Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

A Dollar for Your Thoughts: Feedback-Conditional Rebates on eBay

Luís Cabral, Lingfang (Ivy) Li

MANAGEMENT SCIENCE
Vol. 61, No. 9, September 2015, pp. 2052–2063
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)
http://dx.doi.org/10.1287/mnsc.2014.2074

A Dollar for Your Thoughts: Feedback-Conditional Rebates on eBay

Luís Cabral
Stern School of Business, New York University, New York, New York 10012, luis.cabral@nyu.edu

Lingfang (Ivy) Li
School of Management, Fudan University, Shanghai 200433, China, lingfangivy.li@gmail.com

We run a series of controlled field experiments on eBay where buyers are rewarded for providing feedback. Our results provide little support for the hypothesis of buyers' rational economic behavior: the likelihood of feedback barely increases as we increase feedback rebate values; also, the speed of feedback, bid levels, and the number of bids are all insensitive to rebate values. By contrast, we find evidence consistent with reciprocal buyer behavior. Lower transaction quality leads to a higher probability of negative feedback as well as a speeding up of such negative feedback. However, when transaction quality is low (as measured by slow shipping), offering a rebate significantly decreases the likelihood of negative feedback. All in all, our results are consistent with the hypothesis that buyers reciprocate the sellers' "good deeds" (feedback rebate, high transaction quality) with more frequent and more favorable feedback. As a result, sellers can "buy" feedback, but such feedback is likely to be biased.

Data, as supplemental material, are available at http://dx.doi.org/10.1287/mnsc.2014.2074.

Keywords: feedback; rebates; field experiment
History: Received September 18, 2012; accepted August 24, 2014, by Teck-Hua Ho, behavioral economics. Published online in *Articles in Advance* March 13, 2015.

1. Introduction

Studies of eBay buyer behavior have shown that a seller's past experience, as measured by the number and quality of past feedback postings, is an important determinant of a seller's success, both in terms of number of bidders and size of bids (Cabral 2012). This is important for the seller, insofar as better reputation leads to greater sales; it is also important for the platform owner, to the extent that better information leads to more sales, which in turn lead to higher revenues from selling fees. All in all, customer feedback is an important component of a firm's strategy, be it a seller or an intermediary.

In this paper, we consider one possible seller strategy for eliciting buyer feedback: to provide buyers with a rebate conditional on rating the quality of their transaction. The rebate will be provided as long as a feedback rating is left, regardless of whether the feedback is positive or negative.[1] By establishing two eBay sellers who auction the same homogeneous good (a USB pen drive), we are able to run a series of controlled field experiments where we vary the degree to which buyers are rewarded for feedback, as well as a component of transaction quality (speed of shipment).

Our experiment is motivated by a series of hypotheses derived from two paradigms of buyer behavior. The first, which we refer to as the *homo economicus* paradigm, implies that buyers are more likely to provide feedback when feedback is rewarded, to do so more quickly, and to bid higher in anticipation of a feedback reward. The second, which we refer to as the *homo reciprocus* paradigm, implies that buyers are more likely to give favorable feedback (less negative, more positive) when sellers offer a feedback rebate or higher transaction quality.

Our field experiment addresses these and other related questions. We find relatively little evidence for the *homo economicus* paradigm: the likelihood of feedback increases marginally as we increase feedback rebate values, and the speed of feedback seems relatively insensitive to the size of the feedback rebate, as do bid levels and the number of bids.

By contrast, we find fairly strong evidence consistent with the *homo reciprocus* paradigm. In particular, we find that when transaction quality is low (as measured by shipment delay), offering a rebate decreases the likelihood of a negative feedback message. We also find that lowering transaction quality leads to a higher probability of negative feedback as well as a

[1] In other words, the rebate is unbiased because it is not offered for one type of feedback only.

speeding up of such negative feedback. All of these results are consistent with the general hypothesis that buyers reciprocate the sellers' "good deeds" (feedback rebate, high transaction quality) with more frequent, more favorable feedback.

In sum, our results suggest that a seller can buy more buyer feedback, but such feedback is likely to be biased, although the rebate is provided regardless of whether the feedback is positive or negative. By contrast, contrary to economic theory predictions, we show that feedback rewards have no effect on bidding behavior, either the number of bidders or the average bid. This can be explained by buyer myopia, inattention, or incredulity with respect to the seller's feedback reward offer. It also suggests that the cost of obtaining more customer feedback is greater than a rational, forward-looking model would predict.

Our focus on feedback bias may seem misplaced: if sellers can get more positive feedback by offering feedback rewards, then that's all that matters, one might argue. We reply to that objection in two ways. First, although the "direct" effect of positive feedback is clearly positive, to the extent that inaccurate feedback may reduce buyers' confidence, sellers may also be harmed by biased feedback. If only one seller offers feedback rebates, this effect is unlikely to be very important. However, in a situation where feedback rebates are more general, the cost of inaccurate feedback may be significant.

This leads us to the second point: from the perspective of a platform owner (e.g., eBay or Taobao), feedback bias matters: if feedback is not accurate, then the value of the feedback system is lower, and so is buyers' trust.[2] Our results cannot directly address the effect of a system-wide feedback rebate; this would require eBay itself to run the field experiment. In this sense, our approach is a second best. Still, we believe our results shed some light on how buyers respond to a rebate mechanism.

More generally, our results point to the potential negative effect of reciprocity. Buyers may be "too good for their own good," as it were: in their eagerness to reciprocate an act of kindness (feedback rebate) they treat the seller better than the transaction quality justifies, thus contributing to the lowering of the overall value of the feedback system, which in the process may harm buyers, the platform owner, and possibly even sellers.

[2] To be fair, if bias is systematic (i.e., every seller's feedback is inflated by the same level), then the loss in informational value may not be very high. However, it is possible that greater bias is also associated with noisier feedback (this will be the case if feedback inflation levels vary across sellers).

1.1. Related Literature

A growing theoretical and empirical literature shows that a seller's reputation history is an important determinant of a seller's success, especially in online markets.[3] In online markets, reputation systems usually rely on voluntary feedback from the parties involved. This creates a problem of public good underprovision, for which various solutions have been proposed. For example, Miller et al. (2005) and Jurca and Faltings (2007) propose truth-eliciting incentive schemes to induce buyers to report and do so honestly. These mechanisms either require buyers or the market-maker (e.g., eBay) to bear the reporting cost. Given the importance of customer feedback, sellers have an interest in encouraging buyers to post a review, so sellers may have more incentive to bear the reporting cost than buyers or the market.

Addressing the issue of incentive provision, Li (2010a) proposes and theoretically analyzes a rebate mechanism in an online auction market: sellers have the option of committing by providing a rebate (not necessarily in monetary form) to cover the buyer's reporting cost, regardless of whether the feedback is positive or negative. In theory, this rebate mechanism plays the dual role of incentivizing buyers to leave feedback and providing a device for sellers to signal quality or effort to cooperate. In equilibrium, buyers avoid sellers who do not choose the rebate option and incorporate the rebate amount into their bids.

Although providing monetary incentives is not the only way of inducing the desired buyer behavior, it is still one of the easier strategies for sellers to implement.[4] Using monetary rebates as incentives for feedback, Li and Xiao (2014) conduct a laboratory experiment to examine the effect of the rebate mechanism on market efficiency in a listed-price market.[5] They find that a seller's rebate offer increases the buyers' propensity to report in good transactions but not in bad transactions; market efficiency under the rebate

[3] See Shapiro (1983), Avery et al. (1999), Dellarocas (2003), Bolton et al. (2004), Houser and Wooders (2006), Jin and Kato (2006), Resnick et al. (2006), Cabral and Hortaçsu (2010), and Grosskopf and Sarin (2010).

[4] For instance, Abeler et al. (2010) find that apology is more effective than monetary incentives in making buyers withdraw their negative feedback on eBay. Chen et al. (2010a) run a field experiment on MovieLens.com and find that effective personalized social information can increase the level of public goods provision. Using data from a major online travel agency in China, Gu and Ye (2014) study how the online management responses affect customers' feedback. Alternative ways of motivating agents are also explored in the work of Bénabou and Tirole (2003), Ariely et al. (2009), Chen et al. (2010b), and Wang (2010).

[5] In Li and Xiao (2014), market efficiency is measured by number of efficient trades (i.e., the case where the buyer bought and the seller shipped the product) as well as earnings of buyers and sellers for each treatment.

mechanism increases with the probability that sellers will provide a rebate; and the dollar rebate does not affect the honesty of the feedback reports.

In this paper, we run a field experiment on eBay to test one possible seller strategy: to provide buyers with a monetary rebate conditional on rating the quality of their transaction. We are interested in investigating whether paying for feedback induces buyers to give more feedback, whether buyers bid higher when there is a rebate, and most importantly, whether the nature of buyer feedback is altered by the offer of a monetary reward by the seller. This is a natural next step with respect to Li (2010a), a theory paper, and Li and Xiao (2014), a laboratory experiment paper. With respect to laboratory experiments, field experiments have the virtue of applying to a real-world situation rather than a laboratory setting. Compared to other papers on field experiments, such as Abeler et al. (2010), this paper tests a mechanism hitherto not considered (in field experiments), namely conditional rebates.

1.2. Road Map

This paper is organized as follows. Section 2 derives theoretical predictions in the form of testable hypotheses. In §3, we describe the experimental design. The results are presented in §4, which includes both basic tabulations and regression results. Finally, §5 concludes the paper.

2. Theory and Hypotheses

Economics and psychology provide behavioral paradigms with specific predictions regarding bidding and feedback patterns. In this section, we consider two stylized models corresponding to the paradigms of economic rationality and reciprocal behavior: *homo economicus* (HE) and *homo reciprocus* (HR). We should state from the outset that the two models below are not mutually exclusive. In fact, actual observed behavior will feature a little bit of each model.

Let us first consider the paradigm most closely related to economic rationality, *homo economicus*. This paradigm postulates that buyers make choices so as to maximize utility minus monetary and transactions costs. Suppose that giving feedback implies a cost c, where c is buyer specific (and possibly time specific as well, as we will consider later in this section). Then we would expect the buyer to provide feedback if and only if $r > c$, where r is the rebate value, which leads to the following hypothesis.

Hypothesis HE1. *The likelihood that feedback is given is increasing in rebate value.*

Regarding the nature of the feedback, economic rationality has little to say. One possibility is that the buyer gives random feedback. Another possibility is that the buyer gives true feedback, that is, feedback that accurately corresponds to the quality of the transaction.[6] Either way, as far as economic rationality is concerned, the nature of feedback should be uncorrelated with the rebate value.

Suppose that the buyer's cost of providing feedback is time dependent and given by c_t, which follows some stochastic process. A rational buyer will then implement an optimal stopping time at which to give feedback. In this context, a higher value of r implies not only a higher likelihood of feedback (at some point) but also an earlier stopping time, as described specifically in the following hypothesis.

Hypothesis HE2. *The delay between delivery time and feedback time is decreasing in rebate value.*

Finally, the economic rationality approach also has something to say about bidding behavior. Since conditional rebates are announced before the auction takes place, a rational bidder, havingplaced the highest bid and won the auction in question, anticipates having the option of receiving a feedback rebate. This option is worth

$$z(r) = \max\{r - \min\{c_t\}_{t=0}^{T}\},$$

where T is the number of days when feedback can be given. Assuming that the ex ante distribution of values of c_t is sufficiently dispersed, we have $0 < z(r) < r$. It follows that an economically rational bidder should anticipate this option and increase the bid by an amount $z(r)$. Noting that $z(r)$ is increasing in r, this implies that bids should be increasing in r. To the extent that there is an auction entry cost that varies from bidder to bidder, the above argument also implies that the number of bidders should be increasing in r.

Hypothesis HE3. *Bid values and the number of bidders are increasing in rebate value.*

We now turn to the reciprocal behavior paradigm. It has been theorized and tested (both in the laboratory and in the field) that reciprocity plays an important role in human behavior.[7] In the present context, we propose a simple model of feedback behavior that incorporates this behavioral pattern. The main building block is the assumption that buyers receive utility

[6] This is the assumption made by Cabral and Hortaçsu (2010), for example.

[7] For example, Fehr et al. (1993) test a gift-exchange game in the laboratory and find that higher wages offered by an employer lead to considerably more effort provision (where effort is costly). In a field experiment, Gneezy and List (2006) find that positive reciprocity vanishes over time. Also in a field experiment, Falk (2007) finds that the relative frequency of donations increases by 17% if a small gift is provided to potential donors (75% for a large gift). The role of reciprocity is also explored in the work of Fehr and Gächter (2000b) and Dellarocas and Wood (2008).

from giving feedback and that this utility is a function of transaction quality as well as the rebate value. The idea is that the quality of a transaction as well as the feedback rebate are seen as "gifts" by the seller, whereas feedback is the buyer's way to reciprocate such gifts.

Specifically, we assume that the buyer receives utility u_P and u_N from giving positive and negative feedback, respectively, where

$$\begin{aligned} u_P &= \alpha_0(q-\bar{q}) + \beta_0 r - c, \\ u_N &= \alpha_1(\bar{q}-q) - \beta_1 r - c, \end{aligned} \tag{1}$$

where α_i, β_i are parameters, q is transaction quality, $\bar{q}$ is reference quality level, r is the rebate level, and c is the cost of giving feedback. The idea is that q and r are "good deeds" done by the seller for the buyer. The more of these good deeds the buyer receives, the more the buyer is willing to reciprocate with a "good" act, namely a positive feedback message, and the less the buyer is willing to reciprocate with a "bad" act, namely a negative feedback message. Finally, the buyer receives $u_\varnothing = 0$ from not giving any feedback.

Homo reciprocus chooses action x that maximizes u_x, where $x \in \{P, N, \varnothing\}$. Specifically, *homo reciprocus* follows the feedback procedure:

$$\begin{aligned} &\text{positive feedback} && \text{if } u_P > u_N,\ u_P > 0, \\ &\text{negative feedback} && \text{if } u_N > u_P,\ u_N > 0, \\ &\text{no feedback} && \text{if } u_P < 0,\ u_N < 0. \end{aligned} \tag{2}$$

In this context, an increase in r has an ambiguous effect on the probability of (some) feedback: on the one hand, it increases u_P and so increases the probability of positive feedback; on the other hand, it decreases u_N and so decreases the probability of negative feedback.

Although it is not possible to make a clear prediction regarding the likelihood of feedback, the *homo reciprocus* model implies a clear prediction regarding the nature of feedback. Since u_P is increasing in r and u_N is decreasing in r, we have the following hypothesis.

Hypothesis HR1. *The relative likelihood of negative feedback with respect to positive feedback is decreasing in r.*

So far we have considered the effects of a higher reward for feedback. As mentioned earlier, in our field experiment we also control for transaction quality q, in the form of fast or slow shipping time. From (2), we see that a lower value of q increases u_N and decreases u_P, implying the following.

Hypothesis HR2. *The relative likelihood of negative feedback with respect to positive feedback is higher when transaction quality is lower.*

Table 1 Summary of Theoretical Hypotheses Regarding the Effect of an Increase in Conditional Feedback Rebate

	Rational	Reciprocal
Increase in rebate value		
Probability of feedback	+	
Nature of feedback (N/P)		−
Speed of feedback	+	
Bid level	+	
Decrease in transaction quality		
Nature of feedback (N/P)		+
Speed of negative feedback		+

Notes. "Nature of feedback" refers to the relative weight of positive and negative feedback, whereas "Speed of feedback" refers to an inverse measure of the number of days until the buyer provides feedback.

As for the speed of feedback, the *homo reciprocus* model is as ambiguous as it is regarding the probability of some feedback. When the reward for feedback is increased, we expect positive feedback to arrive earlier and negative feedback to arrive later, but there is no clear prediction regarding the arrival of some feedback.

Regarding transaction quality, (2) implies that a decrease in q leads to an increase in u_N, which by (2) implies an increase in the likelihood of negative feedback. We single out this prediction because previous literature has shown that there is such a thing as "demand for justice."[8] In the present context, this implies that the coefficient α_1 is particularly high: if buyers feel that they were poorly treated in the transaction (low q), then they get a high utility from reciprocating such bad behavior, and do so by providing negative feedback. In addition to Hypothesis HR2 (negative feedback is more likely), we also expect a speeding up of such negative feedback.

Hypothesis HR3. *Lower transaction quality leads to faster arrival of feedback.*

Finally, we note that, regarding bidding behavior, the reciprocity model is silent; that is, it implies no specific prediction (unlike the economic rationality model, which, as we saw earlier, predicts higher bids and numbers of bidders).

Table 1 summarizes the main theory predictions in terms of the effect of an increase in conditional feedback rebate, as well as a decrease in transaction quality, the two control variables we use in our field experiment. The two rightmost columns correspond to the two paradigms we consider, economic rationality and reciprocal behavior. A "+" (or "−") sign represents a positive (or negative) effect of a change in the control variable (feedback rebate increase, transaction quality decrease). A blank cell signifies that

[8] See, for example, Brandts and Charness (2003), Charness and Levine (2007), de Quervain et al. (2004), Fehr and Gächter (2000a), and Xiao and Houser (2005).

Figure 1 (Color online) Kingston 2 GB USB Pen Drive

the paradigm in question implies no particular prediction. From Table 1, we see that the HE and HR paradigms are not necessarily incompatible. In fact, we would expect reality to be a combination of both paradigms: we also have economic sense as well as a sense of reciprocity. Still, an interesting question is the relative importance of each paradigm. We next turn to the testing of our theoretical hypotheses.

3. Field Experiment Design

To examine the effects of monetary rebates on feedback behavior, we sold new Kingston 2 GB USB pen drives on eBay (Figure 1). We chose to sell this particular product because it is a relatively standard product and it is sold by several other sellers.[9] We chose to conduct the field experiment on eBay because it is the world's largest online auction market and has been the object of numerous complementary studies and experiments.[10]

We registered two IDs on eBay and accumulated 75 positive feedback scores on each ID from buying and selling the Kingston pen drive. In this way, before beginning our experiment, we had two sellers with similar, established records. So as to avoid being identified as "experimental" sellers, we operated the two IDs on different days. Considering the large number of sellers of the same object, our two IDs typically did not show on the same search page.

We created several treatments, with characteristics that vary along two dimensions. First, in different treatments we offered different levels of feedback rebate: \$0, \$1, or \$2 per feedback. We clearly stated the rebate amount of \$1 or \$2 in the item listing title. Specifically, we used the title "Brand-New Kingston 2 GB USB Flash Drive" for a no-rebate listing and "Brand-New Kingston 2 GB USB Flash Drive (\$1 rebate available)" for a \$1 rebate listing (and the analogously titled one for a \$2 rebate listing). In the listings with a rebate, we added the sentence:

> Rebate option. Your feedback is important to us; please give us your honest feedback and receive a \$1 credit in your Paypal account.

[9] One difficulty we experienced was that Kingston stopped producing the 2 GB USB drives halfway through our experiment. We tried the best we could to purchase the same model USB drives around the world so as to continue selling the same object.

[10] See, for example, Dellarocas and Wood (2008), Brown and Morgan (2006), Resnick et al. (2006), Brown et al. (2010), and Li (2010b).

Table 2 Field Experiment Treatments

	No rebate	\$1 rebate	\$2 rebate
Fast shipment	F0	F1	F2
Slow shipment	S0	S1	S2

Table 3 Field Experiment Chronology

Phase	Dates	Seller ID	Type	# obs.
1	2010/02/23–2010/05/08	1	F1*, F0	18, 12
		2	F1*, F0	15, 15
2	2010/06/10–2010/08/19	1	F1, F0	14, 16
		2	F1, F0	15, 15
3	2010/10/06–2011/01/11	1	F2, F0	15, 15
		2	S1, S0	15, 16
4	2011/04/06–2011/09/01	1	S2, S0	16,[a]4

[a] In phase 4, our research assistant (RA) mistakenly put \$1 instead of \$2 for one transaction, so among the 16 \$2 rebate transactions, one is actually a \$1 rebate transaction. In our data analysis, we treat it as a \$1 rebate transaction.

The second dimension that distinguishes different treatments is transaction quality. Specifically, we provided the identical USB drives with different speeds of shipment. For a "fast" transaction, we shipped the USB drive immediately upon receiving payment. For a "slow" transaction, we shipped the USB drive two weeks after sale.

Together, rebate value $(0, 1, 2)$ and shipment speed (F, S) create six different possibilities, as listed in Table 2. For example, treatment F1 corresponds to fast shipping speed and a \$1 rebate, whereas treatment S0 corresponds to slow shipping and no rebate offered.

Our experiment may be chronologically divided into four phases, as listed in Table 3. In phase 1, our research assistant (RA) mistakenly offered a \$1 rebate right after receiving payment and *before* buyer feedback was received. In §4, we discuss how we treated this data. By the time we started phase 2, both of our eBay IDs had more than 100positive feedback scores. In phase 2 we offered the same feedback reward as in phase 1, but we consistently made such rebate conditional on receiving buyer feedback. Thus, phase 2 corresponds to the F1 treatment (and the F0 treatment as well, for we only offered a feedback reward on some transactions).

In phase 3, our two hitherto similar sellers took different paths. For seller ID1 we switched to an F2 treatment; that is, we increased the value of the rebate from \$1 to \$2. For seller ID2 we switched to a two-week shipment while keeping the \$1 rebate. Finally, in phase 3 seller ID1 switched from fast to slow shipment while keeping the \$2 rebate.[11]

[11] Some buyers filed complaints with eBay regarding slow shipping by our seller ID2, so that it no longer met the minimal "detailed seller rating requirements" in the "seller performance standards."

Table 4 Descriptive Statistics (Phases 1–4)

Variable	*N*	Mean	SD	Min	Max
Rebate	201	0.687	0.719	0	2
Price	201	3.823	1.573	1.04	11
Bidder count	201	4.000	1.338	2	8
Bid count	201	5.841	2.227	2	14
Bidder score	201	350	700	0	6,454
Bidder positive perc.	201	0.983	0.107	0	1
Seller score	201	112	22	75	151
Seller positive perc.	201	0.998	0.007	0.956	1
Feedback	201	0.682	0.555	−1	1
Feedback lag (days)[a]	155	13.355	9.575	0	54

[a]For feedback-received transactions only.

Table 4 displays basic descriptive statistics of the data generated by the experiment's various phases. We should also mention that, of the 201 completed transactions, in only 3 instances did a buyer repeat a purchase from the same seller (there are a total of 17 repeat buyer sales in our data set, but most correspond to the same buyer purchasing from different sellers).

4. Results

In this section, we present the results from our experiment. We group them into several subsections, roughly following the list of hypotheses outlined in §2. As a preliminary subsection, we compare phase 1 (where feedback rewards were mistakenly given before feedback comments were made) with the remaining phases. Then, we investigate whether offering a feedback rebate induces buyers to give feedback more frequently. Next we look at the nature of feedback, that is, whether it becomes more favorable to the seller (conditional on transaction quality). Following that, we look at the timing of feedback. Finally, we focus on the bidding stage, both in terms of bid level and number of bidders.

4.1. Show Me the Money

As mentioned earlier, a communications error led our RA to mistakenly offer a \$1 rebate *before* receiving feedback during an initial stage of our experiment. We denoted this phase as phase 1 and initially decided not to use it for our statistical tests (to the extent that we are interested in the effect of *conditional* feedback rebates). However, we decided this was a good opportunity to "turn lemons into lemonade": our unintentional mistake provides a test of whether or not the conditionality of the feedback rebate (on actually receiving feedback) plays a role.

Both phases 1 and 2 contain 60 fast shipping transactions. In both cases we have 12 transactions with no feedback and the remaining 48 with positive feedback. In other words, in terms of frequency and type of feedback, the two phases are identical: the conditionality of feedback seems to have no effect on incentives to give feedback. Moreover, restricting our observations to the 48 transactions when feedback was given, we see that the mean number of days before feedback is given is marginally lower when the rebate is given after feedback. Specifically, we use a one-tailed t-test and find $p = 0.0610$.

Together, these results suggest that incentives do not play an important role, either because buyers are unaware or incredulous of the seller's offer, or because reciprocity considerations play a bigger role than economic incentives. According to Bénabou and Tirole (2003, p. 489),

> A central tenet of economics is that individuals respond to incentives. For psychologists and sociologists, in contrast, rewards and punishments are often counterproductive, because they undermine intrinsic motivation.

Our study suggests that the effects of conditional feedback go beyond those of economic incentives; on the contrary, they may consist primarily of motivational incentives. In fact, the analysis that follows largely confirms this suspicion, by providing ample support for the three HR hypotheses but very little for the three HE hypotheses.

Moreover, given the regularity of the effect of rewards on behavior, independent of whether they are given conditionally or unconditionally, in our regression analysis we aggregate data from all phases.

4.2. A Penny for Your Thoughts

The first research question in which we are interested is whether paying for feedback induces buyers to give feedback more frequently. Table 5 tabulates the frequency of feedback for different types of feedback policy. We restrict ourselves to fast shipment transactions, so as to control for quality. In this table, we limit ourselves to transactions in phase 2 by both sellers and phase 3 by seller 1, a total of 90 observations. As can be seen, the percentage of transactions where feedback is given increases from 76.09% to 79.31% as we move from no rebate to a \$1 conditional rebate, and from 79.31% to 93.33% as we switch from a \$1

Table 5 Feedback Behavior Fast Shipment Transactions

Feedback reward →	None		\$1		\$2		Total	
Outcome ↓	#	%	#	%	#	%	#	%
No feedback given	11	24	6	21	1	7	18	20
Positive feedback given	35	76	23	79	14	93	72	80
Total	46	100	29	100	15	100	90	100

As a result, ID2 was not operating during phase 4. As a result of the difficulty of getting the 2 GB USB pen drives as mentioned in Footnote 9, we have only 20 observations in phase 4.

Table 6 Regression Analysis of Feedback Behavior

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	*Any feedback*		*Negative feedback*		*Days until feedback*	
Regression type	Logit	Logit	OLS	OLS	Tobit	Tobit
$1 *rebate* ($\beta_1$)	0.065	0.015	−0.046	−0.012	−0.481	−1.591
	(0.067)	(0.076)	(0.038)	(0.041)	(1.744)	(1.951)
$2 *rebate* ($\beta_2$)	0.103	0.285	−0.121*	−0.003	0.658	1.234
	(0.123)	(0.183)	(0.069)	(0.081)	(3.204)	(3.855)
Slow shipping (β_3)	−0.341**	−0.336**	0.240**	0.406***	−10.63**	−12.54***
	(0.150)	(0.150)	(0.095)	(0.105)	(4.111)	(4.608)
$1 *rebate* × *Slow shipping* (β_4)		0.173		−0.213**		5.307
		(0.156)		(0.095)		(4.293)
$2 *rebate* × *Slow shipping* (β_5)		−0.440		−0.394		−1.278
		(0.288)		(0.143)***		(6.909)
Seller's score (β_6)	−0.053***	−0.052***	−0.018*	−0.019**	−0.931**	−0.923**
	(0.015)	(0.015)	(0.010)	(0.010)	(0.417)	(0.416)
Seller's perfect record (β_7)	0.078	−0.010	0.383***	0.403***	5.534	4.420
	(0.179)	(0.185)	(0.108)	(0.108)	(4.894)	(5.025)
Date (β_8)	0.013***	0.013***	0.006**	0.006**	0.224**	0.223**
	(0.003)	(0.004)	(0.003)	(0.003)	(0.113)	(0.113)
Seller and phase FE	Y	Y	Y	Y	Y	Y
N	201	201	155	155	201	201
Adj. R^2			0.256	0.302	0.0143	0.0156
Pseudo R^2	0.099	0.118				

Notes. Standard errors are in parentheses. The data used in columns (3) and (4) exclude no-feedback transactions, and the coefficients and standard errors reported in columns (1) and (2) are marginal effects at the means. FE, fixed effects.
* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

to $2 conditional rebate. In terms of statistical significance (t-test), the mean value of propensity to leave feedback in the $1 rebate treatment is not significantly higher than in the $0 rebate treatment ($p = 0.3747$), whereas it is significantly higher in the $2 rebate treatment than in the $0 rebate treatment ($p = 0.0747$). The results hold regardless of whether or not we include data from in the analysis.

In other words, our tabulation results suggest that feedback rewards induce higher feedback rates in fast shipping transactions, although weakly: when we offer $1 for customer feedback, we do not observe a significant increase in the feedback rate; a $2 reward, however, leads to an increase in the feedback rate that is economically and statistically significant.

We next turn to regression analysis. Motivated by our results regarding unconditional rebates, we include phase 1 in the data set we use for multivariate regression analysis. However, we control for phase fixed effects to allow for the possibility that there is an effect we were unable to measure in our earlier data tabulation.

Columns (1) and (2) of Table 6 show the results of a logit regression where the dependent variable is the dummy *Any feedback* (any feedback was given in the transaction). In this regression, we pool all of our 201 observations.[12] Marginal effects, evaluated at the independent variable mean, are reported (delta method standard errors are reported in parentheses). Of particular interest, the coefficients on feedback rebate (either $1 or $2) are positive but not statistically significant.

Together with the simple tabulations, the regression results provide fairly weak support for Hypothesis HE1, which predicts that feedback rewards lead to more feedback being given. We find that rebates do not make buyers more likely to leave feedback because we fail to reject the hypothesis that the coefficients associated with rebates are all equal to zero ($\beta_1 = \beta_4 = 0$ and $\beta_2 = \beta_5 = 0$).[13] The results suggest that the sign of the variation on rebates is consistent with Hypothesis HE1, but the coefficient is not statistically significant.

The regressions reported in Table 6 also show that slow shipping matters: we reject the hypothesis that the coefficients associated with slow shipping are all equal to zero ($\beta_3 = \beta_4 = \beta_5 = 0$).[14] For example, in column (2), the estimation of β_3 is -0.336, statistically significant at the 5% level. This means that, under the no-rebate treatment, holding all other variables at the mean, the probability of leaving feedback is 33.6% lower in slow shipping cases than in fast shipping cases.

Finally, although our theory makes no specific prediction regarding the effect of the seller's score (or the

[12] In §4.6, we return to the issues of alternative samples to consider for regression analysis.

[13] $p = 0.3695$ for the first test, $p = 0.2506$ for the second one.

[14] $p = 0.0396$.

dummy *Seller's perfect record*), the results in Table 6 (columns (1)–(4)) suggest that a seller with a better score is less likely to receive some feedback and more likely to receive negative feedback. We find this result somewhat puzzling. One referee suggests that this may be related to buyer expectations: buyers hold higher expectations of sellers with a better record and as a result are more likely to be disappointed and give negative feedback if they are not completely satisfied.

4.3. Can Buy Me Love

The next question of interest is whether the *nature* of buyer feedback is altered by the fact that a reward is offered by the seller. In other words, we now inquire whether feedback rebates work as "bribes," whereby reciprocity-minded buyers feel compelled to provide better feedback (more positive, less negative) than a reward-free experience would lead them to do. This corresponds to our Hypothesis HR1.

Within the set of fast shipment transactions, our results are inconclusive because all feedback was positive. This limitation of our first treatments led us to consider a second set of treatments. We purposely shipped our USB drive more slowly (exactly 14 days after receiving payment) and repeated the shift from no feedback to $1 and $2 conditional feedback (treatments S1 and S2, respectively).

Table 7 displays the results from this new treatment. As before, higher feedback rewards increase feedback frequency (from 60% to approximately 75%). The novel test that Table 7 allows is for whether paying for feedback alters the *nature* of feedback. We observe that as we shift from no rebate to a $1 or $2 rebate the percentage of positive feedback within transactions with feedback increases: from 58% to 75% to 91%, as the reward for feedback increases from $0 to $1 to $2. In terms of statistical significance (t-test), the mean value of propensity to leave feedback in the $1 treatment is not significantly higher than in the $0 treatment ($p = 0.3096$), whereas it is significantly higher in the $2 rebate treatment than in the $0 rebate treatment ($p = 0.0770$).

Next, we return to regression analysis to address the question of the nature of feedback. In the second set of regressions (columns (3) and (4) in Table 6), we run a linear probability model on whether feedback is negative.[15]

This regression allows us to detect biases in the nature of feedback. From the results in column (4) of Table 6, we find that, under the fast shipping treatment, neither a $1 nor a $2 rebate has an effect on the probability of negative feedback (in comparison to the no-rebate case); that is, neither β_1 nor β_2 are significantly different from 0. However, we find that, under the slow shipping treatment, both a $1 and a $2 rebate have a significant effect on the probability of receiving negative feedback (both are significant at the 1% level).[16]

To see whether rebates matter overall, we run joint tests and reject the hypothesis that the coefficients associated with rebates are equal to 0.[17] To compare the difference between the effects of the $1 and $2 rebates, we run joint tests and cannot reject the hypothesis that the difference is 0.[18]

To get an idea of the size of the feedback rebate coefficients, we note that, that under the slow shipping treatment, the probability of leaving negative feedback is 22.5 percentage points lower when a $1 rebate is offered and 39.7 percentage points when a $2 rebate is offered.[19] These results suggest that, consistent with Hypothesis HR1, the relative likelihood of negative feedback with respect to positive feedback is decreasing in rebate.

In sum, tabulation and regression results suggest that, even if you can buy feedback, you may not be able to buy unbiased feedback: reciprocity comes into play and what the seller ultimately buys is positive feedback, not honest feedback.

As previously mentioned, in one set of treatments we purposely lowered the value of q (by slowing down shipment). This allows us to test Hypothesis HR2, namely the hypothesis that the nature of feedback changes with q: a lower q leads to more negative and less positive feedback.

Table 7 Feedback Behavior in Slow Shipment Transactions

Feedback reward →	None		$1		$2	
Outcome ↓	#	%	#	%	#	%
No feedback given	8	40	4	25	4	27
Some feedback given	12	60	12	75	11	73
Total	20	100	16	100	15	100
Positive	7	58	9	75	10	91
Negative	2	17	2	17	0	0
Neutral	3	25	1	8	1	9

[15] Following previous evidence that neutral feedback is commonly interpreted as negative, we pool neutral and negative feedback messages into one single category, which we call "negative." See, for example, Cabral and Hortaçsu (2010). *Negative feedback* in columns (3) and (4) is defined as a dummy variable that equals 1 if feedback is negative or neutral and 0 if feedback is positive. Since fast shipping perfectly predicts nonnegative feedback, many observations are dropped when we run a logit model. For this reason, we use linear probability model instead of a logit model.

[16] For the $1 rebate case, $\beta_1 + \beta_4 = 0$, $p = 0.010$. For the $2 rebate case, $\beta_2 + \beta_5 = 0$, $p = 0.001$.

[17] For the $1 rebate case, $\beta_1 = \beta_4 = 0$, $p = 0.034$. For the $2 rebate case, $\beta_2 = \beta_5 = 0$, $p = 0.004$.

[18] Under the fast shipping treatment, $\beta_2 = \beta_1$, $p = 0.9253$. Under the slow shipping treatment, $\beta_2 + \beta_5 = \beta_1 + \beta_4$, $p = 0.2136$.

[19] That is, $-0.213 - 0.012 = -0.225$ and $-0.394 - 0.003 = -0.397$.

Table 7 provides the first piece of evidence consistent with Hypothesis HR2: whereas our fast shipping transactions received exclusively positive feedback, our slow shipping transactions received four negative feedback comments (the total number of negative and neutral feedback comments are nine). Using a t-test, we confirm that the negative feedback rate following a slow shipment transaction is statistically different from the negative feedback rate following a fast shipment transaction ($p = 0.0002$).[20, 21]

The results from tabulation are confirmed by regression analysis. From the results in columns (3) and (4) of Table 6, we find that slow shipping has a significant effect on the probability of getting negative feedback with respect to positive feedback.[22] From the results in column (4), we find that under the no-rebate treatment, the probability of getting negative feedback is 40.6% higher for slow shipping transactions than for fast shipping transactions. Under the \$1 rebate treatment, the probability of getting negative feedback is 19.3% ($0.406 - 0.213 = 0.193$) higher for slow shipping transactions than for fast shipping transactions.[23] Under the \$2 rebate treatment, the probability of getting negative feedback is not significantly higher for slow shipping transactions than for fast shipping transactions.[24] These results provide support for Hypothesis HR2: the relative likelihood of negative feedback with respect to positive feedback is higher when transaction quality is lower (at least for transactions other than those in the \$2 rebate treatment).

4.4. The Avengers

Do feedback payments change the timing of feedback? To the extent that buyers have a positive time discount, we should expect that feedback rebates lead economically minded buyers to leave feedback quicker; this is the thrust of Hypothesis HE2. Figure 2 presents summary data that addresses this possibility. We create the variable *Days to received feedback* (*DTRF*) by computing the difference between *Date feedback left* (*DFL*) and *Auction end date* (*AED*). In the case of delayed auctions, we subtract an extra 14 days to account for the delay in shipping with respect to fast shipment transactions. In other words, we create the variable so as to measure the lag between the moment the buyer receives the good and the moment the buyer provides feedback. To summarize,

$$DTRF = \begin{cases} DFL - AED & \text{if prompt shipping,} \\ DFL - AED - 14 & \text{if 14 day shipping,} \end{cases}$$

We split the data according to our six treatments (three levels of feedback reward and two levels of transaction quality). Regarding the level of feedback, Figure 2 suggests a very weak effect. This prediction is confirmed by regression analysis. The regressions in columns (5) and (6) in Table 6 show that the effect of feedback rebates on the speed of feedback is not statistically significant.[25] Overall, the empirical evidence provides little support for Hypothesis HE2.

A second prediction regarding speed of feedback relates to the change in transaction quality. As mentioned earlier, previous work has shown that there is such a thing as demand for justice (see §2). In the present context, this would lead us to expect that delayed transactions create a greater demand for feedback (especially negative feedback), and that this would be given quicker. This is the thrust of our Hypothesis HR3.

Figure 2 suggests that Hypothesis HR3 does indeed hold. In fact, the average number of days for leaving negative and neutral feedback is 4.7, whereas the average number of days for leaving positive feedback (in phases 2–4) is 11.2 days. When we include all phases, the average numbers of days are very similar: 4.7 and 10.5, respectively. These tabulation results are confirmed by regression analysis: column (5) in Table 6 shows that the coefficient on the dummy variable *Slow shipping* is -10.63; in other words, on average, slow shipment transactions lead buyers to give

Figure 2 (Color online) Average Number of Days Before Feedback Is Received

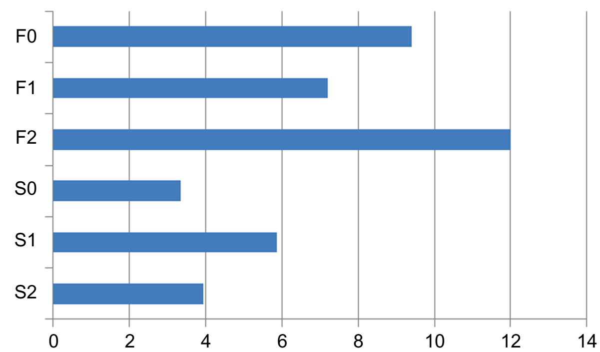

[20] The same is true if we consider both neutral and negative feedback as negative feedback ($p = 0.0000$).

[21] As an aside, we note that the lack of negative feedback for normal transactions should not be ascribed to "fear of retaliation," as several authors have previously argued. First, we (the seller) always give positive feedback promptly, so buyers have no reason to fear retaliation. Second, we do observe an increase in negative feedback as we reduce service quality.

[22] For instance, from the results in column (4), the test of $\beta_3 = \beta_4 = \beta_5 = 0$ yields $p = 0.0006$.

[23] For $\beta_3 + \beta_4 = 0$, $p = 0.073$.

[24] For $\beta_3 + \beta_5 = 0$, $p = 0.935$.

[25] We fail to reject the hypothesis that the coefficients associated with rebate are all equal to 0: for $\beta_1 = \beta_4 = 0$, $p = 0.448$; and for $\beta_2 = \beta_5 = 0$, $p = 0.950$.

Table 8 Feedback Reward and Bidding Behavior

Rebate ($)	0	1	2
Average price	3.44	3.51	3.15
Bidder count	3.65	3.82	3.53
Bid count	5.26	5.53	4.97
N	66	45	30

feedback with a delay (with respect to shipping date) that is approximately 10 days shorter than for fast shipment transactions.

4.5. Free Lunch

Different bidders have different costs of providing feedback. As suggested by the results above, even when no rebate is given, a considerable fraction of buyers do provide feedback. Therefore, when the seller pays any buyer who is willing to provide feedback, the marginal effect of such a payment is small in the sense that only a few buyers switch from not giving feedback to giving feedback. For this reason, it might seem that obtaining those marginal feedback comments comes at a very high cost (that is, a large amount of feedback rebates are given for a small number of extra feedback comments). However, an economically minded, forward-looking buyer with zero feedback cost—one who would give feedback regardless of the rebate—should anticipate a gain of $1 or $2 (as the case may be) and should reflect this anticipated gain in the bid. Even if the cost of giving feedback is not zero, to the extent that it is lower than the feedback reward, a rational buyer should expect a net gain from bidding in an auction with the promise of feedback reward. This is the thrust of Hypothesis HE3, the hypothesis that greater rebates for feedback should lead to higher bids (and a higher number of bidders).

Table 8 displays average price, bidder count, and number of bids for each treatment in phases 2–4.[26] The data suggest that, contrary to the (economic) theory prediction, there is not much difference in bidding behavior resulting from rebate promises.

This is confirmed by regression analysis, the results of which are reported in Table 9. Differently from the previous set of regressions, we do not include phase fixed effects since buyers do not know they are in treatment F or S when they make purchase decisions. In particular, the fact that a rebate was (mistakenly) given before feedback was received (in phase 1) should have no effect on bidding behavior. We do keep, however, seller fixed effects. The first regression looks at the determinants of sale price. Contrary to Hypothesis HE3 (and in accordance with the tabulation results) we see that feedback rebates have no significant effect on price. The regression coefficients have the right sign but are not statistically significant. We do observe, however, a significant negative coefficient on the variable seller score, similar to the first set of equations. In this case, the sign of the coefficient is particularly striking because standard reputation theories would predict it to be positive. However, the coefficient size is very small, a mere 6 cents of the dollar, or approximately 2% of the sale price.[27] Finally, the coefficient of the variable *Seller's perfect record*, which theory would predict to have a positive sign, is not significantly different from zero (and has a negative sign).

A similar pattern is observed in the second and third regressions (number of bidders and number of bids), where the only statistically significant coefficient is that of seller score—and with a negative sign.

Table 9 Regression Analysis of Bidding Behavior (All OLS Regressions)

Dependent variable	Price	Number of bidders	Number of bids
$1 *rebate*	0.050	−0.018	0.107
	(0.227)	(0.199)	(0.335)
$2 *rebate*	0.029	−0.059	−0.617
	(0.389)	(0.34)	(0.573)
Seller score	−0.062**	−0.065***	−0.125***
	(0.027)	(0.023)	(0.040)
Seller's perfect record	0.163	0.540	1.105
	(0.570)	(0.498)	(0.839)
Date	0.006	0.009*	0.019**
	(0.005)	(0.005)	(0.008)
Seller FE	Y	Y	Y
Constant	9.568	9.335	15.37
	(1.895)	(1.655)	(2.791)
N	201	201	201
Adj. R^2	0.147	0.1	0.116

Note. Seller FE, seller fixed effects.
** $p < 0.05$; *** $p < 0.01$.

[26] The results do not change if we include the data from phase 1 as well.

[27] Over time the demand for the 2 GB USB pen drives decreased. Moreover, negative feedback began to be observed as slow shipping transactions took place. As a result, one would expect price and number of bidders to decline. We did a robustness check by only using data from phases 1 and 2 (during which the demand was about the same and no slow shipping transactions took place). We found that the coefficient of seller score is positive and significant at the 10% level in the OLS regression on price and not significant in the other two regressions. We repeated the test in phases 3 and 4 (when slow shipping transactions were introduced and the 2 GB USB pen drives were no longer produced by Kingston). For this subsample we not find any significant effect of seller scores. All in all, we believe that the negative coefficient in the overall sample corresponds to bidders' observing negative feedback in phases 3 and 4, which in turn affected their bidding behavior in a way that is picked up by the seller score variable.

Since the seller score steadily increases over time, one might think that this variable is measuring something other than the buyers' estimate of the seller's value. However, we include the calendar date in all regressions, and this variable seems not to be significant. The negative effect of seller score on bids remains a puzzle to us (see also the discussion in §5).

Overall, the results show that the rebate variables have no significant effect on bidding behavior. In sum, we find no evidence for Hypothesis HE3. We find this result important because it shows that the cost of a feedback program, in terms of dollars per comment, may be considerably higher than what economic theory would predict.

In principle, there are several interpretations for the absence of an effect of feedback rewards on bidding behavior. One is that buyers are myopic in the sense that, at the time of bidding, they do not take into account the future savings provided by the feedback rebate. An alternative explanation is that buyers are incredulous about the feedback rebate offers. Still another alternative explanation is that buyers are simply unaware of the feedback rebate promise.

4.6. Robustness Analysis

We ran a series of additional regressions to test the robustness of our results. We repeated the regressions in Table 6 using data only from phases 2–4, and the general results still hold. The results suggest, nevertheless, that slow shipping is an important determinant of negative feedback, which further confirms Hypothesis HR2. Next, we ran linear probability models for the dependent variable negative or nonnegative feedback by considering (a) both positive feedback and no feedback as nonnegative feedback or (b) only no feedback as nonnegative feedback. We found that in both cases a \$2 rebate continues to have a significant effect on lowering the probability of getting negative feedback under the slow shipping treatment, and that the effect of a \$1 rebate is not significant. Further, we ran a linear probability model on the dependent variable leaving positive or nonpositive feedback by considering (a) both negative feedback and no feedback as nonnegative feedback, as well as (b) only no feedback as nonpositive feedback. We found that a \$1 rebate still has a significant effect on increasing the probability of getting positive feedback under the slow shipping treatment, and that a \$2 rebate has no significant effect.

5. Discussion and Concluding Remarks

Table 1 summarizes the hypotheses we set out to test. They are based on two alternative paradigms of buyer behavior, one based on economic rationality (*homo economicus*), one based on psychological behavior (*homo reciprocus*). Overall, the evidence favors the latter paradigm. We find very weak evidence and no evidence, respectively, for Hypotheses HE1 and HE2, the hypotheses that feedback rewards lead to more (HE1) and quicker (HE2) feedback. We also find no evidence for Hypothesis HE3, the hypothesis that feedback rewards affect bidding behavior (size of bids, number of bidders, or number of bids).

By contrast, we find fairly strong evidence in favor of the *homo reciprocus* paradigm. Specifically, we find that increasing feedback rewards changes the nature of feedback, with a greater share of positive feedback and a lower share of negative feedback being given (Hypothesis HR1), that lowering transaction quality (slower shipment) leads to a lower share of positive feedback and a greater share of negative feedback being given (Hypothesis HR2), and that lowering transaction quality speeds up the arrival of negative feedback (Hypothesis HR3).

All in all, our results suggest that it is possible to increase the feedback rate by giving conditional feedback rebates. However, this is a rather costly way of obtaining feedback. Moreover, it is likely that the nature of feedback will be considerably affected by rebates: you can buy feedback but you cannot buy unbiased feedback.

On March 2012, China's leading online trade platform (with 500 million registered users), Taobao, started a conditional feedback reward system. This development came to us as a surprise and was implemented well after we designed and ran our field experiment. Although it refers to a different trade platform (Taobao as opposed to eBay) and was designed in a slightly different way, the new Taobao scheme provides additional motivation and relevance for our work.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/mnsc.2014.2074.

Acknowledgments

The authors thank Matthew Beachy and Adam Zhanqing Hong for their excellent research assistance. The authors are very grateful to the following people for their comments on an earlier draft of this paper: an associate editor and three referees; Yuanyuan Chen, Ninghua Du, Feng Huang, Haichun Ye, Ke Zhang, Xiaolan Zhou, and Yahong Zhou; conference participants at the European Association for Research in Industrial Economics, Peking University's Industrial Organization and Management Strategy workshop, Shanghai Microeconomics workshop, and Xiamen University's Wang Yanan Institute for Studies in Economics; and seminar participants at the Southwestern University of Finance and Economics. Lingfang Li gratefully acknowledges the financial support of the National Natural Science Foundation of China [Grants 71402031 and 71490721], the University of Hong Kong–Fudan University IMBA Joint

Research Fund [JRF, App. No. JRF13_1201], and the Innovation Program of the Shanghai Municipal Education Commission [Grant 13ZS055].

References

Abeler J, Calaki J, Andree K, Basek C (2010) The power of apology. *Econom. Lett.* 107(2):233–235.

Ariely D, Bracha A, Meier S (2009) Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *Amer. Econom. Rev.* 99(1):544–555.

Avery C, Resnick P, Zeckhauser R (1999) The market for evaluations. *Amer. Econom. Rev.* 89(3):564–584.

Bénabou R, Tirole J (2003) Intrinsic and extrinsic motivation. *Rev. Econom. Stud.* 70(3):489–520.

Bolton GE, Katok E, Ockenfels A (2004) How effective are electronic reputation mechanisms? An experimental investigation. *Management Sci.* 50(11):1587–1602.

Brandts J, Charness G (2003) Truth or consequences: An experiment. *Management Sci.* 49(1):116–130.

Brown J, Morgan J (2006) Reputation in online markets: The market for trust. *Calif. Management Rev.* 49(1):61–81.

Brown J, Hossain T, Morgan J (2010) Shrouded attributes and information suppression: Evidence from field experiments. *Quart. J. Econom.* 125(2):859–876.

Cabral L (2012) Reputation on the Internet. Peitz M, Waldfogel J, eds. *The Oxford Handbook of the Digital Economy* (Oxford University Press, New York), 343–354.

Cabral L, Hortaçsu A (2010) The dynamics of seller reputation: Theory and evidence from eBay. *J. Indust. Econom.* 58(1):54–78.

Charness G, Levine D (2007) Intention and stochastic outcomes: An experimental study. *Econom. J.* 117(522):1051–1072.

Chen Y, Ho T-H, Kim Y-M (2010b) Knowledge market design: A field experiment on Google Answers. *J. Public Econom. Theory* 12(4):641–664.

Chen Y, Harper M, Konstan J, Li SX (2010a) Social comparisons and contributions to online communities: A field experiment on Movielens. *Amer. Econom. Rev.* 100(4):1358–1398.

Dellarocas C (2003) The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Sci.* 49(10):1407–1424.

Dellarocas C, Wood CA (2008) The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. *Management Sci.* 54(3):460–476.

de Quervain DJ-F, Fischbacher U, Treyer V, Schellhammer M, Schnyder U, Buck A, Fehr E (2004) The neural basis of altruistic punishment. *Science* 305(5688):1254–1258.

Falk A (2007) Gift exchange in the field. *Econometrica* 75(5): 1501–1511.

Fehr E, Gächter S (2000a) Cooperation and punishment in public goods experiments. *Amer. Econom. Rev.* 90(4):980–994.

Fehr E, Gächter S (2000b) Fairness and retaliation: The economics of reciprocity. *J. Econom. Perspect.* 14(3):159–181.

Fehr E, Kirchsteiger G, Riedl A (1993) Does fairness prevent market clearing? An experimental investigation. *Quart. J. Econom.* 108(2):437–459.

Gneezy U, List J (2006) Putting behavioral economics to work: Testing for gift exchange in labor markets using field experiments. *Econometrica* 74(5):1365–1384.

Grosskopf B, Sarin R (2010) Is reputation good or bad? An experiment. *Amer. Econom. Rev.* 100(5):2187–2204.

Gu B, Ye Q (2014) First step in social media: Measuring the influence of online management responses on customer satisfaction. *Production Oper. Management* 23(4):570–582.

Houser D, Wooders J (2006) Reputation in auctions: Theory and evidence from eBay. *J. Econom. Management Strategy* 15(2): 353–369.

Jin GZ, Kato A (2006) Price, quality and reputation: Evidence from an online field experiment. *RAND J. Econom.* 37(4):983–1005.

Jurca R, Faltings B (2007) Collusion resistant, incentive compatible feedback payments. *Proc. 8th ACM Conf. Electronic Commerce* (Association for Computing Machinery, New York), 200–209.

Li L (Ivy) (2010a) Reputation, trust, and rebates: How online auction markets can improve their feedback mechanisms. *J. Econom. Management Strategy* 19(2):303–331.

Li L (Ivy) (2010b) What is the cost of venting? Evidence from eBay. *Econom. Lett.* 108(2):215–218.

Li L (Ivy), Xiao E (2014) Money talks? An experimental study of rebate in reputation system design. *Management Sci.* 60(8): 2054–2072.

Miller N, Resnick P, Zeckhauser R (2005) Eliciting honest feedback: The peer prediction method. *Management Sci.* 51(9): 1359–1373.

Resnick P, Zeckhauser R, Swanson J, Lockwood L (2006) The value of reputation on eBay: A controlled experiment. *Experiment. Econom.* 9(2):79–101.

Shapiro C (1983) Premiums for high quality products as returns to reputations. *Quart. J. Econom.* 98(4):659–680.

Xiao E, Houser D (2005) Emotion expression in human punishment behavior. *Proc. Natl. Acad. Sci. USA* 102(20): 7398–7401.

Wang Z (2010) Anonymity, social image, and the competition for volunteers: A case study of the online market for reviews. *B.E. J. Econom. Anal. Policy* 10(1), Article 44.