

Market Transparency, Adverse Selection, and Moral Hazard

Online Appendix

Tobias J. Klein

Christian Lambertz

Konrad O. Stahl*

July 2015

Abstract

This online appendix consists of three parts. Online Appendix [B](#) provides details on the data collection. Online Appendix [C](#) discusses additional competing explanations. An additional robustness check is performed in Online Appendix [D](#).

*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: T.J.Klein@uvt.nl.

B Details on Data Collection

As described in Section 3, we drew sellers from five product categories. These are given in Table B1. We then followed these sellers over time, doing our best to keep track of changes of the user names. There are months in which we were not able to collect data on the first day of the month for technical reasons, and therefore have missing values in our data set, as visualized in Figure 2, for instance. It is important to keep in mind, however, that DSR scores are moving averages, and thus contain information on DSR ratings received in the previous 12 months. This is also illustrated in Figure B1. It shows the points in time at which we have collected data, and the corresponding 12 month periods the DSR scores are calculated for. The data collected in the 12 months after one month with missing data will therefore be useful to infer the flow of missing DSR ratings. We do so indirectly when performing regressions, as in Section 4. In order to calculate the corresponding weights, we assume that the ratings that were received between two points in time were spread evenly over the months.

We could directly re-calculate some information for classic ratings. This is because in addition to numbers on how much positive, negative, and neutral feedback a user has received over the last month, eBay also reports these numbers for the last six months, and for the last twelve months. Say we have collected data for some user on the first of January. We then have numbers for December, but also for July to December, and for January to December. Even without additional observations, we can already infer from those the numbers for January to June, January to November, and July to November.

Combining numbers for overlapping periods from different data collections, we were able to calculate numbers on classical positive, negative, and neutral feedback received for months when we could not collect data. The lightly colored lower bars in Figure B2 show in which months there was no collection. We could not recover DSR numbers for these months, as eBay does not report DSR statistics in a similar fashion. But for classical feedback, we were able to calculate 12-month-figures for all but one of these months, as reflected in the upper bars of that same figure.¹

We use these re-constructed data to analyze how the percentage of positive feedback evolves over time. For this, we do not use the percentage share prominently reported by eBay, but

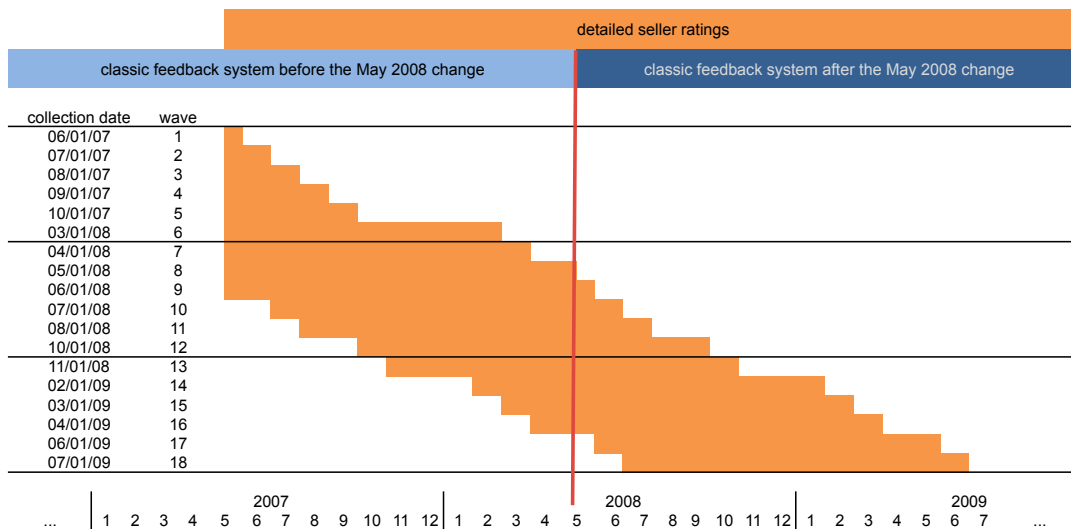
¹We were able to re-construct the number of classic ratings received in each of the months in the figure, but not for all months prior to the sample period. Therefore, there is one month of missing data in the figure.

Table B1: 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

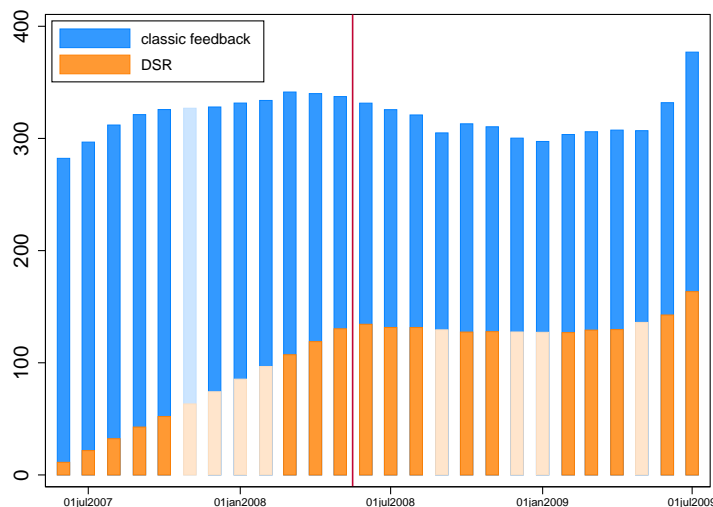
Notes: As of February 2008.

Figure B1: 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 B2: Number of classic feedbacks and DSRs received in the previous 12 months



Notes: Blue bars show the average number of classic feedback received in the previous year, orange bars the average number of DSRs received. Both averages are across sellers. Number of DSRs was counted as zero when less than 10 DSRs were received in the previous 12 months, respectively. Vertical line depicts the time of the May 2008 change to the classic feedback system. Parts of the bars that are in lighter color come from linear interpolations of the original data, before averaging across users.

calculate it from the raw numbers for positive, negative, and neutral feedback. This does not only provide us with numbers for months when we could not collect data. It is also necessary to get comparable numbers, whereas eBay repeatedly made changes to how they calculate the percentage share they report; or to what period it refers. In our analyses, it is consistently calculated as positive divided by the total number (positive plus negative plus neutral) of ratings received in the same period.

Figure B2 also shows how many DSRs were received in the preceding year and over time. It increases until May 2008 because DSRs were only introduced in May 2007. Since then, the number of DSRs received in the previous 12 months is stable over time. Moving forward one more month means that the number of DSR received thirteen months in the past does now not enter anymore and is replaced by the number of DSRs received in the previous month. The observation that the number of DSRs received in the previous 12 months does not change over time therefore means that the number of DSRs received in every single month is stable over time. The figure also shows that on average (across users), more feedbacks are received than DSRs. The reason for this is that DSRs can only be received when acting as a seller and when a classic rating is left (but not *vice versa*). Moreover, 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.

C Additional Competing Explanations for the Increases in DSRs

In addition to the three competing explanations discussed in Section 6.2, an alternative, fourth 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), eBay introduced a number of factors, by which it determined the sequence of listings. One of these factors was the DSR score. The introduction of BM thus provided an incentive for sellers to improve their performance.² 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.

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 percent, and more than 4 percent 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 C2 shows that the number of DSRs received after the introduction of BM

²That 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 2015.)

Table C1: 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: See notes to Table 2. Here, we additionally exclude the 10 percent worst sellers, as measured by their DSR score on March 1, 2008.

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 underlie the results in Table 2 and dropping the 10 percent worst sellers yields Table C1. The results are very similar, thus supporting our claim that our analysis is not affected by the introduction of BM.

This leads us to a fifth competing explanation. At the same time at which the BM ranking scheme was announced, eBay declared that it would provide fee discounts to PowerSellers with favorable DSR ratings (see <http://pages.ebay.com/sell/update08/overview/index.html>, last accessed June 2015). In particular, from July 2008 onwards PowerSellers received a 5 percent final value fee discount if they had received DSRs of 4.6 and above in the last 30 days, and a 15 percent final value fee discount if they had received DSRs of 4.8 or more in the last 30 days. Importantly, these incentives are only provided to a small group of (potential) high-volume PowerSellers. According to Table 1, only 7 percent of the sellers in our sample were PowerSellers on June 1, 2007, and therefore we do not expect this to overturn our main results. To assess this more formally, we excluded all those sellers from the sample who were observed to be a PowerSeller at least once. Then, we re-ran the regression underlying the results in column (1)

Table C2: Effect of Best Match and feedback change by decile of DSR rating

	decile of DSR score on March 1, 2008									
	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: Upper part of this table shows 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. 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 percent of the DSR in March, 2008 were received after the introduction of BM and 50 percent 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 10, 5 and 1 percent level is indicated by *, ** and ***, respectively.

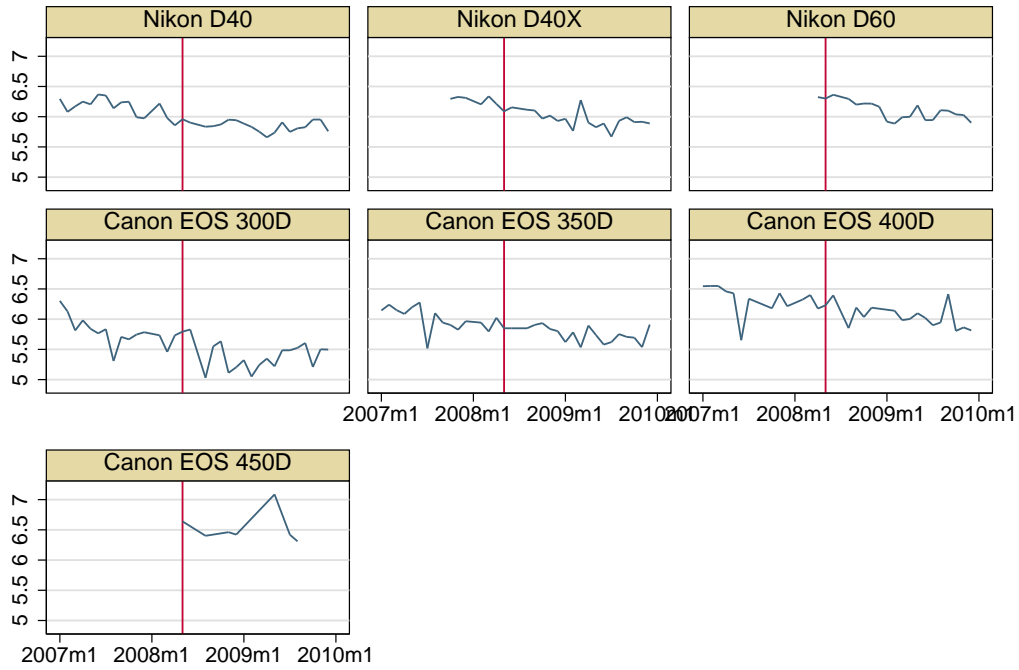
of Table 2. Based on 40,393 observations for 3,500 sellers we find that the effect of the feedback change is 0.0517, from a baseline level of 4.7079 and with a standard error of 0.0034, which is still highly significant and of very similar magnitude as reported for the whole sample.

A sixth alternative reason why ratings could have increased could be that buyer demand has shifted from auctions to fixed-price offerings, as documented by [Einav, Farronato, Levin, and Sundaresan \(2013\)](#), in particular in their Figure 1. As one can see there, however, the decrease was gradual at least until September 2008, while our Figure 2 shows that DSRs increased already before that, right after the change to the feedback mechanism.³ Moreover, and more importantly, one would not expect that a change in format should have an effect on buyer satisfaction as measured in our paper. The reason is that the DSR score we used is the average over the four DSR scores in the rating categories item description, communication, shipping charges and shipping speed. Arguably, none of them is related to whether or not the item has been offered in an auction. After all, our interpretation as based on our simple model equally applies to auctions.

Seventh and finally, one might argue that the time period around the change was one of great macroeconomic turmoil, and that this may have had important effects that we attribute to the change of the system. To begin, Figure 2 shows that before the feedback change there was no time trend in the DSR ratings, even though there was macroeconomic turmoil at that time. This already indicates that DSR ratings would not have increased had the classic feedback system not been changed. Towards further evidence against the influence of macroeconomic turmoil we collected data on average prices of a selection of camera models from the website Pixelpeeper. Figure C1 shows respective time series of the log of the average price on eBay against time. It already suggests that time effects did not play a role. To analyze this more formally, we have regressed log average prices on a full set of time and camera model indicators, as well as on indicators for the total number of models in the market and the number of models of the same brand. Figure C2 shows the fitted time trend and confirms that indeed, there has not been any structural break in the time effects. We complement this evidence with the number of searches for the camera models over time. Figure C3 does not show any structural breaks either. Taken together, we take this as evidence that even though the time period was one of great macroeconomic turmoil, this does not seem to have affected transactions on eBay in

³See their Section 7 for an explanation why the jump in September 2008 is mechanical.

Figure C1: Average auction prices for selected camera models




Graphs by model

Notes: Figure shows the log of average auction prices for entry-level digital single-lens reflex (DSLR) camera models from Nikon and Canon in North America. Collected from <http://www.pixel-peeper.com/>. Vertical line indicates the change to the feedback system in May 2008.

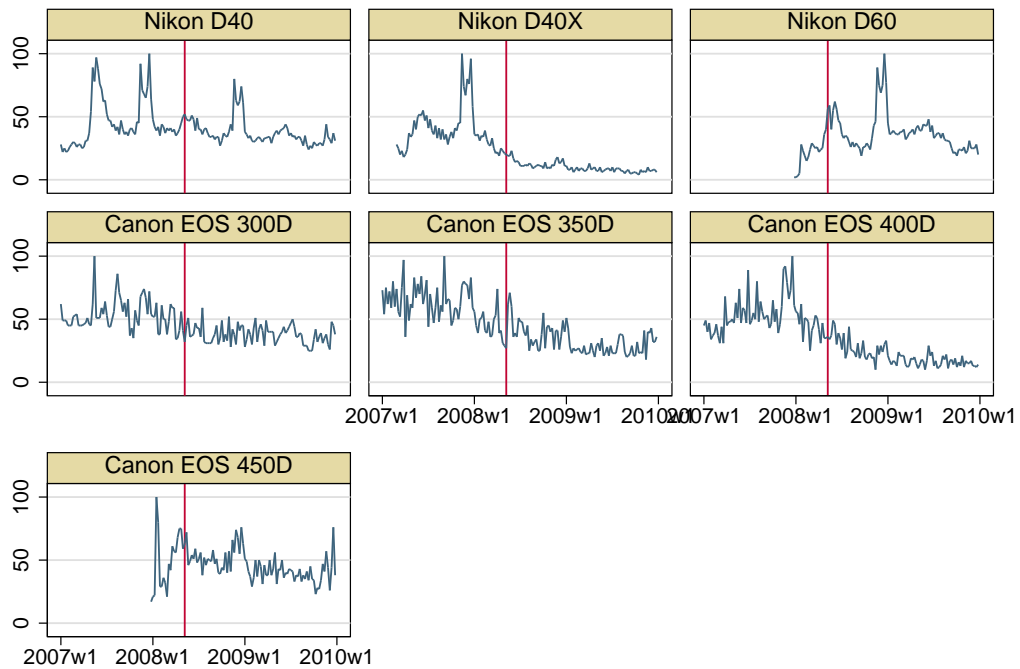
drastic ways. Therefore, macroeconomic developments are likely not an alternative explanation for our findings either.

Figure C2: Time trend in camera prices



Notes: Figure constructed from the data shown in Figure C1 and shows the time fixed effects controlled for camera dummies, time-on-the-market dummies for each model, number of models of the same brand on the market, and total number of models on the market. Bars indicate pointwise 95 percent confidence intervals, based on the reported cluster-robust standard errors, clustered at the level of the camera models. Vertical line indicates the change to the feedback system in May 2008.

Figure C3: Google searches for camera models



Graphs by model

Notes: Figure shows time profiles of the number of Google searches for different camera models in North America, available on a weekly basis, downloaded from the Google Trends website. Search terms used are "nikon d40," "nikon d40x," "nikon d60," "canon 300d," "canon 350d," "canon 400d," and "canon 450d." Number of searches rescaled by Google Trends so that highest number is 100 (percent) for each model. Vertical line indicates the change to the feedback system in May 2008.

D Delay between Transaction and Rating

We have performed one additional, more technical robustness check. It is related to the fact that we don't know the exact date of when eBay enacted the change whose effects we are reporting here. To see that this may be a problem in principle, 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 inasmuch unanticipated. Nevertheless, a feedback could have been left for that transaction in the second half of May 2008. In our analysis, we have assumed that half of the DSRs received in May 2008 corresponded to transactions conducted after the feedback change. Disregarding this reporting delay, we attribute the ratings after the change all to transactions thereafter, and with this tend to underestimate the effect of the feedback change. We do not expect this to have big effects, however, because the delay is likely to be small relative to the length of our observation period.

To assess whether this is indeed the case, we check the robustness of our results to changes in the delay we implicitly assume. We don't have a record on delays between announcement, 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 percent are left after 2 weeks, and almost 90 percent after 4 weeks.⁴ Based on this, we re-did the analysis relating to BM and the feedback change in Section 6.2, assuming that out of all DSRs received in March 2008, 75 percent of the transactions took place after the introduction of BM. Moreover, we assumed that out of all DSRs received in May 2008, 25 percent of the transactions took place after the change to the feedback system. The results were very similar. We also re-did the analysis underlying Table 2, assuming that 25 percent of the transactions took place after the change to the feedback system. Table [D1](#) shows the results. They are very similar.

⁴For 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. After May 2008, these strategic considerations were not important anymore because sellers could not retaliate negative feedbacks anymore.

Table D1: Effects with time delay

	(1) full sample	(2) small window	(3) time trend	(4) DSR < 4.75	(5) DSR \geq 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	5,224	4,919	5,224	2,337	2,337
number observations	67,376	30,488	67,376	31,260	33,508

Notes: See notes to Table 2. The difference between the two tables is that here, we assume that only 25 percent of the feedbacks received in May 2008 correspond to transactions that took place after the change.

References

- EINAV, L., C. FARRONATO, J. LEVIN, AND N. SUNDARESAN (2013): “Sales mechanisms in online markets: What happened to internet auctions?,” NBER Working Paper No. 19021, National Bureau of Economic Research, Boston, MA, U.S.A.
- FEIRING, D. (2009): “Best match made simple,” available at <http://www.auctioninsights.info/bestmatchbook>.
- 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.