 12$ and $\tau \geq t$. Otherwise, it is given by the number of ratings received in τ divided by the total number of ratings received between period $t - 12$ and $t - 1$. Hence $\sum_{\tau=t-12}^{t-1} w_{it}^\tau = 1$ and

$$DSR_{it} = \sum_{\tau=t-12}^{t-1} w_{it}^\tau \cdot d sr_{i\tau}. \quad (1)$$

We want to estimate how $d sr_{i\tau}$ changed after May 2008. That is, we are interested in estimating the parameter β in

$$d sr_{i\tau} = \alpha + \beta POST_{i\tau} + \alpha_i + \varepsilon_{i\tau},$$

where $POST_{i\tau}$ takes on the value 1 after the change, and zero otherwise. The change occurred between the 1st of May and the 1st of June, 2008, and therefore we code $POST_{i\tau} = 1$ if τ is

¹⁷The change occurred in mid-May 2008. Hence, the DSR score at the beginning of June, 2008 contains no DSRs left before the change because it is calculated from the ratings received in the preceding 12 months. Conversely, the DSR score at the beginning of May, 2008 contains no ratings received after the change. Figure 8 in Appendix A shows at which points in time data were collected and depicts over which periods, respectively, the DSR scores were calculated.

¹⁸Unfortunately, we were not able to collect data for more than one year after the change, because eBay started to ask users to manually enter words that were hidden in pictures when more than a small number of pages were downloaded from their server. Otherwise, we would be able to assess whether the curve indeed flattens out one year after the change. The remarkable fact, however, is that the scores start increasing rapidly and immediately after the change.

equal to July 2008, or later, and $POST_{i\tau} = 0.5$ if τ is equal to June 2008. This means that we assume that half of the ratings received in May 2008 correspond to transactions taking place after the change.¹⁹ α_i is an individual fixed effect with mean zero and $\varepsilon_{i\tau}$ is an individual- and time-specific error term. We cannot estimate β directly by regressing $dsr_{i\tau}$ on $POST_{i\tau}$ because $dsr_{i\tau}$ is not observed. However, by (1), the reported DSR score is the weighted average rating received in the preceding 12 months, so that

$$DSR_{it} = \alpha + \beta \left(\sum_{\tau=t-12}^{t-1} w_i^\tau \cdot POST_{i\tau} \right) + \alpha_i + \left(\sum_{\tau=t-12}^{t-1} w_i^\tau \cdot \varepsilon_{i\tau} \right).$$

$\sum_{\tau=t-12}^{t-1} w_i^\tau \cdot POST_{i\tau}$ is the fraction of DSRs received after the 2008 change of the system. Hence, we can estimate α and β by performing a fixed effects regression of the reported DSR score on a constant term and that fraction. We can control for time trends in a similar way.²⁰

It is important to control for fixed effects in this context because the DSR score is only observable for a selected sample of sellers, namely those who were involved in enough transactions so that the the DSR score was displayed. Otherwise, results may be biased; for example because the DSR score of worse sellers with lower α_i 's would be less likely to be observed before the change because by then they would not have received enough ratings. Thereby, we also control for seller exit when studying effects on staying sellers' behavior. Controlling for fixed effects is akin to following sellers over time and seeing how the DSR score changed, knowing the fraction of the ratings that were received after the feedback change. This is generally important because we are interested in the change in the flow of DSRs that is due to the change of the May 2008 change of the feedback mechanism.

Our results will turn out to be robust to controlling for a time trend, however. In light of Figure 2 this is not surprising, as it already shows that there was no time trend in the reported DSR scores before May 2008. After that, DSR scores increase slowly over time, but we will show that this is mainly driven by the fact that DSR scores are averages over DSRs received in the

¹⁹This is conservative in the sense that, if anything, it would bias our results downwards because we would partly attribute a positive effect to the time prior to the change. Then, we would (slightly) underestimate the effect of the change. See also the discussion in Section 7.3 and the robustness check in Section 7.4.

²⁰For two separate time trends, the regressors are weighted average times before and after the change. When we subtract the time of the change from those, respectively, then the coefficient on the indicator for the time after the change is still the immediate effect of the change. The change in the trend can be seen as part of the effect. We will also make a distinction between a short-run effect and a long-run effect. For this, the regressors will be the fraction of ratings received until the end of September 2008, and thereafter.

Table 2: Effect of the May 2008 change

	(1) full sample	(2) small window	(3) time trend	(4) DSR < 4.75	(5) DSR ≥ 4.75
average DSR before change	4.7061*** (0.0007)	4.7030*** (0.0005)	4.7149*** (0.0034)	4.5912*** (0.0011)	4.8138*** (0.0006)
effect of feedback change	0.0581*** (0.0024)	0.0414*** (0.0047)		0.0904*** (0.0044)	0.0316*** (0.0021)
effect of feedback change until September 2008			0.0169** (0.0083)		
effect of feedback change after September 2008			0.0589*** (0.0184)		
linear time trend before change			0.0009** (0.0004)		
linear time trend after change			0.0007 (0.0019)		
fixed effects	yes	yes	yes	yes	yes
R^2	0.0580	0.0131	0.0603	0.0809	0.0466
number sellers	5,225	4,919	5,225	2,337	2,337
number observations	67,373	30,488	67,373	31,260	33,508

Notes: This table shows the results of regressions of the average DSR score, averaged over the four categories, on a constant term and the fraction of feedbacks received after May 2008. We assume that half of the feedbacks in May 2008 were received before the change and the other half after the change. In specification (2), we do not include observations after October 2008 and before March 2008. In specification (3) we distinguish between the effect until September 2008 and after that date, and also account for a piecewise linear time trend. Specification (4) includes only those sellers who had a DSR score below the median of 4.75 in May 2008 and (5) only those above the median. One observation is a seller-wave combination. Throughout, we control for fixed effects. The R^2 is the within- R^2 . Standard errors are cluster-robust at the seller level and significance at the 5 and 1 percent level is indicated by ** and ***, respectively.

previous year, and the fraction of DSRs received after May 2008 increased gradually over time. Consequently, the DSR scores will also only increase gradually, even if the flow of DSRs jumps up and remains unchanged at a higher level after the change.

Table 2 shows the regression results using DSR scores averaged over the four detailed scores of all sellers and using all 18 waves of data. We first look at the first three columns and will discuss the last two below. In specification (1), we use the whole sample and find an effect of 0.0581. In specification (2), we restrict the data set to the time from March 1 to October 1, 2008; hence there are only 30,488 observations. We do so to estimate the effect locally, because this allows us to see how much of this global effect is due to an immediate response by sellers. The estimated effect is equal to 0.0414, which suggests that most of the effect occurs from mid-May to October 1, 2008. In specification (3), we instead allow for a piecewise linear time trend over the entire observation window. We find that the time trend before the change is very small and

Figure 3: Evolution for two different groups

Notes: The figure shows how the average DSR score changed over time, with sellers split into those who had a DSR score above the median of 4.75 prior to the May 2008 change, and those who had a score below that. The circles and squares are linearly interpolated values for the periods in which we did not collect data. The error bars depict pointwise 95% confidence intervals.

not significantly different from zero after the change. The effect of the change is estimated to be a short-run effect of 0.0169, until the end of September 2008, and a bigger effect of 0.0589 after that.

To assess the magnitude of the effect, it is useful to express the numbers in terms of quantiles of the distribution of DSR scores among sellers prior to the May 2008 change. According to the results in the first column, the average DSR before the change is 4.7061, and after the change, it is $4.7061 + 0.0581 = 4.7642$. This corresponds to roughly the 40 and 60 percent quantiles of the distribution of ratings prior to the change, respectively. Hence, the May 2008 change has led to a *significant and sizable increase in seller effort*.

We now look at how this increase is differentiated between sellers with low, and high DSR before the change. Towards that, we split our sample at the median DSR of 4.75 just before the change between high and low ranked sellers before the May 2008 change. Figure 3 gives the picture. The increase in DSR score is stronger for sellers with below-median score *ex ante*. Furthermore, the difference in the increase between the above-median and the below-median

sellers is significantly different from zero. The last two columns of Table 2 report the corresponding estimates, again controlling for seller fixed effects. The difference between the effect for above- and below-median sellers is significantly different from zero. We obtain similar results when we perform regressions for those two different groups only for a smaller time window, as in specification (2), or control for time trends, as in specification (3). In the second part of Table 6 in Appendix B discussed later within robustness checks, we show the effects of the feedback change by decile of sellers' DSR rating. Note the decline in the magnitude and significance of the effect, with increasing decile.

Recall again that the system change was not with respect to DSR, but with respect to the classic reporting mechanism. The anonymous and unilateral DSR were established one year before the May 2008 change whose consequences we consider here, and they remained anonymous and unilateral thereafter. Already with the DSR introduced in May 2007, buyers had been able to express their true valuation of seller performance without fear of retaliation by that seller. By looking at the effect of changing the non-anonymous established reporting mechanism, we pick up only an additional effect. It is remarkable that this effect shows up as clearly as documented above.

In all, the empirical evidence provides support of our hypothesis that abandoning negative buyer rating by sellers—and thereby reducing impediments against negative buyer rating—has led to substantive seller reactions. Buyer ratings—especially of *ex ante* poorly rated sellers—improved significantly, which, as we will argue in the ensuing section, results from an improvement in the behavior especially of poorly performing sellers, and with this, a reduction in moral hazard.

5.2 Seller Exit

Recall that the May 2008 change provided additional means to buyers to voice negative experiences without fear of negative seller reaction. This should also have incentivized sellers performing poorly before the change to leave the market. To empirically assess this, we looked at the exit rate of the 10 percent worst sellers, as measured by their DSR rating on May 1, 2008. Figure 4 shows how the hazard rate into inactivity and the fraction of individuals who have

Figure 4: Exit from the market

Notes: The bars show the fraction of users that have become inactive in the previous month, among the 10 percent poorest performing sellers as measured by their DSR rating on May 1, 2008, with the percentage on the right axis. The dots denote the cumulative fraction of users who have become inactive since June 1, 2007, with the hazard rate on the left axis.

become inactive changed over time.²¹ We see that overall, many sellers leave over time, both before and after the change. By June 1, 2009, almost 80% of those sellers have become inactive. It is interesting to contrast this to the pattern for the 10 percent best sellers, again measured by their DSR rating on May 1, 2008. For them, we find that only about 45% have become inactive by that time. This shows on the one hand that eBay is a very dynamic environment, and suggests on the other hand that better performing sellers are more likely to stay in the market, as compared to poorly performing sellers.²²

The most important finding, however, is that *the May 2008 change did not trigger any significant increase in the exit rate of sellers*. We also formally tested whether the hazard rate was different before and after the change. Towards this, we excluded the month of June because the change occurred mid-Mai, and tested whether the hazard rate in July and August was equal

²¹The stronger increase in that fraction towards the end of that window is due to the truncation bias naturally reflected in our observations, by which we incorrectly sort infrequent sellers into the set of exiting firms towards the end of the window.

²²Recall that we follow a sample of sellers over time, and hence we cannot study entry.

to the hazard rate in April and May. The difference between the two is estimated to be 0.003 with a standard error of 0.002, so it is not significantly different from zero. If we focus on the 10 percent worst sellers, as measured by their DSR rating on May 1, 2008, we find that the difference is estimated to be -.0247 with a standard error of 0.0148, so also for those sellers the estimated change in the hazard rate is not significantly different from zero.²³

6 A Simple Explanatory Paradigm

Figure 5: Sequence of decisions in a typical eBay transaction

In this section, we develop the explanation of these results preferred by us, and in the ensuing section, we discuss it against alternative explanations. Our preferred explanation is developed with help of a toy model involving one stage in an infinitely repeated seller-buyer game. Rather than concentrating on auctions that have substantively lost market share in internet markets, we specify our toy model with respect to a standard transaction involving a fixed price offer by one seller, and one transaction with some interested buyer. The explanation concentrates on the effects of removing the hold-up on the typical buyer's evaluation, that before May 2008 was generated in the classic feedback system by the fact that the seller could retaliate any negative rating by the buyer.

More precisely, the sequence of decisions in such a typical eBay transaction is condensed in the time line contained in Figure 5. We focus on a rating sequence involving the seller's rating of the buyer *after* the buyer's rating of the seller, with the following justification. In [Klein, Lambertz, Spagnolo, and Stahl \(2006\)](#), we found that the seller rated his counterpart before the

²³We have also checked whether exit rates increased for other groups after the feedback change, but found no evidence for that.

buyer did so in only 37 per cent of all cases in which the rating was mutual, and that in this sequence, a positive rating by the seller was followed by a negative buyer rating in less than one per cent, indicating that the hold up situation we feel to be at the root of the phenomenon analyzed here is not prevalent in this case.

We want to retain utmost simplicity, and begin by assuming that the good to be traded by a seller S can take on one of two qualities, q_h and q_l with $q_h > q_l$, selected by nature and privately known to the seller. The seller is endowed with publicly known reputation capital consisting of two DSRs, namely *accuracy of item description* denoted by k^d , and *quality of shipping* denoted by k^s , both taking on values on $[0, 5]$. When deciding to offer the good at cost normalized to zero, he considers whether to announce it at its true quality q_i and the appropriate price p_i , $i \in \{l, h\}$, which he always does if $q = q_h$; or to shirk with some probability if $q = q_l$, by announcing the low quality good as of high quality, q_h , with an associated high price p_h .

The buyer B observes the quality-price tuple announced by the seller, which we denote by $[\hat{q}_j, \hat{p}_j]$, $j \in \{l, h\}$, as well as the seller's reputation vector $[k^d, k^s]$. On their basis she forms an expected utility $\mathbb{E}u[\hat{q}_i, \hat{p}_i, k^d, k^s]$, $i \in \{l, h\}$. Natural assumptions on this utility are that it increases in the first, the third and the fourth argument, and decreases in the second. She decides to buy the item if $\mathbb{E}u[\hat{q}_i, \hat{p}_i, k^d, k^s] \geq \tilde{u}$, where $\tilde{u}$ is her exogenously specified outside option. In the case she buys, S decides whether to engage in effort $e \in \{0, \bar{e}\}$ towards delivery, where effort is provided at unit cost. In all, the seller's stage pay off is $p_i - e \geq 0$. Finally, B receives the good, observes the accuracy of the item description and the shipping quality, and rates S in terms of these two dimensions. This results in a revision of both, k^d and k^s , that enters next period as seller S ' reputation capital.

Before May 2008, the sequence of decisions involving such a transaction was typically concluded by the additional step indicated in Figure 5, in which the seller rated the buyer along the classic scale, and this rating was publicly observed. We assume that the buyer derives value from being rated well, e.g. because she intends to use such ratings as reputation capital when entering as a seller in the future. Decisions are supposed to be taken rationally, that is, with backward induction in that simple stage game that is repeated infinitely.

Towards results from this toy model, consider first the sequence of decisions before the May 2008 change. Since the buyer derives value from being rated well, seller S can opportunistically condition his rating on the buyer's rating observed by him, by giving a negative mark if the

buyer does so. Recall that the anonymous DSR cannot be given without a non-anonymous classical rating. Retaliation by the seller implies that a buyer with strong reputational concerns is captive to the seller's rating, and thus forced to rate S positively in the classic feedback necessarily preceding the DSR, no matter the seller's decisions taken before. In this case, if nature selects q_l , the opportunistic seller can shirk with probability 1 on the buyer, by offering the low quality good at high price p_h and by not taking any effort to deliver the good—yet still receives a positive contribution to his reputation capital.

By contrast, with eBay's change in the rating mechanism effective May 2008, even the buyer with reputational concerns can give an unbiased negative rating without fearing retaliation. There is abundant evidence that a seller, who intends to stay in the market, must be very concerned with his reputation because he can sell more rapidly, and at higher prices. The May 2008 change then implies that such a seller tends to accurately describe the item even if of low quality; to then quote an appropriately low price; and to take effort in delivering the item. All of this results in good buyer marks and consequently an increase in the seller's reputation capital. Hence, with the May 2008 change, the seller intending to stay in the market is incentivized to behave properly: if nature has selected the good to be of low quality, he announces it properly and wages high effort towards its delivery, implying a stage payoff $p_l - \bar{e}$ that possibly is even negative.

Consider now the special case of a seller who is endowed with low reputation capital before the May 2008 change. He faces two alternatives: either to exit the market—but before then profitably depleting his reputation capital, by shirking, i.e. selling the low quality good at high price and by not providing costly effort towards delivery, resulting in payoff $p_h > 0$; or alternatively to forego that short run rent and to continue operating in the market—but then to provide goods in a way that his reputation capital increases very strongly, resulting in a possibly negative payoff $p_l - \bar{e}$, which later on allows him to increase his price for any delivered good without losing the buyer.

In all, on the basis of this toy model, the May 2008 change results in two main effects: first, *a reduction in moral hazard exercised by sellers who intend to stay in the market*; and second, *a reduction in adverse selection exercised by poorly rated sellers*. Moral hazard is reduced via an increasingly accurate item description and increased delivery effort of the sellers remaining in the market and results in an improved buyer evaluation; adverse selection is improved via an

increase in the exit rate of sellers poorly rated before the May 2008 change. Alternatively, if these sellers continue to stay in the market, we should see a stronger than average contribution to the reduction in moral hazard, towards an improvement in their reputation capital.²⁴

Rather than observing both effects, we only observe in our results a significant improvement in buyer ratings, which we interpret as a factual improvement of seller behavior; and do not observe any increase in the exit rate of poorly rated sellers—yet a particularly strong increase in their ratings after the May 2008 change. Why does the poorly performing sellers’ reaction to the May 2008 change appear to be so asymmetric, against exit, and for improved performance in place? We attribute this to the fact that *improving on the performance—as reflected in the buyers’ evaluations—tends to be less costly for sellers, relative to exiting from the market*. Clearly, giving up on shirking, by correctly describing and selling a low quality good at low price, involves the opportunity cost of foregoing a possibly large rent of selling that good at a high price. Yet that rent must be held against the depletion of the reputation record. By contrast, describing the item accurately involves a very small cost, as does the effort involved in careful and speedy delivery, with all of this contributing to the rent that can be obtained later from improved reputation.

7 Checking for Competing Explanations and Robustness

In Section 5, we have shown that removing negative seller ratings of buyers, and with it potential retaliation to negative buyer ratings in eBay’s classic feedback system has resulted in a significant improvement in buyer ratings, especially for sellers that previously were rated poorly; and no change in sellers’ exit behavior, especially that of the poorly rated ones. In the previous Section 6, we gave an explanation that is entirely consistent with these results. In this section we ask whether our explanation is based on unfounded assumptions, and whether our results could have been generated from other causes. We finally discuss the robustness of our assumptions on the timing of the change and delays between transaction and rating.

Returning to the first aspect, our explanation is based on the assumption that the DSRs indicate the true value of the transaction. We have no way of a direct proof of this because eBay,

²⁴As to detail, if buyers’ feedbacks are delayed, then we predict from this simple paradigm a downward jump in the feedback score right after the May 2008 change, resulting from the fact that before that change, sellers exercised moral hazard in transactions rated negatively by the buyers right after the change, whereas in transactions after the change, sellers would strategically anticipate unbiased buyer rating.

as a pure intermediary. does not record any data on the transaction. Yet we do have additional evidence supporting our assumption. In the following subsection 7.1, we show that, in line with standard predictions on end game behavior, the ratings on sellers prior to exiting the market degenerate significantly, relative to those on continuing sellers. As buyers cannot predict the exit of those sellers, we see this as a clear indication of a good correspondence between buyer rating and the transaction itself.

As to alternative causes, the results could have been generated by a change in the buyers' rating behavior (buyers could become less critical in rating sellers), or by a change in the stratum of buyers filing a report (for reasons not known to us). As our data are anonymous, we again cannot directly exclude these two causes. Yet we again can report on a strong indirect indication, on the basis of classic feedbacks. Unlike in the DSR, the seller can identify, in the classic feedback system, the feedback given by a particular buyer on a particular transaction—but cannot retaliate a negative buyer feedback after the May 2008 change. Since there is a delay between transaction and rating, transactions taking place right before the May 2008 change must have been rated right after the change. If buyers react as we expect them to react, then there should be a negative downward jump in the classic rating right after the change, thereafter (more than) compensated by improved seller behavior after the change. In subsection 7.2 we document exactly this.

In principle, the effects we report on could also have been generated from other measures taken by eBay in the observation window. In fact, three months before the May 2008 change, eBay introduced a what is called “Best Match” scheme, under which poorly rated sellers were positioned later in the rank order in which eBay presents offers on a particular item. In subsection 7.3, we show that it is unlikely that this scheme has caused the reported results.

Finally, in subsection 7.4 we check whether changes in our assumptions on the timing of the change and on the timing of ratings relative to transactions affect our results, and show that they remain unchanged.

7.1 Behavior prior to Exit

In an anonymous market such as the one under discussion, we expect rational agents to change their behavior just before leaving the market. In particular, we expect them to exert more moral hazard. In view of this, we consider it as an indication supporting our claim that buyers correctly

Figure 6: Behavior prior to inactivity

Notes: In this figure we compare the ratio between the DSR and the DSR three months before for exiting users (depicted by the squares) to that of the stayers (depicted by the dots). We used linear interpolation in case we did not collect data for the latter DSR score when calculating the ratio. Circles and squares are linearly interpolated values for the periods in which we did not collect data. The error bars depict pointwise 95% confidence intervals.

value the transaction via the DSR if we see buyer ratings on exiting sellers to degenerate.

In Figure 6, we compare the continuing and the exiting sellers' DSR scores, relative to those obtained three months earlier. Whereas the relative DSR scores of the continuing sellers remain essentially unchanged in the time window considered, the exiting sellers' DSR relative scores go down on average. The error bars seem to suggest that this difference is not significantly different from zero. As before, however, the confidence intervals are point-wise. The ratio is significantly different between those individuals exiting and those staying when we pool over the time periods. The corresponding regression with standard errors clustered at the seller level shows that the point estimate of the intercept, which is the average over the dots for the stayers in the figure is 1.001, with a standard error of 0.0001.

More importantly, the coefficient on an indicator for becoming inactive is -0.005 with a standard error of 0.0008. This means that the ratio is significantly lower for individuals who retire from the market, indicating that performance trends downwards before retirement. The

standard deviation of the ratio in a given wave, e.g. May 2008, is 0.0077, so the effect is equal to 65% of this, which arguably is non-negligible. In all, this reinforces the results derived by [Cabral and Hortascu \(2010\)](#) we have referred to in the literature section 2.

7.2 Effect on Classic Ratings

Recall that the classic feedback given to a specific transaction was (and is) identifiable by the seller (and the observer). Recall also that our analysis so far was concentrating on the subsample of the users with enough transactions as sellers so they obtained DSR records. Here we look at the effect of the May 2008 change on the classic feedback ratings received by all users sampled, buyers and sellers. Towards conducting the analysis separately for these two groups, we first classify all users sampled as being foremost sellers or buyers on eBay. We based our classification on the ratio between DSR and (cumulative) classic feedbacks received by May 1, 2008. Specifically, we classify those 25% users with the highest ratio as foremost sellers and the 25% with the lowest ratio as foremost buyers.

In Figure 7 we compare the percentage of positive feedbacks obtained for the two subpopulations in the observation window. It shows very clearly that effective May 2008, the percentage of positive feedbacks dropped for users identified as foremost sellers, but remained unchanged for those identified as foremost buyers. Our explanation is as follows: Typically there is a delay between any transaction and its evaluation by the buyer. As the May 2008 change was not anticipated by the sellers, they did, on average, exhibit more shirking and carelessness in delivery right before, rather than right after the May 2008 change. Effective this very date, however, buyers could leave negative classic ratings without the risk of retaliation by the seller. In the left figure, we therefore observe a downward jump in buyer ratings right after the May 2008 change, i.e. before sellers could react to that change. The left figure suggests also that, as ratings increase, they do react as time goes on.²⁵

Table 3 contains the corresponding formal tests. The four columns on the left contain results for foremost sellers, and the four columns on the right results for foremost buyers, following our classification. Each column corresponds to a different bandwidth for the kernel regressions, and in the rows we show results for a local linear regression and a local quadratic regression. Figure

²⁵Keep in mind, however, that some of these ratings could still have been received by those users in their role as buyers.

Figure 7: Effect on classic feedbacks

Notes: The left figure shows how the percentage positive feedbacks changed over time. The lines are fitted values of local quadratic regressions and the shaded area shows pointwise asymptotic 95 percent confidence intervals, respectively. We used the Epanechnikov kernel with a bandwidth of 200. The dots are averages per wave. The solid vertical line depicts the change to the classic feedback mechanism. The left figure is for those eBay users who were foremost sellers, and the right figure for those who were foremost buyers. To classify users, we used the ratio between DSR and classic feedbacks received on May 1, 2008. In particular, we classify those 25% users with the highest ratio as foremost sellers and the 25% with the lowest ratio as foremost buyers.

7 suggests that a bandwidth of 200 together fits the data well when we use a local quadratic specification. The corresponding estimate for sellers is $-.727$. It is significant at the 1 percent level.²⁶ Observe finally that the estimated effect for buyers is never significantly different from zero. In view of the fact that this detail is developed from the full sample (and not only from the subsample of poor performers); that it completely fits the explanation given within our toy model; and that no alternative explanation to the one given above is in sight, this result very

²⁶The pointwise confidence interval in Figure 7 would suggest that the effect of the change is not significantly different from zero at the 1% level. Any conclusion based on this would be misleading, however. The proper test is based on the estimated size of the jump and the standard error for that size: the variance of the difference between the two curves is not the sum of the two variances, but smaller, because the covariance between the two is positive.

Table 3: Effect on classic feedbacks

bandwidth	foremost sellers				foremost buyers			
	50	100	200	300	50	100	200	300
local linear	-.369 (.470)	-.542* (.296)	-.328 (.227)	-.256 (.204)	.052 (.219)	.053 (.177)	-.002 (.122)	.030 (.125)
local quadratic	-.490 (.638)	-.408 (.497)	-.727*** (.296)	-.762*** (.349)	-.085 (.378)	.098 (.230)	-.049 (.187)	-.081 (.158)

Notes: This table shows estimated effects of the feedback change on the percentage positive classic ratings received by users who were either foremost sellers or buyers. These were obtained by performing kernel regressions. We used an Epanetchnikov kernel. The cells contain estimates for local linear and local quadratic regressions and the respective standard errors in parentheses. Each column corresponds to a different bandwidth. To classify users, we used the ratio between DSR and classic feedbacks received on May 1, 2008. In particular, we classify those 25% users with the highest ratio as foremost sellers and the 25% with the lowest ratio as foremost buyers. This leads to 22,762 observations for the first group and 26262 for the second group, coming from 1168 and 1169 users, respectively. Bootstrapped standard errors are cluster-robust at the seller level. Significance at the 10 and 1 percent level is indicated by * and ***, respectively.

strongly supports the explanation we gave to our results in Section 6 as against the stories introduced at the beginning of this section.

7.3 Best Match

An alternative explanation of our results could be another change in eBay’s allocation mechanism implemented within our observation window. Three months prior to the change whose effects are discussed here, eBay changed the order in which listings were displayed when buyers on eBay searched for an item. Before that, offers were simply ranked by the time remaining until the offer was closed. Under the new ranking scheme, called “Best Match” (BM), a number of factors determined which listings appeared first. One of these factors was the DSR score, and therefore the introduction of BM provided an incentive for sellers to improve their performance.²⁷ The ranking scheme was modified several times since. The exact way of ranking listings is a trade secret highly guarded by eBay, as is e.g. Google’s search algorithm. We now assess whether the introduction of BM could have geared our results.

²⁷That change was obviously motivated by the increased attractiveness of the fixed price over the auction format to sellers: *A related reason was, we introduced the fixed price format of listings. They could be 30 days, 60 days, and 90 days. And when you have fixed price listings that can be live on the site for 30 days, 60 days, 90 days, “time ending soonest” which was a sort on eBay, no longer made sense for those types of listings. You have a 30-day listing that might only come up to the top of the results 30 days after it was listed. So we had this problem, lots of fixed price inventory, 30 days and 60 days.* (Taken from http://files.meetup.com/1537023/Best_Match_Transcript.doc, last accessed in June 2013.)

Table 4: Effect of the May 2008 change without 10 percent worst sellers

	(1) full sample	(2) small window	(3) time trend	(4) DSR < 4.75	(5) DSR ≥ 4.75
average DSR before change	4.7400*** (0.0006)	4.7426*** (0.0004)	4.7545*** (0.0030)	4.6465*** (0.0010)	4.8129*** (0.0006)
effect of feedback change	0.0535*** (0.0022)	0.0382*** (0.0038)		0.0844*** (0.0041)	0.0306*** (0.0021)
effect of feedback change until September 2008			0.0016 (0.0074)		
effect of feedback change after September 2008			0.0682*** (0.0174)		
linear time trend before change			0.0016*** (0.0004)		
linear time trend after change			-0.0013 (0.0018)		
fixed effects	yes	yes	yes	yes	yes
R^2	0.0703	0.0165	0.0748	0.1067	0.0437
number sellers	4,047	4,047	4,047	1,794	2,253
number observations	58,004	26,358	58,004	25,390	32,614

Notes: This table is the same as Table 2, except that we exclude the 10 percent worst sellers, as measured by their DSR score on March 1, 2007.

Feiring (2009, 3rd ed, p.16) reports that within the time window of our analysis, the ranking induced under the BM scheme affected only the very poorest sellers, namely those for whom *Item as Described* and *Communication, Shipping Time*, or *Shipping and Handling Charges* were ranked only 1 or 2 (out of 5) stars in more than 3 per cent, and more than 4 per cent of their transactions, respectively. We concentrate our robustness check on these. We will show first, that this is a small group of sellers, and second that excluding them from our analysis does leave our results essentially unaffected.

As the first order effect of introducing BM, we expect the sellers with relatively poor records to realize fewer transactions, and correspondingly obtain significantly fewer DSRs. So we looked at shifts in the number of DSRs received post March 2008 by percentile of sellers distributed by DSR scores. Table 6 in Appendix B shows that the number of DSRs received after the introduction of BM decreased significantly only for the 10 percent poorest sellers (the effect is -5.76 from a level of 44.44 ratings per month before that, with a standard error of 1.50). Re-doing the regressions that underly the results in Table 2 and dropping the 10 percent worst sellers yields Table 4. It is immediate that the results that are very similar, thus supporting our claim that our analysis is not affected by the introduction of BM.

7.4 Other Robustness checks

We don't know the exact date of when eBay enacted the change whose effects we are reporting here. For simplicity of our analysis, we also considered coincident the timing of a transaction and its rating.²⁸ The effects of moving the date of the change and introducing delays between transaction and rating are clearly confounding—also with delays in the sellers' and the buyers' perception of the date of change. In view of the coarseness of our data collection, we will discuss them within one sensitivity analysis.

Suppose for now that both sellers and buyers would know the exact date of the change (around May 15, 2008), and the seller changed his behavior in a transaction right after the change. That transaction was probably completed by May 25 and this was also the time at which the buyer left a DSR for him. Conversely, if a transaction took place before May 15, 2008, then the seller was not able to react to the feedback change, as it was unanticipated. Nevertheless, a feedback could have been left for that transaction in the second half of May 2008. So far, we have assumed that half of the DSRs received in May 2008 corresponded to transactions conducted after the feedback change. If the effect of the feedback change was positive, then we tend to underestimate it, because we overestimate the fraction of ratings left after the change. We do not expect this to have big effects, because the delay is likely to be small relative to the length of our observation period.

We don't have a record on delays between transaction and feedback. Yet Figure 2 in [Klein, Lambertz, Spagnolo, and Stahl \(2006\)](#) shows the distribution of the time between the end of the auction and the moment at which the first feedback was left. The vast majority of feedbacks is positive and for those about 60% are left after 2 weeks and almost 90% are left after 4 weeks.²⁹ Based on this, we re-did the analysis of Section 7.3, assuming that out of all DSRs received in March 2008, 75% of the transactions took place after the introduction of BM. Moreover, we assumed that out of all DSRs received in May 2008, 25% of the transactions took place after the change to the feedback system.

Table 7 in Appendix B shows the results. They are very similar to the ones reported in

²⁸Net of the analysis of classic feedback discussed in subsection 7.2 by which we further supported our explanation.

²⁹For negative feedbacks, the distribution is shifted to the right. [Klein, Lambertz, Spagnolo, and Stahl \(2006\)](#) argue that this may be due to strategic considerations: both parties had an incentive to wait with their first rating if it was negative, because then it was less likely to be retaliated.

Table 2 and 4.

8 Conclusion

In anonymous markets, buyers (and sellers) must rely on reports of each others' performance in order to efficiently execute transactions. In this paper, we use changes in the mechanism by which buyers can report on seller performance, to extract effects, which we interpret as seller adverse selection and seller moral hazard, and separate between the two.

Specifically, in May 2008, eBay changed its established non-anonymous feedback system from bilateral to essentially unilateral ratings, by allowing sellers to evaluate buyer behavior only positively, rather than also neutrally or negatively as before. With this, eBay dismissed with buyer fear of seller strategic retaliation to negative feedback given by buyers which, by eBay's own argument, resulted in buyer opportunism when rating sellers.

One year before the change in focus, eBay had introduced unilateral anonymous Detailed Seller Ratings that already allowed buyers to rate sellers without a bias generated by fear of seller retaliation to a negative rating—but retained the classic rating that, because non-anonymous, could be opportunistically biased. We use the unbiased Detailed Seller Rating as the basis for checking the effect of removing buyer reporting bias via the May 2008 change in the classic rating mechanism.

We show that this change resulted in improved seller ratings by buyers, but no exit of poorly rated sellers. In fact, the poorly rated sellers' ratings improved more than average. Towards the explanation of these results preferred by us, we develop a toy model that focusses on the effects of this natural experiment, from which we derive first, a reduction in seller moral hazard in preparing and executing the transaction; and second, a reduction in seller adverse selection, as generated by the increased exit of poorly performing sellers from the market. Within this context, we can interpret the absence of an effect on seller adverse selection by the relatively low cost of improving on moral hazard.

We check this explanation against a number of competing ones coming to mind, and develop tests supporting the one given by us.

Our results are derived within the context of an anonymous internet market—a type of market that obviously increases in importance, and is likely to dominate in transactions volume

classical markets in the foreseeable future. From a business policy point of view, our analysis provides for a fine example in which relatively small changes in the design of an information mechanism can have a significant effect. From the point of view of academic research, our result is, to the best of our knowledge, the first in which the effects of reducing buyer–seller informational asymmetries on adverse selection and moral hazard are clearly separated and directly juxtaposed.

References

- ABBRING, J., P. CHIAPPORI, AND J. PINQUET (2003): “Moral hazard and dynamic insurance data,” *Journal of the European Economic Association*, 1(4), 767–820.
- ABBRING, J., J. HECKMAN, P. CHIAPPORI, AND J. PINQUET (2003): “Adverse selection and moral hazard in insurance: Can dynamic data help to distinguish?,” *Journal of the European Economic Association*, 1(2-3), 512–521.
- AKERLOF, G. A. . (1970): “The Market for Lemons: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, 84, 488–500.
- ANDERSON, M., AND J. MAGRUDER (2012): “Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database,” *Economic Journal*, 122(563), 957–989.
- ARON-DINE, A., L. EINAV, A. FINKELSTEIN, AND M. CULLEN (2012): “Moral hazard in health insurance: How important is forward looking behavior?,” NBER Working Paper No. 17802, National Bureau of Economic Research, Cambridge, MA, U.S.A.
- AVERY, C., P. RESNICK, AND R. ZECKHAUSER (1999): “The market for evaluations,” *American Economic Review*, 89(3), 564–584.
- BAJARI, P., H. HONG, AND A. KHWAJA (2006): “Moral hazard, adverse selection and health expenditures: A semiparametric analysis,” NBER Working Paper 12445, National Bureau of Economic Research, Cambridge, MA, U.S.A.
- BAJARI, P., AND A. HORTASCSU (2004): “Economic Insights from Internet Auctions,” *Journal of Economic Literature*, 42(6), 457–489.
- BOLTON, G., B. GREINER, AND A. OCKENFELS (forthcoming): “Engineering trust-reciprocity in the production of reputation information,” *Management Sc.*
- CABRAL, L. (2012): *Reputation on the Internet* chap. 13, pp. 343–354. Oxford University Press, New York, NY, U.S.A.
- CABRAL, L. M., AND A. HORTASCSU (2010): “The Dynamics of Seller Reputation: Theory and Evidence from eBay,” *Journal of Industrial Economics*, 58(1), 54–78.

- CHIAPPORI, P. (2000): “Econometric models of insurance under asymmetric information,” in *Handbook of Insurance*, ed. by G. Dionne, Huebner International Series on Risk, Insurance, and Economic Security, pp. 365–393. Kluwer, Boston, MA, U.S.A.
- CHIAPPORI, P., AND B. SALANIÉ (2000): “Testing for asymmetric information in insurance markets,” *Journal of Political Economy*, 108(1), 56–78.
- DELLAROCAS, C., AND C. WOOD (2008): “The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias,” *Management Science*, 54(3), 460–476.
- DRANOVE, D., AND G. Z. JIN (2010): “Quality Disclosure and Certification: Theory and Practice,” *Journal of Economic Literature*, 48(4), 935–963.
- EINAV, L., A. FINKELSTEIN, S. RYAN, P. SCHRIMPF, AND M. CULLEN (2011): “Selection on moral hazard in health insurance,” NBER Working Paper No. 16969, National Bureau of Economic Research, Cambridge, MA, U.S.A.
- FANG, H., M. KEANE, AND D. SILVERMAN (2006): “Sources of advantageous selection: Evidence from the Medigap insurance market,” NBER Working Paper No. 12289, National Bureau of Economic Research, Cambridge, MA, U.S.A.
- FEIRING, D. (2009): “Best Match Made Simple,” available at <http://www.auctioninsights.info/bestmatchbook>.
- FINKELSTEIN, A., AND J. POTERBA (2004): “Adverse selection in insurance markets: Policyholder evidence from the UK annuity market,” *Journal of Political Economy*, 112(1), 183–208.
- HOUSER, D., AND J. WOODERS (2006): “Reputation in Auctions: Theory, and Evidence from eBay,” *Journal of Economics and Management Strategy*, 15(2)(2), 353–370.
- JIN, G., AND A. KATO (2007): “Dividing Online and Offline: A Case Study,” *Review of Economic Studies*, pp. 981–1004.
- (2008): “Price, quality, and reputation: Evidence from an online field experiment,” *The RAND Journal of Economics*, 37(4), 983–1005.
- KLEIN, T., C. LAMBERTZ, G. SPAGNOLO, AND K. STAHL (2006): “Last minute feedback,” C.E.P.R. Discussion Paper No. 5693, Centre for Economic Policy Research, London, U.K.

- KLEIN, T. J., C. LAMBERTZ, G. SPAGNOLO, AND K. O. STAHL (2009): “The Actual Structure of eBay’s Feedback Mechanism and Early Evidence on the Effect of Recent Changes,” *International Journal of Electronic Business*, 7(3), 301–320.
- LIVINGSTON, J. (2005): “How Valuable is a Good Reputation: A Sample Selection Model of Internet Auctions,” *Review of Economics and Statistics*, 87(3)(3), 453–465.
- LUCKING-REILEY, D., D. B. N. P., AND D. REEVES (2007): “Pennies from eBay: The Determinants of Price in Online Auctions,” *Journal of Industrial Economics*, 55(2), 223–233.
- MELNIK, M., AND J. ALM (2002): “Does a Seller’s Reputation Matter? Evidence from eBay Auctions,” *Journal of Industrial Economics*, 50(3), 337–349.
- RESNICK, P., AND R. ZECKHAUSER (2002): “Trust among strangers in Internet transactions: Empirical analysis of eBay’s reputation system,” in *The Economics of the Internet and E-Commerce*, ed. by M. Baye, vol. 11, pp. 127–157. Elsevier Science, Amsterdam, The Netherlands.
- RESNICK, P., R. ZECKHAUSER, J. SWANSON, AND K. LOCKWOOD (2006): “The Value of Reputation on eBay: A Controlled Experiment,” *Experimental Economics*, 9(2), 79–101.

A Further Details on the Data Collection

As described in Section 4, we drew sellers from 5 product categories. They are given in Table 5.

Recall that DSR scores are averages of DSRs received in the previous 12 months. Figure 8 shows the points in time at which we have collected data and the corresponding 12 month periods the DSR scores are calculated for.

Figure 9 shows how many DSRs and classic feedbacks were received in the preceding year and over time. The number of DSRs received in the previous 12 months increases until May 2008 because DSRs were only introduced in May 2007. Since then, the number is relatively stable. The figure shows that on average (across users), more feedbacks are received than DSRs. The reason for this is that for this figure we have counted DSRs as zero when they were not displayed and DSRs are only displayed if at least 10 DSRs were left in the previous 12 months.

Table 5: Product categories

Product categories
Home → All Categories → Computers & Networking → Laptops, Notebooks
Home → All Categories → Consumer Electronics → Apple iPod, MP3 Players
Home → All Categories → Toys & Hobbies → Model RR, Trains
Home → All Categories → Collectibles → Trading Cards
Home → All Categories → Home & Garden → Food & Wine
<i>Notes:</i> As of February 2008.

Figure 8: Data collection

Notes: The shaded area depicts the time interval DSR scores are covering. The change to the classic feedback system was implemented during the month of May 2008.

Figure 9: Number of classic feedbacks and DSRs received in the previous 12 months

Notes: The blue bars show the average number of classic feedback received in the previous year and the orange bars show the average number of DSRs received. Both averages are across sellers. The number of DSRs was counted as zero when less than 10 DSRs were received in the previous 12 months, respectively. The vertical line depicts the time of the May 2008 change to the classic feedback system. The parts of the bars that are in lighter color are not directly computed from data but come from linear interpolations of the original data, before averaging across users.

B Additional Tables and Figures

Figure 10: Average DSR score by category

Notes: This figure shows how the different DSRs changed over time. Figure 2 shows how the average of those four DSRs changed over time and also reports error bars for that average. They are suppressed here for the sake of clarity. As in the aforementioned figures, the dots are averages across users. Circles are linearly interpolated values for the periods in which we did not collect data. The vertical line denotes the time of the May 2008 change to the classic feedback mechanism. The horizontal dashed lines visualize that the dots are averages over the 12 months prior to the point in time at which the DSRs are displayed.

Table 6: Effect of Best Match and feedback change by decile of DSR rating

	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
number sellers	459	443	450	478	423	451	554	347	451	450
number observations	5,115	6,091	6,177	6,888	6,234	6,700	8,180	5,122	6,554	6,058
<i>effects on the number of DSRs received</i>										
average number DSR before March 2008	44.44*** (3.48)	69.00*** (3.16)	54.17*** (1.50)	41.82*** (3.93)	47.64*** (1.41)	31.97*** (1.15)	34.76*** (1.00)	28.71*** (0.94)	31.53*** (0.85)	18.29*** (0.71)
effect of Best Match	-5.76*** (1.50)	-0.09 (1.33)	-0.72 (0.67)	-2.02* (1.22)	-0.24 (0.48)	0.10 (0.43)	0.26 (0.45)	0.02 (0.38)	0.10 (0.36)	0.18 (0.30)
effect of feedback change	3.09** (1.26)	-2.13 (1.39)	0.14 (0.74)	0.50 (0.70)	0.15 (0.46)	-0.17 (0.36)	-0.56 (0.38)	-0.11 (0.35)	-0.20 (0.33)	-0.15 (0.24)
fixed effects	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
R^2	0.04	0.02	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00
<i>effects on the DSR ratings received</i>										
average DSR received before March 2008	4.32*** (0.01)	4.55*** (0.01)	4.63*** (0.00)	4.68*** (0.00)	4.72*** (0.00)	4.75*** (0.00)	4.78*** (0.00)	4.81*** (0.00)	4.84*** (0.00)	4.89*** (0.00)
effect of Best Match	-0.03 (0.05)	-0.03 (0.04)	0.00 (0.02)	0.00 (0.02)	0.04*** (0.01)	0.02 (0.01)	0.00 (0.01)	0.00 (0.01)	-0.01 (0.01)	-0.01 (0.01)
effect of feedback change	0.20*** (0.05)	0.14*** (0.04)	0.09*** (0.02)	0.07*** (0.02)	0.03*** (0.01)	0.04*** (0.01)	0.05*** (0.01)	0.03*** (0.01)	0.02* (0.01)	-0.01 (0.01)
fixed effects	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
R^2	0.08	0.08	0.13	0.11	0.19	0.15	0.11	0.06	0.01	0.02

Notes: The upper part of this table shows the results of regressions of the number of DSR received per month on the fraction of months in which BM and the new feedback system were in place, respectively. The lower part shows regressions of the average DSR score, averaged over the four categories, on a constant term and the fraction of feedbacks received after the introduction of BM and the May 2008 change to the feedback mechanism, respectively. We assumed that 75% of the DSR in March, 2008 were received after the introduction of BM and 25% of the DSR received in May 2008 were received after the change to the classic feedback mechanism. One observation is a seller-wave combination. There are 2,337 sellers. Throughout, we control for fixed effects. The R^2 is the within- R^2 . Standard errors are cluster-robust at the seller level and significance at the 1 percent level is indicated ***.

Table 7: Effects with time delay

	(1) full sample	(2) small window	(3) time trend	(4) DSR < 4.75	(5) DSR ≥ 4.75
average DSR before change	4.7067*** (0.0006)	4.7034*** (0.0004)	4.7150*** (0.0035)	4.5919*** (0.0010)	4.8143*** (0.0006)
effect of feedback change	0.0589*** (0.0024)	0.0435*** (0.0052)		0.0921*** (0.0044)	0.0318*** (0.0021)
effect of feedback change until September 2008			0.0183** (0.0081)		
effect of feedback change after September 2008			0.0652*** (0.0180)		
linear time trend before change			0.0009** (0.0004)		
linear time trend after change			0.0001 (0.0019)		
fixed effects	yes	yes	yes	yes	yes
R^2	0.0583	0.0125	0.0606	0.0820	0.0459
number sellers	5,225	4,919	5,225	2,337	2,337
number observations	67,373	30,488	67,373	31,260	33,508

Notes: This table shows the results of regressions of the average DSR score, averaged over the four categories, on a constant term and the fraction of feedbacks received after May 2008. We assume that half of the feedbacks in May 2008 were received before the change and the other half after the change. In specification (2), we do not include observations after October 2008 and before March 2008. In specification (3) we distinguish between the effect until September 2008 and after that date, and also account for a piecewise linear time trend. Specification (4) includes only those sellers who had a DSR score below the median of 4.75 in May 2008 and (5) only those above the median. One observation is a seller-wave combination. Throughout, we control for fixed effects. The R^2 is the within- R^2 . Standard errors are cluster-robust at the seller level and significance at the 5 and 1 percent level is indicated by ** and ***, respectively.