The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment*

Chris Nosko†

Steven Tadelis‡

PRELIMINARY AND INCOMPLETE

July 10, 2014

Abstract

We argue that reputation mechanisms used by platform markets suffer from two inherent problems. First, buyers may draw conclusions about the quality of the platform from any single transaction, causing an externality across sellers. Second, buyers who abandon the platform without leaving feedback will cause seller reputations to be biased. Using data from ebay, we document both of these problems and argue that platforms will benefit from actively screening and promoting higher quality sellers. Exploiting the bias in feedback, we create a measure of seller quality and demonstrate the benefits of our approach through a controlled experiment that prioritizes better quality sellers to a random subset of buyers. We thus highlight the importance of reputational externalities in platform markets and chart an agenda that aims to create more realistic models of such markets.

*We grateful to many employees and executives at ebay without whom this research could not have been possible. We thank Andrei Hagiu and Glen Weyl for very helpful comments.

†University of Chicago Booth School of Business and ebay Research Labs. Email: cnosko@chicagobooth.edu

‡UC Berkeley Haas Schol of Business, ebay Research Labs, and NBER. Email: stadelis@berkeley.edu

1 Introduction

Decentralized marketplaces constitute some of the most fundamental building blocks of economic activity. Ebay, one of the first internet success stories, morphed from a used-goods auction site into one of the largest platform markets with over sixty billion dollars of merchandise traded in 2012. Online platform markets are rapidly growing with ebay's global expansion, Amazon's Marketplaces, and Taubau.com in China to name a few.

The anonymity of traders on platform markets raises concerns about market failures that result from asymmetric information. As a result, ecommerce marketplaces use some sort of decentralized "reputation" mechanisms. A concern, however, is that small decentralized sellers in platform markets do not internalize the impact of their actions on the marketplace as a whole. In particular, if buyers are uncertain about the distribution of seller-quality on a platform then one poor outcome may cause a buyer to update his beliefs about the quality of *all* sellers on the platform, resulting in a *reputational externality* across sellers. Furthermore, if some buyers leave the platform after a disappointing transaction without leaving feedback then the platform's reputation mechanism will be positively biased.

We study the challenges faced by market platforms in the presence of reputational externalities and biased feedback. We explore the limits of reputation mechanisms in the face of these problems, their impacts on the marketplace, and ways in which a platform designer can mitigate these adverse impacts. Thus, our paper offers three contributions to the literatures on the design of platform markets and on reputation mechanisms..

First, using data from ebay that records the *actual* behavior of buyers, we demonstrate that the externality we describe exists and that feedback is disproportionately absent for transactions that cause buyers to leave the platform. Buyers are not only unwilling to repeatedly transact with a seller who executed a bad transaction, but they are also less willing to return to the platform as a whole. Second, we propose a mechanism to mitigate the externality problem in which good-quality sellers are prioritized in search results. To better identify higher quality sellers, we exploit the absence of negative feedback to create a new measure of seller quality. This approach suggests that the design of online marketplaces may benefit from an "intermediate" ground between a *laissez faire* marketplace in which every seller is treated equally, and a heavily regulated marketplace that aggressively screens sellers. Last but not least, we conduct a field experiment where we change search result for a randomly chosen subset of buyers using our quality measure and we find that this approach increases the quality of transactions and the retention of buyers.

We begin our analysis by suggesting a simple conceptual framework of buyer behavior in online marketplaces. We proceed to construct a longitudinal dataset using ebay transactions that follow a cohort of new buyers who joined ebay in 2011 and include all their transactions through May 2014. The data include every transaction made by this cohort, including characteristics of the item, its price, and the item's seller. The goal is to proceed and measure how the quality of a transaction affects the future behavior of buyers on the platform.

We then show that the standard measure of a seller's quality, his reputation feedback, is highly skewed and omits valuable information. The "percent positive" (PP) measure for each seller is computed by dividing the number of transactions with positive feedback by the number of all transactions with any feedback for that seller. In our dataset, PP has a mean of 99.3% and a median of 100%, consistent with other studies that use ebay data (see, e.g., Dellarocas and Wood, 2005). This suggests that a central challenge in this paper is to construct a measure of seller quality that more accurately reflects a seller's true quality.[1] We construct a new quality measure that we call "effective percent positive" (EPP) that is computed by dividing the number of transactions with positive feedback by the total number of all transactions for that seller, including those with no feedback. In our dataset, EPP has a mean of 69% and a median of 73%, offering much more variability than PP.

We use our conceptual framework to empirically analyze the actual behavior of the cohort of buyers with respect to how current transactions affect their future behavior on ebay. In particular, we study the effect of a seller's EPP in a *current* transaction on the buyer's propensity to *continue* buying on ebay. As our framework suggests, a buyer who has a

[1]Beyond the observation that feedback is highly skewed, to measure the effect of quality on buyer behavior we cannot use measures of quality that are observable to the buyer because buyers will adjust their purchase behavior, which might affect seller actions. For instance, an observable lower quality seller might list a product for a lower price, and a buyer might knowingly take a risk on that item.

positive experience on ebay will be more likely to continue to transact on ebay again in the future and vice versa. Furthermore, Bayesian buyers should be less sensitive to their current transaction quality as they have more experiences on ebay. That is, the response to a negative experience early in a buyer's tenure on ebay should be more severe than later in his purchasing sequence. We confirm these hypotheses using the data and prove that EPP is indeed a better measure of seller quality than other available measures.

To establish the effect of improving buyer experiences by prioritizing higher quality sellers, and to alleviate any concerns about selection and endogeneity between buyers and a seller's EPP, we report results from a controlled experiment on ebay that implements our suggested approach by incorporating EPP into ebay's search-ranking algorithm. We select a random sample of ebay buyers who, when searching for goods on ebay, will be shown a list of products that promotes the EPP measure of the sellers compared to the control group in which this is not done. The results confirm the conclusions from the regression analysis described above and shows that treated buyers who have been exposed to higher EPP sellers are significantly more likely to return and purchase again on ebay compared to the control group of buyers.

Rather than display EPP instead of PP, we suggest that online marketplaces use measures like EPP in more opaque ways that improve a buyer's experience indirectly through the marketplace's search rank algorithm for two reasons. First, different buyers may interpret the same information in different ways. For one buyer a score of 88% might be satisfactory, while for another it is not, without having a clear understanding of how such a score translates into actual experiences. In theory, every rational expectations model of reputation has buyers being fully informed about the relationship between scores and outcomes, but in practice, and especially for less experienced buyers, such a mapping is unlikely to exist. Second, by using EPP instead of PP, sellers will most likely harass buyers who do not leave feedback in order to manipulate this new measure of seller quality. This would then cause a bias in EPP, and other measures of seller quality will need to be inferred from other parts of the data.

A growing empirical literature has focuses attention on ebay's reputation system, including Bajari and Hortacsu (2004), Bolton et al. (2013), Cabral and Hortacsu (2010), Dellarocas

(2003) and Klein et al. (2014) to name a few. Within this literature, the paper closest to ours is Dellarocas and Wood (2005) who reveal the problem of skewed feedback and propose an econometric method to uncover a seller's actual reputation.[2] There is a large theoretical literature on the role of reputation in facilitating trade between parties (see Bar-Isaac and Tadelis (2008) for a survey). Tirole (1996) proposes a model of collective reputations in which a group's reputation aggregates the reputation of it's members, yet his focus is on reputation persistence and the way in which the incentives of group members are influenced by the group's reputation, instead of the extrernality problem between group members.

We advocate for ways in which platform markets can design policies to augment their reputation mechanisms. The paper closest to ours in this regard is Hui et al. (2014) who study the role of ebay's "buyer protection" warranty on buyer welfare in conjunction with ebay's reputation and certification mechanisms. This differs from our approach of using the platform's search algorithm for controlling the exposure of seller quality to buyers. We conclude by discussing a more general agenda for studying reputation and quality in the design of platform markets.

2 Conceptual Framework

We wish to distinguish between two possible scenarios for a platform like ebay. The first is that buyers see the platform as a means of gaining access to sellers, but they neither consider characteristics of the platform itself, nor do they believe that sellers on the other side of the platform represent of the platform as a whole. In this case, there are no externalities across sellers. A buyer updates on the quality of the seller that he interacted with, and if the transaction goes badly, he may not deal with that seller again, but this does not affect the buyer's willingness to transact with other sellers on the platform.

In the second scenario the buyer uses outcomes of individual transactions to form beliefs about the whole platform. To consider this, imagine a dynamic Bayesian decision problem of

[2]Their method relies on the two-way feature of ebay's reputation mechanism. However, in 2008 ebay changed the mechanism so that buyers can no longer receive negative feedback, implying that their mechanism can no longer be used on ebay data.

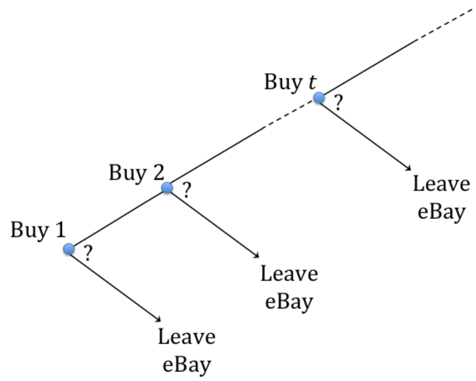

Figure 1: A Buyer's Dynamic Bayesian Decision Problem

a buyer who arrives at the ecommerce marketplace for the first time and is contemplating whether or not to purchase an item. His decision to purchase will rely on three basic elements: first, how much does he enjoy the site (user) experience; second, what are his expectations about the quality of the transaction; last, conditional on his belief, how price competitive is the site compared to other comparable marketplaces. If he decides to purchase, then after he receives the item he will update his beliefs about the quality of the site, and decide whether or not to purchase again, and so on, as depicted in Figure 1 .

Buyers can use a seller's feedback to form expectations about the quality of the seller, and by association, the marketplace overall. Every time the buyer makes a purchase, he collects an observation through which he updates his prior belief about the site's and the seller's expected quality. If the experiences were bad enough, he will update his belief about quality downward enough so as to decide to leave the site altogether. If, however, his experience was good, he will update his posterior in a positive way and continue to purchase from other sellers on the marketplace platform.

This framework of Bayesian updating also implies that the more transactions a buyer has made, the tighter will be his posterior, and this in turn implies that the effect of early experiences will be much more influential on the next purchase decision than an additional later experience. [3] It follows, therefore, that if a buyer experiences a relatively bad transaction earlier in the dynamic decision problem, then he is more likely to leave the marketplace than if he experiences the same transaction after several good experiences. This simple observation will form the basis for the central analysis on buyer behavior in Section 5.

3 Reputation and Transaction Quality at ebay

A large literature argues that reputation mechanisms mitigate inefficiencies in markets with asymmetric information. By publicly revealing ratings from past transactions, sellers are punished for delivering bad quality through the loss of future business from other market participants. Sellers therefore have an incentive to not defraud buyers and ensure that transactions go smoothly even if they are unlikely to interact with that specific buyer again. Mechanisms that operate with these principles in mind have been credited with sustaining markets such as long distance trade during the Middle Ages (Greif, 1989) and are often cited as reasons that online anonymous markets, such as ebay, AirBnB and others, were able to come into existence in the first place (Dellarocas, 2003).

A well-functioning reputation mechanism allows buyers to correctly infer the likelihood of a transaction going well without having past experience with any particular seller. In practice, however, the extent to which reputation mechanism reveal a seller's quality depends on two important assumptions. First, that the public information correctly mirrors the quality of past transactions, and second, that buyers correctly interpret reputation information. If either of these assumptions fail then buyers will inaccurately infer seller quality.

Ebay's reputation mechanism is often described as a resounding success for two reasons. First, ebay exists as a successful business despite the complete anonymity of the marketplace. The reputation mechanism seems to be all that is standing in the way of a collapse of trust.[4]

[3] This heuristic framework can easily be formalized using a standard dynamic model of a Bayesian decision maker that faces a distribution of quality with a well defined prior on the distribution of quality. Due to the well-understood nature of this dynamic problem, it would be redundant to offer the formal model.

[4] As Dellarocas (2003) puts it, "ebay's impressive commercial success seems to indicate that its feedback mechanism has succeeded in achieving its primary objective."

Second, many observable reputation characteristics correlate with our prior notions of the directional movement that these measures should induce. For instance, sellers with higher reputation scores and more transactions receive higher prices for their products. Similarly, reputation seems to matter more for higher priced goods than for lower priced goods.[5]

Rather than explore the returns from reputation on ebay, we focus attention on the extent to which the observable reputation measures are a good indication of transaction quality, and, if not, to what extent can a better measure of seller quality be created. Once a more accurate measure of quality is established, we will explore whether the second function of a reputation system – to give ex-ante incentives to sellers to ensure high quality – is operating in the current system (i.e., externalities across sellers are insignificant), or whether a platform should intervene and use levers other than the reputation system to increase platform quality.

When buyers complete a transaction, they are given the option of leaving either a positive, negative, or neutral feedback score, or leave no feedback. About 65% percent of buyers leave feedback on a transaction. Ebay uses this information to provide two observable seller reputation measures. The first, percent positive (PP), is defined as the seller's number of positive feedbacks divided by the sum of his number of positives, neutrals and negatives.[6] The second is feedback score, and is a summed value of the number of positive feedbacks minus the number of negative feedbacks. Both of these measures are displayed when a user views detailed item information. Figure 2 shows what this looks like for two different sellers. Seller A as a percent positive score of 96.9 and a seller feedback score of 317, while seller B has a percent positive of 99.5 and a seller feedback score of 44949.

Using the observable ebay reputation measures to examine the effect of seller quality on future outcomes is problematic for two reasons. First, buyers select the sellers they purchase from, leading to bias in any estimates of long term benefits from self-selecting into different quality sellers. For example is that price may adjust to reflect lower quality reputations.

[5]See Bajari and Hortacsu (2004) and Hortacsu and Cabral (2010) for more on these facts.

[6]To be precise, these numbers only look back at the last 12 months of a transaction for a seller and exclude repeat feedback from the same buyer for purchases done within the same calendar week.

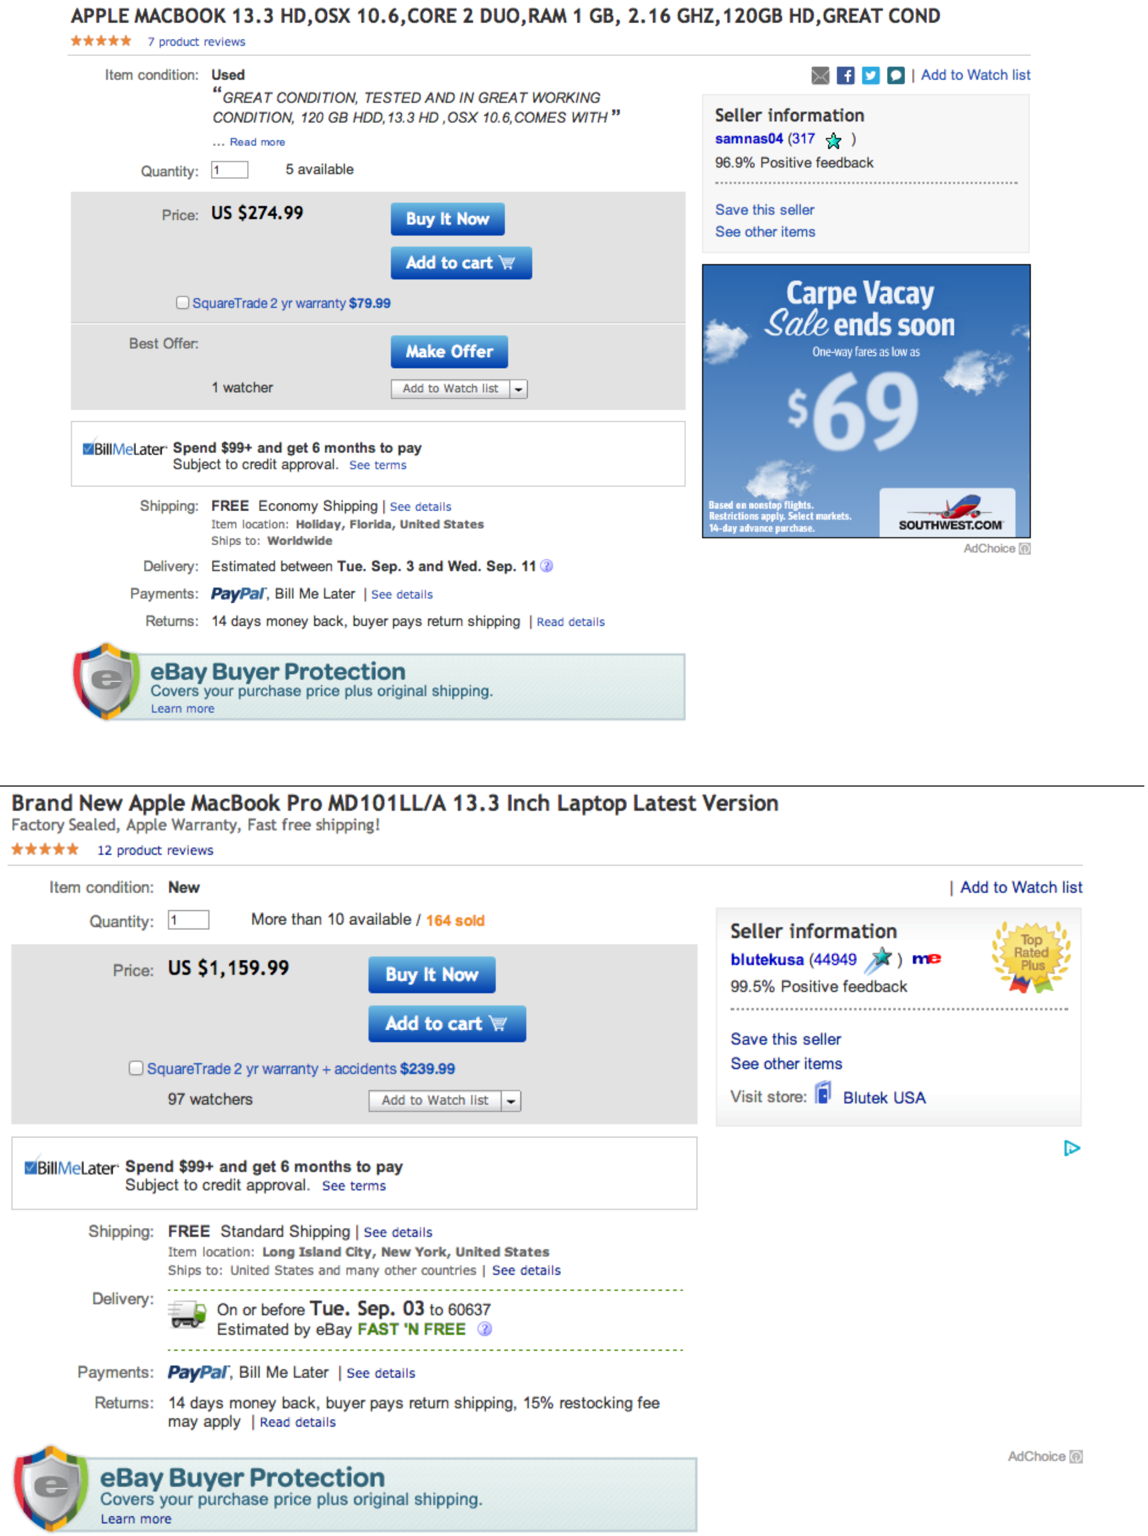

Figure 2: Seller reputation information as displayed to buyers

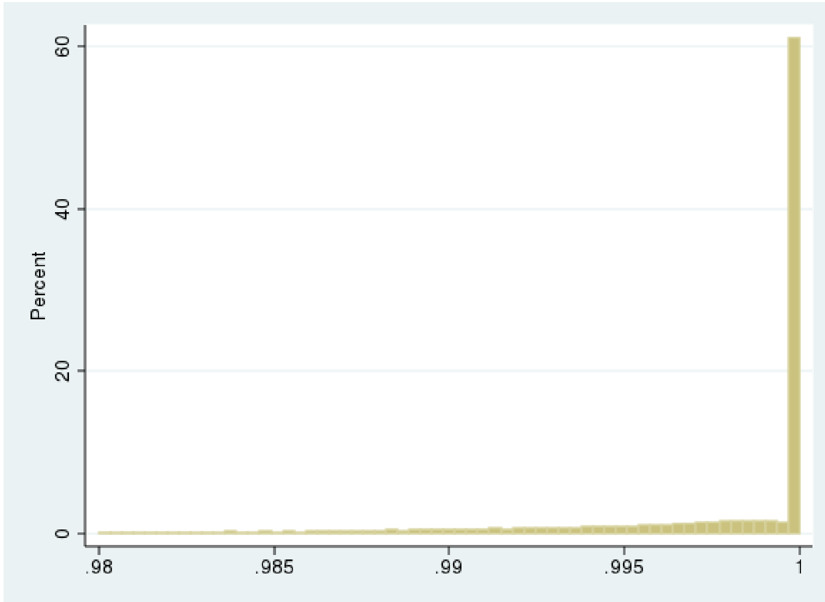

Figure 3: Percent Positive of Sellers

Second, the PP measure is highly skewed and hard to parse. Figure 3 displays the histogram of seller PP from a cross section of sellers in October 2011. The X-axis starts at 98%, which is the tenth percentile.[7] The median seller has a score of 100%. This could be indicative of a reputation system that works extremely well – bad sellers exit when their score falls low enough, leading to a high positive selection. Unfortunately, this is not the case. Out of over 44 million transactions completed in October of 2011 on ebay's U.S. marketplace, only 0.39% had negative feedback, while at the same time, over 1% had an actual dispute ticket opened within the ebay system (a step that takes substantially more effort on a buyer's part than leaving negative feedback). This indicates that there are a substantial number of transactions that went badly for which negative feedback was not left.[8]

Another problem is that many buyers may have trouble interpreting the numbers they are presented with. Naively, one may think that a score of 98% is excellent (in some sort of

[7]Our main dataset contains a panel of buyers and tracks their behavior over time. Here, we display a cross sectional view from a separate dataset of all transactions that occurred on the U.S. site in October of 2011. There are some newer sellers with a percent positive of 0. Leaving them in would skew the graph such that the detail at the top end could not be seen.

[8]This in fact proves the central tenet in Dellarocas and Wood (2005) who claim that silence in the feedback system includes many transactions for which buyers had bad experiences but chose not to report them.

absolute scale). In reality, a score of 98% places a seller below the tenth percentile of the distribution. This effect may be especially pronounced for new users who may not have seen enough sellers to be able to judge the scale accurately, a point to which we return below.

Arguably, buyers do not leave negative feedbacks because it is not anonymous and sellers historically reacted by reciprocating.[9] Anecdotal evidence shows that sellers sometimes react badly to negative feedback, often harassing buyers in an attempt to get them to change it.[10] In part because of this, a social norm has developed around not leaving negative feedback.

We proceed to construct a measure of *unobservable* seller quality based on the idea that buyers who experience a bad transaction are less likely to leave negative feedback, and are silent instead, while buyers who experience a good transaction are more likely to leave positive feedback.[11] To operationalize this, we measure the propensity of positive feedback for any given seller. Controlling for observable feedback measures (PP and feedback score), we conjecture that a seller with a lower propensity of positive feedback will be more likely to deliver a worse experience.

Walking through a specific example, consider two sellers, A who had 120 transactions and and B who had 150. Assume that both received one negative feedback, and 99 positive feedbacks. Using ebay's PP measure, both have a PP of 99% ($\frac{99}{99+1}$). However, seller A had only 20 silent transactions with no feedback while seller B had 50 silent transactions. We define "effective" PP (EPP) as the number of positive feedback divided by total transactions, in which case seller A has an EPP of 82.5% while seller B has an EPP of only 66% and is a

[9] Up until 2008, both parties could leave negative feedback, and after that sellers can only leave positive feedback or no feedback. There is a long history of reciprocal feedback behavior before the 2008 change as documented by Bolton et al. (2013).

[10] In one case, a seller called the buyer and threatened him after his negative feedback. ("ebay Shopper Says He Was Harassed By Seller," http://www.thedenverchannel.com/lifestyle/technology/ebay-shopper-says-he-was-harassed-by-seller). In another case, a buyer was sued for leaving negative feedback ("ebay buyer sued for defamation after leaving negative feedback on auction site," http://www.dailymail.co.uk/news/article-1265490/ebay-buyer-sued-defamation-leaving-negative-feedback-auction-site.html.)

[11] For example, consider a set-up in which there is a distribution of "public mindedness" among individuals that compels them to enjoy leaving feedback for the benefit of future buyers. If the costs and benefits of leaving feedback would not depend on the quality of the transaction, then the feedback left should be unbiased. However, if the cost of leaving truthful feedback is higher for bad transactions due to the harassment costs, then such a skew in feedback will result.

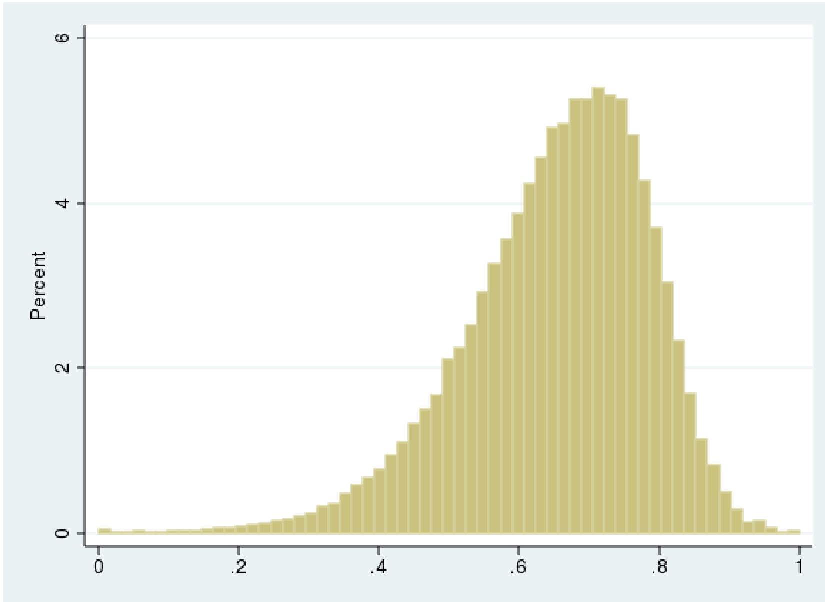

Figure 4: Histogram of Sellers' Effective Percent Positive Scores

worse seller on average. Importantly, ebay does not display the total number of transactions a seller has completed and buyers cannot therefore back-out a seller's EPP score.

Figure 4 displays a histogram of EPP scores for the cross section of sellers mentioned above. The mean of this distribution is .64 and the median is .67. Importantly, unlike percent positive, there is substantial spread in the distribution.

To verify that EPP contains information about buyers' experiences, we define a bad buyer experience (BBE) as one in which the buyer either left negative feedback, opened a dispute with ebay, or left low stars on the detailed seller ratings. Out of the 15,384,439 transactions in our full dataset, 521,473 (3.39 percent) of them resulted in BBEs. Despite the belief that buyers are reluctant to report bad experiences either because of seller wrath or because of the time cost of going through the ebay dispute process, table 1 explores whether their is a relationship between BBEs and seller quality as measured both by observables and by EPP. We run a probit regression of BBE on seller quality scores and controls for price, category, and purchase type (auction or fixed price). All of the coefficients are highly statistically significant. The coefficient of interest on EPP is negative and highly significant, indicating that transacting with higher EPP sellers decreases the probability that a BBE will occur, consistent with EPP being a measure of seller quality.[12]

Table 1: Probit of Bad Buyer Experiences and EPP

BBEG Flag (0/1)	
Seller Feedback Score	-6.55e-08***
	5.21e-09
Percent Positive Dummy	
excluded: $0 < .994$	
$\geq .994 < 1$	-0.146***
	0.00284
$= 1$	-0.182***
	0.00336
EPP	-0.695***
	0.0134
Item Price	0.00113***
	0.0000260
Seller Standards Dummy	
excluded: Below Standard	
Standard	0.0448***
	0.00588
Above Standard	-0.129***
	0.00551
ETRS	-0.278***
	0.00580
Constant	-1.278***
	0.0135
N	12,814,847

Regression includes controls for (coefficients not displayed) auction type (auction, fixed price), item category, item condition (new, used, refurbished), and buyer transaction number

Even though EPP is unobservable to buyers, perhaps buyers observe signals that are correlated with EPP, questioning its exogeneity. To explore this we consider whether more experienced buyers differentially select into transacting with higher EPP sellers. Table 2 shows the results of an OLS regression where EPP is the left hand side variable. The variable Buyer Transaction Number is the number of transactions that a buyer has completed on ebay.

[12]Many of the regressions do not use the full dataset because certain variables (such as new vs. used or product category) are missing in the ebay database for some observations. Instead of limiting the variables we use, we exclude these observations from the regressions.

If more experienced buyers were transacting with higher EPP sellers, we would be worried about the selection story. That does not appear to be the case. Although the variable is statistically significant (with the wrong sign), its magnitude is negligible.

Table 2: Selection into EPP

	b/se
Buyer Transaction Number	-0.00000405***
	4.28e-08
Seller Feedback Score	0.000000116***
	3.74e-10
Percent Positive Dummy	
excluded: 0 < .994	
$\geq .994 < 1$	0.0711***
	0.0000619
= 1	0.0842***
	0.0000949
Seller Standards Dummy	
excluded: Below Standard	
Standard	-0.0184***
	0.000164
Above Standard	-0.00996***
	0.000144
ETRS	-0.0126***
	0.000145
Item Price	-0.0000959***
	0.000000626
Constant	0.634***
	0.000182
N	12,814,870

Regression includes controls for (coefficients not displayed) auction type (auction, fixed price), item category, item condition (new, used, refurbished), and the number of transactions a seller has completed up to the focal observation

4 Data

We collected data for a cohort of new users to the U.S. site of ebay.com who joined the platform anytime in 2011. We select users who signed up for a new account and purchased

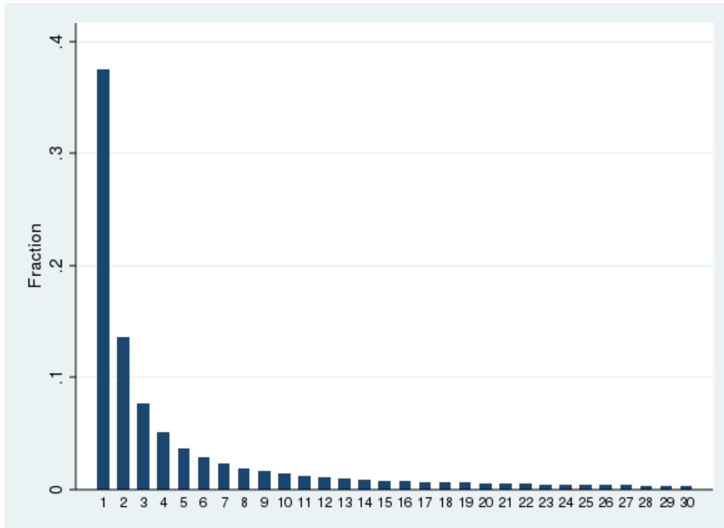

Figure 5: Histogram of Total Transactions by Buyer

an item within 30 days of sign-up.[13] The 2011 cohort includes over nine million new users, making their whole transaction history too large for meaningful analysis. As a result, we take a 10% random sample and analyze the behavior of 935,326 buyers.

For each buyer we track all transactions starting with their initial sign up until May 31, 2014. That results in 15,384,439 observations – on average 16 transactions per user. Each observation contains rich information about the transaction, including but not limited to price, item category, title, the seller, whether it was an auction or fixed price, and quantity purchased. There were a total of 1,854,813 sellers associated with all of the buyer transactions. We also collect information on the seller that each buyer transacted with.

Figure 5 is a histogram of the total number of transactions by an individual buyer over the course of his tenure, truncated at 30 transactions. A large percentage of ebay buyers make very few purchases over their life-cycle – a full 38% of new buyers purchase only once, with an additional 14% purchasing twice and moving on. On the other end of the spectrum, there is an extremely large right tail. While the median number of transactions is 2, the mean is 16, the 95th percentile is 65, and the max is 19,359.

[13]Replicating the analysis for the 2008, 2009, and 2010 cohorts yields very similar results.

Basic seller information such as the feedback score, percent positive, and number of transactions the seller had in the past, were collected. We construct a seller EPP score at each separate transaction by looking back at all of the seller's transactions (capped in January of 2005, the earliest data stored) up to the point right before the transaction, and then dividing the number of positive feedbacks by the total number of transactions for that seller. This generates a complete snapshot of the information structure at the point when the buyer was making his decision and, as such, we do not include the focal transaction in the measure. Recall that the buyers cannot observe or infer the EPP measure.

5 Reputational Externalities and Buyer Behavior

We are first interested to separate between two scenarios. In the first, buyers use ebay merely to connect with certain sellers, and no externalities across sellers are present. In the second, buyers consider ebay as a provider of quality services, in which case they infer the quality of the platform from individual transaction and an externality across sellers exists.

5.1 Buyer Behavior

Table 3 is a cross tabulation by buyer of how the total number of transactions relates to the total number of sellers that a buyer interacted with. Out of the 38,149 buyers who had a total of 20-29 transactions during our sample period, 23,367 of them bought from between 20 and 29 different sellers while 116 of them bought all their transactions from a single seller.[14] This shows that buyers tend to deal with large numbers of sellers and therefore suggests that externalities may indeed exist. This is further discussed in the next subsection.

[14] The total number of buyers adds up to 872,348 instead of 935,326 because we have truncated at buyers with 49 total transactions or less for compactness.

Table 3: Total Transactions by Total Number of Sellers for Each Buyer

Total Transactions	Total Number of Sellers 00-01	02-05	06-09	10-19	20-29	30-49	Total
00-01	350,881	0	0	0	0	0	350,881
02-05	27,603	253,032	0	0	0	0	280,635
06-09	1,206	19,374	60,590	0	0	0	81,170
10-19	492	2,802	15,959	64,112	0	0	83,365
20-29	116	386	767	13,513	23,367	0	38,149
30-49	67	207	273	1,810	11,685	24,106	38,148
Total	380,365	275,801	77,589	79,435	35,052	24,106	872,348

5.2 Externalities across Sellers

In order to explore the extent of seller externalities, we ask how transaction quality affects the probability that a buyer will transact on the platform as a whole again, and compare this to the probability that he transacts with the same seller again. If, as we conjectured earlier in Section 2, buyers learn about the platform from the quality of any given transaction, then a higher quality transaction will cause more purchases overall for the platform and create a disconnect between the incentives for an individual seller and the platform as a whole.

Our econometric specifications include probit regression of the following form,

$$y_{ijt} = \alpha_0 + \alpha_1 EPP_{jt} + \beta \cdot \bar{b}_{it} + \gamma \cdot \bar{s}_{jt} + \delta \cdot \bar{d}_t + \varepsilon_{ijt}, \tag{1}$$

where y_{ijt} is an indicator equal to 1 if buyer i bought again on ebay after purchasing transaction t from seller j, EPP_{jt} is the EPP of seller j at transaction t, $\bar{b}_{it}$ is a vector of buyer characteristics (e.g., how many transactions they completed), $\bar{s}_{jt}$ is a vector of seller characteristics (e.g., reputation score and PP), and $\bar{d}_t$ is a vector of transaction characteristics.

Because buyers do not observe EPP and do not act on it (recall Table 2), it can be considered as an exogenous shock to the quality of the seller. The higher the EPP, the more likely it is that the transaction goes well and therefore the more likely the buyer is to return and purchase on ebay – consistent with our conceptual framework outlined in Section 2.

Table 4: Baseline EPP Regressions

	60 Day Return	180 Day Return	Any Return
Seller Feedback Score	-9.12e-08***	-2.07e-08	7.20e-08***
	1.18e-08	1.29e-08	1.44e-08
Percent Positive Dummy			
excluded: 0 < .994			
≥ .994 < 1	-0.0507***	-0.0553***	-0.0524***
	0.00144	0.00160	0.00169
= 1	-0.0804***	-0.0888***	-0.0925***
	0.00218	0.00244	0.00257
EPP	1.014***	1.083***	1.231***
	0.00667	0.00708	0.00743
Item Price	-0.00175***	-0.00165***	-0.00155***
	0.0000144	0.0000139	0.0000140
Seller Standards Dummy			
excluded: Below Standard			
Standard	-0.0895***	-0.0838***	-0.109***
	0.00367	0.00412	0.00440
Above Standard	-0.0669***	-0.0736***	-0.0941***
	0.00321	0.00361	0.00388
ETRS	-0.0911***	-0.0921***	-0.103***
	0.00326	0.00363	0.00390
Seller Number of Trans	2.20e-08**	-1.43e-08	-6.59e-08***
	6.85e-09	7.58e-09	8.49e-09
Used / New Dummy			
excluded: New			
Refurbished	0.0126**	-0.000736	-0.0210***
	0.00425	0.00461	0.00485
Used	-0.00164	-0.0290***	-0.0608***
	0.00207	0.00218	0.00228
Constant	-0.506***	-0.404***	-0.226***
	0.00630	0.00673	0.00706
N	11,883,443	11,883,443	11,883,443

Regression includes controls for (coefficients not displayed) auction type (auction, fixed price), item category, and the transaction number of the buyer at the time of the focal observation.

Table 4 is our baseline regression table. In column 1, $y_{ijt} = 1$ if the buyer returns within 60 days of a given purchase, column 2 is within 180 days, and column 3 is if they had any repeat purchase in the sample period. Standard errors are clustered at the individual level and we control for transaction type and category without reporting the coefficients.

The coefficient estimates are stable across the probit specifications and most results confirm in the expected relationship. Buyers who end up purchasing more (total transactions) are more likely to return after any given purchase.[15] Our measure of seller quality, EPP, is highly statistically and economically significant (magnitudes will be discussed below), indicating that when buyers end up with a higher quality seller, they stick around and purchase more both in terms of frequency and revenue.

Interestingly, the observable seller reputation measures are either unstable (Seller Feedback Score) or negative (Percent Positive). We interpret this as selection: Once we control for seller quality (by including EPP), buyers self select into different types of sellers based on their ability to interpret quality or on their initial intentions. Perhaps, someone who knows ebay and is planning on sticking around may be more willing to take a risk with a lower Percent Positive seller, whereas someone who is weary about ebay and does not plan to return may only choose to buy from 100% PP sellers.

Table 5 repeats the regression but now y_{ijt} is an indicator equal to 1 if buyer i bought again *from seller* j after purchasing transaction t from seller j. As the table clearly indicates, the coefficients are a lot smaller than they are in Table 4. This clearly indicates that the experience a buyer has on ebay is a lot more likely to influence whether he returns to the site, rather than him returning to the same seller, which is a very unlikely event.

5.3 Buyer Learning About the Platform

While there are many reasons that a buyer may make a purchase on ebay and not return in the future, the above regressions help us distinguish between two important ones. The first is selection: People come to ebay looking for a specific item, purchase that item and then have

[15]We include total transactions as a regressor to control for the selection of who ends up staying on ebay for long periods of time because of their taste for what ebay offers.

Table 5: Likelihood of Returning to the Same Seller

	60 Day Return	180 Day Return	Any Return
Seller Feedback Score	-0.00000125***	-0.00000122***	-0.00000114***
	6.24e-08	5.89e-08	5.46e-08
Percent Positive Dummy			
excluded: 0 < .994			
$\geq .994 < 1$	0.0921***	0.0951***	0.0992***
	0.00506	0.00479	0.00458
= 1	-0.144***	-0.152***	-0.160***
	0.00687	0.00662	0.00644
EPP	0.221***	0.263***	0.343***
	0.0307	0.0295	0.0285
Item Price	-0.00183***	-0.00165***	-0.00153***
	0.0000889	0.0000807	0.0000750
Seller Standards Dummy			
excluded: Below Standard			
Standard	-0.378***	-0.387***	-0.391***
	0.00856	0.00832	0.00813
Above Standard	-0.0294***	-0.0220***	-0.0128*
	0.00680	0.00656	0.00639
ETRS	-0.0296***	-0.0170*	0.000411
	0.00739	0.00717	0.00700
Seller Number of Trans	0.000000798***	0.000000808***	0.000000787***
	3.17e-08	3.03e-08	2.84e-08
Used / New Dummy			
excluded: New			
Refurbished	-0.0258	-0.0362	-0.0456*
	0.0265	0.0242	0.0227
Used	-0.292***	-0.301***	-0.307***
	0.0133	0.0131	0.0127
Constant	-1.196***	-1.224***	-1.264***
	0.0249	0.0240	0.0234
N	11,883,452	11,883,452	11,883.452

no need to return in the future. The second, which is new to this paper, is that buyers come to ebay with limited knowledge about the platform and update beliefs over time. Any seller's quality influences their decision to come back to the platform as a whole.

The dynamic updating framework described in Section 2 implies that transaction quality should matter less for more experienced buyers relative to newer ones. Every experience a buyer has on ebay will help him learn about his idiosyncratic match value with the site as get a draw from the seller distribution on the site. As such, a buyer with more experience is more likely, on average, to return to the site *both* because of selection (people who had low match values have left already) and because a bad draw later will have less of an influence on his beliefs about quality because of a more refined posterior belief.

Table 6: Cross Tab: No. of Transaction with 180-day Return Probability

	No Return	Return	Total
01-05	912,616	1,788,930	2,701,546
	33.78%	66.22%	100.00
06-09	135,472	938,428	1,073,900
	12.61%	87.39%	100.00
10-19	132,613	1,650,773	1,783,386
	7.44%	92.56%	100.00
20-29	55,264	1,146,697	1,201,961
	4.60%	95.40%	100.00
30-49	50,728	1,557,494	1,608,222
	3.15%	96.85%	100.00
50-99	41,639	2,082,843	2,124,482
	1.96%	98.04%	100.00
100+	32,002	4,858,940	4,890,942
	0.65%	99.35%	100.00
Total	1,360,334	14,024,105	15,384,439
	8.84	91.16	100.00

Table 6 cross tabulates the buyer's experience measured in the number of transaction against whether or not a buyer returns to purchase on ebay within 180 days. As a buyer becomes more experienced, he is much more likely to return to ebay and purchase, consistent with either selection or learning about the idiosyncratic match with the platform.

Next consider the effect that the quality of a transaction has as a buyer progresses through his tenure on ebay. Table 7 extends our baseline regression in Table 4 above to include the buyer's experience. As the results on the interaction effects between experience (Transaction Cat) and EPP show, more experienced buyers are less effected by the seller's EPP, implying that they are not responding as much to new information.[16] This is consistent with the dynamic Bayes framework that we outline is Section 2.

Table 7: Baseline Regressions including Buyer Transaction Number

	60 Day Return	180 Day Return	Any Return	
	b/se	b/se	b/se	b/se
Seller Feedback Score	-6.82e-08***	-5.26e-08***	-6.82e-08***	
	2.67e-09	3.13e-09	2.67e-09	
Percent Positive	-1.554***	-1.414***	-1.554***	
	0.0210	0.0226	0.0210	
EPP	0.909***	0.862***	0.909***	
	0.00798	0.00792	0.00798	
Item Price	-0.162***	-0.153***	-0.162***	
	0.00120	0.00113	0.00120	
Trans Cat 06-09*EPP	0.379***	0.372***	0.379***	
	0.0132	0.0141	0.0132	
Trans Cat 10-19*EPP	0.387***	0.352***	0.387***	
	0.0121	0.0130	0.0121	
Trans Cat 20-29*EPP	0.354***	0.338***	0.354***	
	0.0146	0.0166	0.0146	
Trans Cat 30-49*EPP	0.272***	0.320***	0.272***	
	0.0144	0.0164	0.0144	
Trans Cat 50-99*EPP	0.131***	0.294***	0.131***	
	0.0148	0.0166	0.0148	
Trans Cat 100-199*EPP	-0.00192	0.343***	-0.00192	
	0.0191	0.0206	0.0191	
Trans Cat 200-999*EPP	-0.137***	0.267***	-0.137***	
	0.0272	0.0267	0.0272	
Trans Cat 1000+*EPP	-0.274**	0.0669	-0.274**	
	0.0964	0.0873	0.0964	
cons	1.309***	1.387***	1.309***	
	0.0199	0.0214	0.0199	
N	15654527	15654527	15654527	

[16]The regression includes the uninteracted experience measures, which are not reported, though they are consistent with Table 6.

6 Using Search to Internalize Seller Quality

Platforms can use several mechanisms to match buyers with better quality sellers. At a draconian extreme, the platform can expel any seller from the platform once it learns that a transaction was less than perfect. This would almost surely result in a selected pool of high quality sellers. In would also strongly incentivize sellers to ensure that transactions went well according to the metric used to determine being kicked off of the site. The downside would be a number of sellers mistakenly getting booted off the platform, along with a reduction in the overall site inventory and size.

Instead, ebay has traditionally adopted a laissez-faire approach to managing sellers on its marketplace, and its ethos rests on buyers and sellers being good "citizens." Along with the reputation system, the approach and policies of ebay can best be described as Caveat Emptor. This has begun to change recently, with ebay taking a more active role in managing the marketplace.[17] These measures include actively seeking to weed out bad quality sellers (to a much smaller extent than the extreme described above) and creating a "buyer protection" program that allows buyers to voice complaints to the platform directly about a transaction for potential reimbursement, rather than having to go through the individual sellers.[18]

We suggest an intermediate approach and prescribe online platforms to leverage their search technology to control buyer experience. At one extreme, a seller can effectively be removed if he never showed up in search results because no buyers would find him. On the other extreme, seller quality can only be considered to a very limited extent in prioritizing search results. We report results from an experiment where we incorporate our measure of seller quality, EPP, into ebay's search algorithm so that higher EPP sellers get displayed

[17] See, for example, "EBay to Get More Involved in Transaction Disputes", http://www.pcworld.com/article/163099/article.html

[18] Every platform implicitly or explicitly takes a stand on how actively they manage their marketplace. In contrast to ebay, Amazon has always extensively pruned its seller pool on Amazon Marketplace, making them jump through hoops to join, holding transaction receipts in escrow, and kicking sellers off the site quite quickly. Similarly, Stubhub (now an ebay company) has been much more careful to control the buyer experience. Buyers purchasing on Stubhub are not even aware of which seller they are purchasing from and all disputes are handled with Stubhub directly. These policies completely negate the need for a reputation system and essentially mean there are no externalities across sellers (all of this is internalized by the platform).

higher on the search results page (holding everything else constant). This form of intervention provides a large amount of variation in the extent to which sellers are prioritized.

For this strategy to work, buyers must face some search cost and be more likely to select an item higher up on the search results page. The literature on search costs has demonstrated correlation between ranking and purchase (or click-through) behavior (Ghose et al., 2013). To the best of our knowledge, the experimental evidence we present is some of the first to show how buyers respond to truly exogenous shifts in search rankings.

Promoting seller quality as we propose may come at the expense of providing more relevant items. Thus, one component of the experiment is to estimate the trade off inherent in manipulating search results to prioritize better quality sellers. On one hand, there could be a longterm benefit from buyers interacting with better quality sellers and returning to the site more often in the future. On the other hand, buyers may be less likely to purchase because they have a harder time finding what they want.

The experiment we conduct allows us to do three things: (1) Answer any lingering doubts about the exogeneity of EPP as an unobserved measure of seller quality; (2) Explore the extent to which consumers respond to search ranking schemes, and hence how effective changes in them might be for platforms wishing to internalize seller quality externalities; and (3) Quantify the downside of using search rankings – the extent to which consumers do not purchase because they are unable to find the product they are looking for.

From December 14th, 2011 though January 2, 2012, 10% of ebay's U.S. site traffic—several million searches per day—was placed into our experimental treatment and exposed to a ranking scheme that differed from the default site algorithm. Because of other site considerations, we had limited control over the weighting that the EPP measure received, a point to which we will return below.

The ideal experiment would treat active users, ensuring that an individual user was always in the treatment group whenever he searched for a product on ebay. However, it is impossible to unambiguously link a site visit to a specific user either because the user visits the site from a computer or browser he has never signed in from before, or because he has deleted his cookies (the way sites keep track of users between visits). This creates the potential for leakage between the treatment and control groups, where a user is sometimes exposed to search results in the treatment group and sometimes in the control (normal) group.

In order to understand the magnitude of this problem, it is necessary to understand how site visits are linked to users. Ebay stores a Globally Unique Identifier (GUID) in a cookie on browsers that visit the site, allowing ebay to track whether the same browser visits the site in the future. Through a complex algorithm, ebay attempts to match GUIDs to user IDs (UIDs) by tracking whether that browser was used to sign into an ebay account at any time (thereby matching GUID to UID). Multiple GUIDs may be linked to the same UID if a user signs in from multiple browsers on the same computer or from multiple computers. Our experiment was run at the GUID level, meaning 10% of active GUIDs were placed into the treatment group. This means that a user could be placed into the treatment group for one, but perhaps not all, of the GUIDs linked to a their user account. Fortunately, we can track this behavior and observe the number of searches that a user made within the treatment and control groups, allowing for a partial correction to the problem.

We collected data on all searches done during the experimental period (including the search query and the items that were displayed to the user), whether or not the GUID was in the treatment or control group, any other site behavior that the user did (clicks on products, etc.), and purchases both during the experimental period and after the experimental period. Tracking purchases beyond the end of the experiment is critical because this is our measure of buyer satisfaction. The test is whether a buyer is more likely to come back and purchases again if randomly assigned to the treatment group, conditional on purchasing in the experimental period.

Because an individual search returns several items (typically 50) on a single page, and each of those items is associated with a seller (and hence, an EPP score), it is instructive to collapse a whole search into one single measure of the quality of sellers in the search. We use a weighted average we call "discounted search EPP", which is done by weighting each item by its position in the search results. Specifically, for any given search we weight the EPP score of each item displayed to the user by the inverse of the item's position on the search results page. This reflects the prevailing belief that items ranked higher up on the page are more visible, and hence play a larger role in the user's decision process.

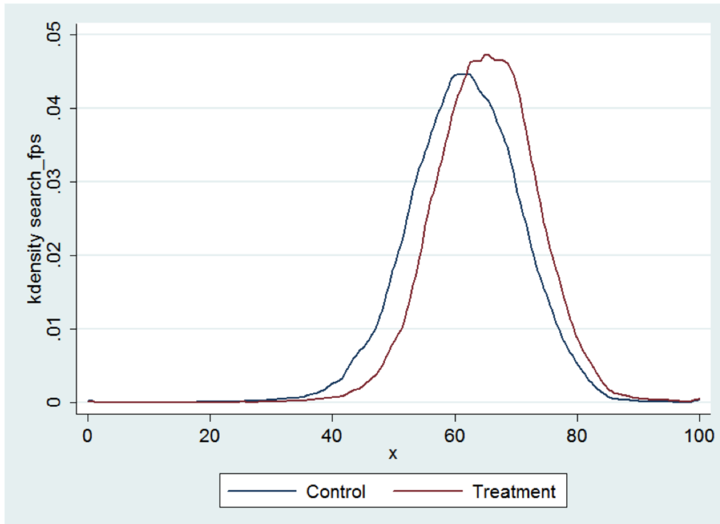

Figure 6: Discounted Search EPP Scores between Treatment and Control Groups

Figure 6 shows a kernel density plot of the Discounted Search EPP scores for all searches in the treatment and control groups. The mean in the control group is 59% and of the treatment group is 64%. The distributions are statistically different from each other, meaning that the experiment did in fact change rankings.

Table 8 shows the results of a t-test of means between the treatment and control group. Of the 14,872,320 purchases that occurred in the control group during the experiment, 78% of the time the user returned and purchased again within three months following the experiment. In the treatment group there were 1,478,987 purchases and users returned 80% of the time. The difference of 1.8 percentage points is highly statistically significant and economically meaningful. By simply re-ordering search results to prioritize better quality sellers, we were able to improve the probability of return.

Table 9 leverages the experimental variation in regression form. We control for everything as in the transactional regressions above (Tables 4 and 7) but use the treatment dummy

Table 8: t-test for difference between treatment and control group

Group	Obs	Mean	Std. Err.	95% Conf. Interval	
Control	14,872,320	.7817077	.0001071	.7814978	.7819176
Treatment	1,478,987	.8002802	.0003287	.7996359	.8009245
combined	16,351,307	.7833876	.0001019	.7831879	.7835873
diff		.0185725	.0003551	.0192686	.0178764

Table 9: Regression of Treatment on Probability of Return

	180 Day Return b/se
Seller Feedback Score	-0.000000355***
	8.51e-09
Percent Positive Dummy	
excluded: 0 < .994	
$\geq .994 < 1$	0.0388***
	0.000924
$= 1$	0.0564***
	0.00154
Treatment Dummy	0.0276***
	0.00139
Item Price	-0.00235***
	0.00000832
Seller Standards Dummy	
excluded: Below Standard	
Standard	-0.0720***
	0.00279
Above Standard	-0.0547***
	0.00232
ETRS	-0.121***
	0.00231
Seller Number of Trans	0.000000230***
	5.30e-09
Used / New Dummy	
excluded: New	
Refurbished	0.153***
	0.00266
Used	0.146***
	0.00136
Constant	0.794***
	0.00274
N	14,207,773

Regression includes controls for (coefficients not displayed) auction type (auction, fixed price), and item category.

instead of EPP as the key variable. Since the treatment was exogenous by definition, this helps rule out any lingering stories of EPP's exogeneity.

Together with the reduced form results of Section 5, the experiment demonstrates three important facts. First, a transaction's quality is an important component of an individual's propensity to return to a platform, over and above an individual's propensity to return to an individual seller. Second, platforms can use an intermediate mechanism – search result rankings – to guide buyers to better quality sellers, alleviating some of the externality issues associated with platform transaction quality. We note that search ranking may be one of many levers that platforms might use. Third, search ranking causally affects buyer purchase decisions. We show this by varying the search ranking of the same search in treatment vs. control groups which gets around the traditional problem of search ranking endogeneity.

7 Discussion

We demonstrated that market platforms face two important challenges. First, there is a reputational externalities across sellers, and second, reputation feedback can be—and in bay's case is—biased for more than one reason. We studied the limits of reputation mechanisms in the face of these problems, their impacts on the marketplace. We then offered an implementable search prioritization strategy that online platforms can use to mitigate the adverse impacts of repetitional externalities and based feedback, and demonstrated its effectiveness through a field experiment on bay's platform..

One nice feature of our search prioritization strategy is that buyers with different search costs will be differentially affected by the intervention. Section 5.3 provides evidence of buyer learning about the platform. If search costs are correlated with experience on the platform, which we suspect they might be, then our intervention naturally separates out new buyers from more experienced ones as the search costs of the latter ought to be lower. Specifically, if experienced buyers are more likely to search extensively (because they have lower search costs driven by more familiarity with ebay), then the intervention should affect them less. This is exactly the strategy a platform would like to implement because it exposes new buyers indirectly to better quality sellers.

We argued that factors such as reputational externalities across sellers play an important role in influencing market platforms, and hence delivering welfare to consumers. The theoretical literature on two-sodded markets has, however, generally ignored this issue. The standard model that Rochet and Tirole (2006) propose for two-sided markets does not allow for externalities between agents on the same side of the market and concentrates on the binary decision of joining a platform or not. While these models may be appropriate for industries such as credit cards (the classic example used in many of these papers), the models cannot capture the complexity of relationships that exist on a large marketplace platforms. In light of the growing importance of online platform markets in the economy, we advocate allocating some focus to models that can incorporate more complex market setups.

References

[1] Bajari, Patrick, and Ali Hortaçsu (2004) "Economic insights from internet auctions," *Journal of economic literature*, 42(2):457–486.

[2] Bar-Isaac, Heski, and Steven Tadelis (2008) "Seller reputation," *Foundations and Trends in Microeconmics,* 4(4):273-351.

[3] Bolton, Gary, Ben Greiner, and Axel Ockenfels (2013) "Engineering trust: reciprocity in the production of reputation information,". *Management Science*, 59(2):265–285.

[4] Cabral, Luis, and Ali Hortacsu (2010) "The dynamics of seller reputation: Evidence from ebay," *The Journal of Industrial Economics*, 58(1):54–78.

[5] Dellarocas, Chrysanthos (2003) "The digitization of word of mouth: Promise and challenges of online feedback mechanisms," *Management science*, 49(10):1407–1424.

[6] Dellarocas, Chrysanthos, and Charles A Wood (2008) "The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias," *Management Science*, 54(3):460–476.

[7] Ghose, Anindya, Panagiotis G Ipeirotis, and Beibei Li (2014) "Examining the impact of ranking on consumer behavior and search engine revenue," *Management Science*, 60(7):1632-1654.

[8] Greif, Avner (1989) "Reputation and coalitions in medieval trade: evidence on the Maghribi traders," *The journal of economic history*, 49(4): 857-882.

[9] Hui, Xiang, Maryam Saeedi, Zeqian Shen and Neel Sundaresan (2014) "From Lemon Markets to Managed Markets: The Evolution of eBay's Reputation System," mimeo.

[10] Klein, Tobias J., Christian Lambertz and Konrad O. Stahl (2014) "Market Transparency, Adverse Selection, and Moral Hazard," mimeo.

[11] Rochet, Jean-Charles, and Jean Tirole (2006) "Two-Sided Markets: A Progress Report," *The RAND Journal of Economics,* 37(3):645-667.

Appendix

In this appendix we explore a series of robustness checks and additional specifications. For all of these regressions, the dependent variable is a binary indicator of whether a buyer purchases again within 180 days. For compactness we do not report the results for 60 days and whether the buyer ever returns to ebay, but results are very similar for those variables.

Table 10 breaks out the main regression for new (column 1) versus used (column 2) products. The EPP coefficient is larger for new products than for used ones. Table 11 displays specifications with seller fixed effects (column 2) and buyer fixed effects (column 3). Note that both because of computational issues and an incidental parameters problem, running non-linear models with such a large number of fixed effects is not feasible. Hence, we run these as linear probability models. In order to compare to a baseline, column 1 reports the coefficients from the linear probability model with the same set of controls as above.

Table 10: New vs. Used

	New	Used
	b/se	b/se
Seller Feedback Score	-2.17e-08	-0.00000107***
	1.29e-08	9.15e-08
Percent Positive Dummy		
excluded: 0 < .994		
≥ .994 < 1	-0.0611***	-0.0280***
	0.00172	0.00485
= 1	-0.104***	-0.0366***
	0.00289	0.00492
EPP	1.134***	0.819***
	0.00815	0.0153
Item Price	-0.00196***	-0.000777***
	0.0000164	0.0000292
Seller Standards Dummy		
excluded: Below Standard		
Standard	-0.0757***	-0.119***
	0.00468	0.00969
Above Standard	-0.0677***	-0.0998***
	0.00394	0.00956
ETRS	-0.0843***	-0.137***
	0.00397	0.00954
Seller Number of Trans	1.62e-08*	0.000000554***
	7.56e-09	4.71e-08
Constant	-0.432***	-0.254***
	0.00762	0.0154
N	9,669,511	1,951,384

Regression includes controls for (coefficients not displayed) auction type (auction, fixed price), item category, and the transaction number of the buyer at the time of the focal observation.

Table 11: Fixed Effects Regressions

	OLS	Seller Fixed Effects	Buyer Fixed Effects
	b/se	b/se	b/se
Seller Feedback Score	7.30e-10	-0.000000122***	-7.34e-09***
	1.07e-09	5.09e-09	9.76e-10
Percent Positive Dummy			
excluded: $0 < .994$			
$\geq .994 < 1$	-0.00875***	0.000774*	-0.000930***
	0.000187	0.000377	0.000163
$= 1$	-0.0118***	0.00683***	-0.00200***
	0.000281	0.000637	0.000242
EPP	0.130***	0.421***	0.0287***
	0.000804	0.00285	0.000749
Item Price	-0.000302***	-0.000272***	-0.000136***
	0.00000181	0.00000284	0.00000178
Seller Standards Dummy			
excluded: Below Standard			
Standard	-0.00744***	0.00402***	-0.00129**
	0.000471	0.000686	0.000404
Above Standard	-0.00751***	-0.00299***	0.000240
	0.000415	0.000505	0.000355
ETRS	-0.0112***	-0.00308***	-0.00104**
	0.000418	0.000529	0.000361
Seller Number of Trans	-3.18e-09***	1.08e-08***	3.73e-09***
	5.71e-10	1.95e-09	5.28e-10
Used / New Dummy			
excluded: New			
Refurbished	-0.00228***	0.00260**	0.00122*
	0.000530	0.000804	0.000475
Used	-0.00113***	0.00465***	0.000947***
	0.000228	0.000452	0.000216
Constant	0.448***	0.269***	1.064***
	0.000788	0.00193	0.000767
N	11,883,455	11,883,455	11,883,455

Regression includes controls for (coefficients not displayed) auction type (auction, fixed price), item category, and the transaction number of the buyer at the time of the focal observation.