Master's Thesis

On Price Setting and Online Ratings

Chair of Economic Theory
Universität Basel

Supervised by:
Professor Dr. Aleksander Berentsen

Author:
Manuel Buchmann
Sevogelstrasse 140
CH-4052 Basel, Switzerland
manuel.buchmann@unibas.ch
Submission Date: June 5, 2014

Abstract

The goal of this thesis is to describe price setting behavior of sellers in an online market with a rating (reputation) system. I first review existing literature on various subjects in the area of online reputation systems. Key findings include that online ratings do matter for buyer's decisions and that they can influence seller and buyer behavior. I then set up a model which, in its basic version, shows that new sellers have an incentive to charge lower prices in early periods in order to build up a reputation. The model is extended to incorporate potential bribing and cheating behavior by the seller, highlighting the role of rating systems in markets with adverse selection and moral hazard.

Contents

1 Introduction

An ever growing portion of bilateral trade happens over the internet, with buyers as well as sellers being anonymous to various degrees. This gives rise to different types of information asymmetry. Consider for example adverse selection, as first described by Akerlof (1970) in the context of a used car market. Without any public knowledge of the participants' trustworthiness or the quality of products and services offered, markets quickly deteriorate in average quality and possibly break down altogether. Untrustworthy participants have an incentive to enter the market due to the ease of cheating and honest participants are driven out or don't even enter such a market. The characteristics of online markets, where participants are largely anonymous and it is often hard to take legal measures, only aggravate this problem and additionally give rise to moral hazard issues. Because there is no personal interaction and no possibility to inspect the goods or services there is an incentive to deliver a lower quality product than advertised, package the goods badly or to not deliver at all (Klein et al., 2013). One possibility to overcome such effects is to implement a rating system in which the participants can publicly rate each other and thus are able to build a reputation on an online platform. The rating system mitigates moral hazard issues by punishing misbehavior and acts as a signaling device for true quality to counteract adverse selection (Dellarocas, 2006). I discuss such rating systems and present known theoretical issues from existing literature in Sections 2.1-2.3.

The effectiveness of such rating systems has been extensively discussed in the empirical literature especially during the mid-2000s. The overwhelming majority of these studies report a statistically significant effect of ratings on prices and probability of sale, among other visible effects. Prominent examples include for example Melnik and Alm (2002), Resnick and Zeckhauser (2002), Lucking-Reiley et al. (2007) or Cabral and Hortacsu (2010). While the lion's share of studies focuses on eBay.com, more recent studies observe similar effects for other platforms as well[1]. I conduct a more thorough investigation of existing empirical literature in Section 2.4.

[1]see for example Anderson and Magruder (2012) or Jolivet et al. (2013).

Based on this evidence I propose a model focusing firstly on a new seller entering a market with an online rating system and the question of how the existence of said system affects his behavior. Intuitively it seems obvious that a new seller without any ratings has to "pay his dues" in terms of setting a lower price to make up for the information uncertainty of the potential buyers. Indeed, even before the internet era, Shapiro (1983) found that in a market where product quality cannot be observed easily "[...] such a seller must sell his product at less than cost: he cannot command those prices associated with high quality items until his reputation is established". Modeling interactions on an online market with free re-entry under new identities as a Prisoner's Dilemma, Friedman and Resnick (2000) find that a "Pay-Your-Dues" strategy will promote cooperation. Przepiorka (2013) models and empirically confirms that a sellers positive ratings positively correlate with both the probability of sale and, more importantly, the price of the good in fixed price offers. Thus the results from Shapiro (1983) still seem to hold in online markets. My model confirms this main observation and adds to the existing literature in 2 ways. Firstly, I differ methodologically from other models by separating the probability of a seller meeting a buyer from the probability of a buyer buying the product. The latter does not depend on a seller's reputation but on the price set by the seller, only the arrival rate depends on the reputation. Secondly, in 2 extensions I model additional buyer behavior that would be expected in markets dominated by the threat of moral hazard or adverse selection. These include sellers choosing to defect based on economic incentives (moral hazard) and low-quality sellers increasing their rating by explicitly or implicitly paying off buyers. Especially the latter has rarely been discussed in theoretical or empirical literature, but anecdotal evidence suggests that it might be an important issue in practice, thus potentially weakening the rating system's role of preventing adverse selection. My model shows that there is an incentive for low-quality sellers to bribe buyers in order to induce higher ratings and that it can be optimal for sellers to cheat on buyers, in some cases even for high-reputation sellers.

2 Rating Systems

2.1 Design of Rating Systems

Online rating systems are very common and widespread. Nearly every major (and minor) website incorporates some kind of rating system, ranging from Facebook-Likes to Google search rankings. The focus of this thesis lies on rating systems that facilitate a trade in goods or services over the internet. Important examples include online auction houses like eBay or the Chinese Taobao, retailers and online marketplaces like Amazon or dedicated websites that focus on the rating of restaurants, local businesses or holiday residences (Yelp, TripAdvisor). These rating systems differ in various characteristics to meet the requirements of the platform they were designed for. For example, on eBay users rate each other reciprocally while on Amazon only buyers give ratings and write reviews. On Amazon, products are reviewed and graded on a ranking of 1-5 stars, these reviews can in turn be rated by other customers as either helpful or not helpful. The average rating of a product is displayed in tenth-stars (for example 4.6/5 stars). Similarly, on Yelp customers can rate restaurants and local businesses on a scale of 1-5 stars, the restaurant's average rating is then displayed in half-star increments. On TripAdvisor customers rate hotels in different quality dimensions on a scale of 1-5, the site displays the average rating for all dimensions in half-stars as well as an overall share of recommendations ("thumbs-up") of the hotel. Again, reviews can be labeled as helpful or not helpful by other customers and active raters are rewarded with a profile that displays their reviews and amount of "helpful" reviews. eBay ratings are only binary (users rate each other either positively, neutrally or negatively) and only short comments can be given. The site displays the share of positive ratings, amount of each rating received as well as a net feedback score.

When designing a rating system, it is important to differentiate between the two main purposes of rating systems. Following Dellarocas (2006), rating systems can serve as *signaling device* to combat adverse selection: A high rating then indicates a high-quality product. A rating system can also be used to mitigate moral hazard, it is then a *sanctioning device*: A high rating serves as reward for cooperation and a low rating as punishment for misbehavior. A typical example of a rating system being used as a signaling device would be Amazon, where there is very little moral hazard and ratings are usually only concerning product quality. On eBay, on the other hand, users mainly rate their trading partners and not product quality. Due to the large presence of moral hazard, this rating system is mainly a sanctioning device (Dellarocas, 2006). An empirical confirmation of this statement is provided by Jin and Kato (2006) who find in an experimental study that high reputation on eBay does infer good-faith sellers but not necessarily high-quality products.[2]

Several unique characteristics of online markets affect the functioning of such systems. For example, Friedman et al. (2007) classify three threats to the integrity of rating systems: Whitewashing, Incorrectly Reported Feedback and Phantom Feedback (or sybil attacks). Whitewashing describes the ability to create a new identity and start fresh. If sellers can whitewash their identities the sanctioning function of ratings is very limited (i.e. sellers can't be sanctioned below their starting reputation). This is definitely an important issue and discussed further in Section 2.3. Incorrectly reported feedback can refer to either dishonest feedback, for example due to fear of retaliation or an incentive to be seen as a "good" rater (see Section 2.1.1) or to the issue that not enough ratings might be given in general (Section 2.2). Finally, a sybil attack or phantom feedback describes a situation in which a seller will create a large number of fake identities to boost the reputation of his main identity. Anecdotal evidence suggests that fake reviews are not rare and grow in importance. For example, The New York Times has reported about businesses hiring workers to post fake positive and negative feedback on Yelp (Segal, 2011) or writers hiring reviewers to post positive reviews of their books on the internet (Streitfeld, 2012*a*). In a well documented incident from 2004, Amazon's Canadian site revealed the names of supposedly anonymous book reviewers as a result of a system glitch; it turned out that many of these reviews were posted by the authors themselves (Smith, 2004). As it

[2] It should be noted that these distinctions between Amazon and eBay are slowly disappearing, as both platforms branch out. For example, more and more small independent sellers sell their products over the Amazon marketplace and large businesses use eBay to sell their products in fixed price offers.

is not possible to observe fake reviews directly, there is only limited empirical evidence on this issue. In recent years some studies have found ways to identify false reviews, however. For example Luca and Zervas (2013) use reviews filtered by Yelp's algorithmic detector as a proxy for fake reviews and Mayzlin et al. (2012) use differences between two hotel rating sites (TripAdvisor and Expedia) as well as characteristics of neighboring hotels to identify fake reviews. Both studies find clear evidence for the presence of fake reviews.[3] There have been attempts to counteract such behavior. For example, Expedia only allows users who booked a service through the website to write reviews. There is however a trade-off between the believability of reviews and the amount of reviews being posted: websites that require verification will attract less reviewers and are thus less popular. One possible compromise would be the example of Amazon, where verified buyers are highlighted, but there is still the question of whether to include unverified reviews into the average rating (Mayzlin et al., 2012).

2.1.1 eBay

The bulk of existing research focuses on eBay[4], mainly due to its economic significance but also the availability of data and the wealth of potential interaction between buyers and sellers (as opposed to, for example, Amazon where buyers only rarely interact with sellers directly and ratings/reviews mostly focus on product quality). eBay offers auctions as well as fixed price offers, so different types of interactions can be observed (see for example Przepiorka (2013)). The classic eBay approach of rating systems is characterized by a reciprocal element: not only do buyers rate sellers, but sellers can rate buyers in return. At first this seems to be an appropriate approach, as buyers often also act as sellers in other auctions and vice-versa. However, such a system opens the door for retaliatory behavior (i.e. giving bad ratings in return for receiving bad ratings). The fear of retaliation would then in theory lead to a reduction in negative ratings given, even if a negative rating would be

[3] More details on these studies are provided in Section 2.4.

[4] It should be noted that in recent years, eBay has become less important in the empirical literature and attention has shifted to other platforms.

appropriate.

Chwelos and Dhar (2007) provide a simple game-theoretic model to illustrate this weakness. They argue that there is an incentive to always rate the other party positively (or not at all) for fear of retaliation, independent of the actual product quality. Similarly, it is always superior to respond to a positive rating with a positive rating because giving too many negative ratings could have a negative impact on future interactions (the user could be seen as a non-cooperative, strict rater which makes him an undesirable business partner).

To show this effect within a simple game theoretic framework Chwelos and Dhar (2007) set up the following "Feedback Game", where P_j are future payoffs from feedback received[5], R_j are future payoffs from feedback provided[6], and p, c, V, E_j are the price of the product, the acquisition cost to the seller, the valuation of the product by the buyer and the cost of effort needed to provide a rating. $+$ denotes giving positive feedback and $-$ giving negative feedback.

		Seller: $+$	Seller: $-$
Buyer	$+$	$P_B + R_B - p + V - E_B$, $P_S + R_S + p - c - E_S$	$-P_B + R_B - p + V - E_B$, $P_S - R_S + p - c - E_S$
	$-$	$P_B - R_B - p + V - E_B$, $-P_S + R_S + p - c - E_S$	$-P_B - R_B - p + V - E_B$, $-P_S - R_S + p - c - E_S$

Figure 1: Feedback Game following Chwelos and Dhar (2007)

Clearly, the unique Nash Equilibrium of this Feedback Game is $(+/+)$, as there is never an incentive for the buyer or for the seller to deviate from this strategy. The rating process is assumed to be purely strategic while the actual transaction is "sunk" and not relevant for the rating given. This means that in the Nash equilibrium of this model, there will never be negative ratings.

[5] defined as $\frac{\beta_j}{1-r_j}$, where β_j is the per-period net benefit of receiving a positive rating and r_j is the discount rate.

[6] defined as $\gamma_j \frac{\beta_j}{1-r_j}$, where γ_j is the percentage gain/loss from providing positive/negative feedback.

The authors also evaluate a model with the option "Silence", i.e. giving no rating at all. There, remaining silent can also be part of a Nash Equilibrium, but again there will never be negative feedback.

This idea of reciprocity and retaliation has been thoroughly discussed in the literature. Empirical evidence for this type of behavior has been reported for example by Resnick and Zeckhauser (2002), Cabral and Hortacsu (2010), Klein et al. (2006), Bolton et al. (2013) or Dellarocas and Wood (2008). They all find an extremely low share of total negative and neutral feedback on eBay (0.9%, 1.4%, 1.8%, 1.4% and 0.72% respectively). Bolton et al. (2013) also report this number for the (eBay-owned) Brazilian online marketplace MercadoLivre, where the rating takes place during a 21-day blind period during which the ratings are not made public. On this website, the share of negative feedback is estimated at 18.7% for ratings given by buyers and 29.2% for ratings given by sellers. They also show that on the website RentACoder.com, following a change of the feedback system from an open to a double-blind system in May 2005, the correlation between buyer and seller feedback reduced drastically. Klein et al. (2006) additionally estimate conditional feedback probabilities and show that the probability of receiving negative feedback after giving negative feedback increases to 36.7% (compared to 0.002% after giving positive feedback). In a different approach, Jian et al. (2010) model three different rating strategies for traders: (1) never giving feedback, (2) giving feedback unconditionally (i.e. always giving feedback, independent from the trading partner) and (3) reciprocate (i.e. giving feedback if and only after the trading partner gives feedback). Using a large dataset from 1999 provided by eBay they estimate their model and find that 38% of the buyers and 47% of the sellers give feedback unconditionally, 23% of the buyers and 20% of the sellers reciprocate and the rest remains silent.

This effect did of course not go unnoticed by eBay, which resulted in their feedback system changing in May of 2008. Sellers are now no longer able to rate buyers neutrally or negatively, only positive ratings (or remaining silent) are allowed. This change is supposed to reduce the buyers' fear of retaliation and thus encourages giving negative ratings when appropriate. As most of the existing research focuses on eBay data from *before* the change of the system I do not go into further detail on this change.[7] The important lesson to learn from the original eBay feedback system is that traders do indeed rate strategically and respond to the design of rating systems.

2.2 Buyer Behavior

While the focus of this thesis lies on the seller-side of the online transaction, it is interesting to consider the buyer-side as well. Often it is just assumed that buyers leave (1) sufficient and (2) honest feedback. But this is not always obvious. Miller et al. (2002) illustrate this point by noting that smart parents may not want to reveal the names of their favorite baby-sitters in an environment where supply of good baby-sitters is limited. Additionally, posting feedback requires some effort from the buyer and no direct gain, so there is a large incentive to free-ride: the provision of feedback can be considered to be a public good (Avery et al., 1999). As seen in the example of eBay in Section 2.1.1, buyers might be reluctant to rate negatively in order to preserve their own reputation or for fear of retaliation. There could also be direct or indirect interaction between the seller and the buyer which influences the rating given (i.e. rewarding the buyer for giving positive feedback). Judging by empirical results, at least not receiving enough feedback does not seem be an important issue in practice. For eBay, Resnick and Zeckhauser (2002) find that buyers will leave feedback 52.1% of the time, while Cabral and Hortacsu (2010) find that feedback is left in 40.7% of transactions. However, as already discussed in the context of eBay, dishonest feedback is likely to be an issue. Anecdotal evidence also suggests that some sellers do try to compensate buyers for giving positive reviews and that buyers largely accept these terms. For example, The New York Times[8] reported of a case in which a seller of tablet leather cases and stun guns on Amazon would not only deliver the ordered product, but also a letter promising a full refund in return for a positive review. This scheme resulted in an overwhelmingly positive review page for the product with barely any critical comments.

[7] I refer to Klein et al. (2013) for a detailed analysis on the effects of the change.

[8] see Streitfeld (2012*b*)

In an attempt to model a rating system that would counteract these issues, Miller et al. (2005) suggest a mechanism (the so-called Peer Prediction Method) that elicits honest feedback from buyers. An exact representation of their proposed model would exceed the scope of this thesis, however the general idea is to pay raters based on how likely it is that their rating would be given by another rater, the so-called reference rater. This payment is supplied by a third buyer whose rating does not affect the transfer paid. Miller et al. (2005) find that in the Nash Equilibrium, feedback is honest and the system breaks even.

2.3 Seller Behavior

As noted by Shapiro (1983), in a market where product quality cannot be directly observed (such as an online market for experience goods) there should be an incentive for sellers to undercharge in earlier periods until they have been able to build up a reputation, which is achieved in online markets by accumulating enough positive ratings. Przepiorka (2013) investigates this issue by first modeling and then empirically confirming it with data from eBay. He sets up the following binary one-shot trust game to model the buyer-seller interaction.

First, the buyer decides whether to buy or not, then, if the buyer does buy, the seller decides whether to ship or not. If no trade takes place both receive payoff P, if the seller ships both receive payoff R and if the buyer buys but the seller does not ship the buyer receives payoff S and the seller receives payoff T. The payoffs are ordered as follows: $T > R > P > S$. Additionally, a buyer is active in the market only once while a seller stays in the market with probability δ (which is heterogeneous and private information). If a trade takes place, the seller gives a rating with probability φ. This rating is positive if the seller ships and negative if the seller does not ship. Sellers can re-enter the market with a new identity at no cost. Naturally, buyers prefer to buy from sellers who received positive ratings in the past (as it turns out, a seller with a positive rating will always ship). Thus, a newcomer to the market will have to charge a lower price to incentivize buyers to consider him. The price discount c is given by equalizing expected payoffs from buying from a newcomer with buying from an established seller:

$$\alpha R + (1 - \alpha) S + c = R \tag{1}$$

where α stands for the probability that a seller ships. The discount c is therefore given as

$$c = (1 - \alpha)(R - S) \tag{2}$$

By comparing sellers' expected payoffs from shipping and not shipping, Przepiorka (2013) shows that a seller has an incentive to always ship if he is a long term type with $\delta > \delta^*$ and to never ship if he is a short term type with $\delta < \delta^*$, with δ^* being given as

$$\delta^* = \frac{T - R}{\varphi c + (1 - \varphi)(T - R)} \tag{3}$$

For the buyer, the probability of a seller shipping now becomes the probability of meeting a long term seller (when interacting with a newcomer, an established seller always ships). Assuming a uniform distribution, this probability becomes $1 - \delta^*$. Plugging in to equation (2) and re-arranging, Przepiorka then finds the equilibrium price discount c^*:

$$c^* = \frac{1}{2\varphi} \left[[(1 - \varphi)^2 (T - R)^2 + 4\varphi (T - R) R]^{1/2} - (1 - \varphi)(T - R) \right] \tag{4}$$

Based on these findings, he sets up three hypotheses: (1) a seller with higher reputation will set a higher price in a fixed price offer, (2) buyers will pay a higher price to a seller with high reputation in auctions and fixed price offers and (3) if the market does not clear, a seller with high reputation will be more likely to sell his item in auctions and fixed price offers. He tests these hypotheses by collecting a large sample of eBay data from offers for SD memory cards in the last quarter of 2006. He performs logit regressions on sales as well as OLS regressions on selling prices, considering both fixed price offers ("Buy-It-Now") and standard auctions. He finds a statistically significant positive effect of positive ratings and negative effect of negative ratings on both probability of sale as well as selling prices. Additionally, he finds that in fixed price offers the prices increase with the reputation of the sellers. This confirms the initial idea that newcomers have to charge lower prices than established sellers.

An important feature of online markets is the ease, with which a new identity can be attained. There is a large incentive to cheat your customers for a few periods and once too many negative ratings are given to just "whitewash" and create a new account. Friedman and Resnick (2000) model this problem and show that there will still be cooperation in such a system, however there will emerge a convention under which new players have to "pay their dues", meaning that they have to accept worse conditions in the first few periods of the game. They show this by setting up a basic infinitely repeated prisoner's dilemma game and add public record keeping (i.e. each player knows the other player's past actions). It is clear that under fixed identifiers (e.g. real names) total cooperation is a sustainable equilibrium. This is not the case if identifiers can be changed freely, then it is always possible to defect and start with a fresh identity without being punished. Friedman and Resnick (2000) therefore suggest the PYD ("paying your dues") strategy: if a newcomer meets a compliant veteran, the newcomer cooperates and the veteran defects. They show that, in this system, everyone except a veteran meeting a newcomer cooperates. Thus the newcomer pays his dues with worse treatment now in return for better treatment once he is a compliant veteran himself.

One way to avoid the problem of whitewashing would be to ask an entry fee from participants to make acquiring a new identity more costly. However, due to the competitive nature of online markets, this is often not possible or viable for the site owners (Friedman et al., 2007).[9] In general, it is usually assumed that traders can acquire new identities with little to no cost. However, there is a trend towards real names becoming more common on the internet (mainly due to the rise of social media and their integration into other services), so this might change in the future.

[9] As an alternative option Friedman et al. (2007) suggest a system where users are required to reveal their real names to site owners and receive one free anonymous identity, having to pay for additional ones.

2.4 Empirical Evidence

After focusing primarily on theoretical issues so far, I now want to summarize some empirical evidence on the noted issues. Most of the available research focuses on eBay auctions, using data from before the change of the rating system in 2008. There is a vast amount of evidence indicating that prices increase with reputation. This is shown for example by Melnik and Alm (2002) for the case of collectible coins, Dewally and Ederington (2006) for comic books or Houser and Wooders (2006) for Pentium processors, among many others. An exception is Resnick and Zeckhauser (2002) who find that reputation does not increase prices of Beanie Babies and MP3 players, but it does increase the probability of sale. Positive effects on the probability of sale are also reported by various other studies (see for example Przepiorka (2013), Livingston (2005) or Dewan and Hsu (2004)). Livingston (2005) additionally finds that the effect of positive ratings decreases with increasing reputation (i.e. a decreasing marginal benefit of ratings). Mixed results have been found with respect to the effect of negative ratings. In an experiment where participants were asked to judge the trustworthiness of constructed eBay seller pages, Ba and Pavlou (2002) find that negative ratings had a larger impact than positive ratings. They cannot confirm this with real data however; they find that while positive ratings increase prices of music, software and electronics products, negative ratings have no significant impact on prices. Eaton (2005) finds that the presence of negative feedback reduces the probability of sale for electric guitars but actually increases their price.[10] Przepiorka (2013) and Mickey (2010) find that negative ratings have a negative impact on prices of electronics, Lucking-Reiley et al. (2007) confirm this for collectible coins and additionally find that negative ratings have a stronger impact on prices than positive ratings. In a different approach, Resnick et al. (2006) set up a controlled experiment in which an established high-reputation seller of vintage postcards sold the same items under his main identity, as well as several new identities without reputation. They find that prices were on average 8.1% higher under the reputable identity than the new identities. Comparing new

[10] One potential explanation for this observation was suggested to be the fact that more active and experienced sellers are more likely to receive negative feedback at some point during their eBay career and potentially have higher reservation prices.

identities with and without existing negative feedback, they find that a small amount of negative feedback does not affect buyers' willingness-to-pay.

While the studies mentioned so far focus on the effect of ratings on the buyer's decision, there is also evidence with regard to seller behavior. For example, Jin and Kato (2006) use another experimental design and buy collectible baseball cards from different eBay sellers to have them professionally graded afterwards. They find that while high reputation sellers have lower default/counterfeit rates than low reputation sellers, the quality of the cards conditional on authentic delivery does not differ with reputation. Additionally they find that offers with high quality claims received higher prices but were more likely to be fraudulent and high reputation sellers are less likely to claim very high card quality. Bruce et al. (2004) observe that a seller who defaulted on a transaction before is more likely to do so again and that sellers who would leave the market soon were more likely to default as well. Similarly, Cabral and Hortacsu (2010) observe that a seller with more negative feedback is more likely to leave the market and just before leaving a seller receives more negative feedback than his lifetime average. There are two alternative interpretations of these observations: Either the seller is planning to exit the market and as a consequence decides to lower his quality/defaults on transactions while still profiting from his reputation, or he exits the market as a result of the negative ratings which were induced exogenously, or possibly just a result of bad luck. Anecdotal evidence points towards the first explanation being prevalent. For example, during their experiment Jin and Kato (2006) encountered two sellers who would intentionally build up a positive reputation to then commit a series of defaults and abandon their accounts soon after.

In recent years researcher's attention has shifted away from eBay however, and alternative rating systems have been empirically investigated. For example Chevalier and Mayzlin (2006) compare book sales on BarnesAndNoble and Amazon and find that a relatively higher rating of a book on one site compared with the other leads to relatively higher sales on that site, thus showing that ratings also have an impact on websites where moral hazard is not the dominant issue. Focusing on the major French online marketplace

PriceMinister and using an exceptionally large and complete dataset obtained directly from PriceMinister, Jolivet et al. (2013) find a strong positive effect of reputation on prices. Employing a regression discontinuity design, Anderson and Magruder (2012) analyze restaurant ratings on Yelp and observe that an extra half star rating causes restaurants in San Francisco to sell out 49% more frequently. Additionally, they look into the issue of sybil-attacks: Robustness checks confirm that restaurants do not manipulate ratings in a confounding, discontinuous manner.[11] Using similar methodology but a different outcome variable, Luca (2011) finds that restaurants in Seattle increase their revenue by 5-9% per extra star on Yelp. He also finds that it is mostly independent restaurants (as opposed to chain restaurants) that benefit from Yelp, indicating that online ratings function as a substitute for classic (offline) reputation. Focusing on Boston area restaurant reviews that have been marked as fraudulent by Yelp's own algorithmic indicator,[12] Luca and Zervas (2013) find that 16% of reviews are identified as fake and that review fraud is likely to be a response to economic incentives, rather than a strategy used by a small number of unethical businesses. In a different approach of trying to identify fake reviews, Mayzlin et al. (2012) exploit an organizational difference between the hotel rating sites TripAdvisor and Expedia. While on TripAdvisor every user can post a review of any hotel, on Expedia only users who booked at least one night through the website can post a review. Thus the cost of posting disingenuous reviews should be higher on Expedia than on TripAdvisor. The authors define several hotel characteristics that are likely to affect the likelihood of hotels posting fake reviews[13] and compare reviews of such hotels on the two sites, using a slightly unconventional difference in differences approach. For example they would compare the ratio of low re-

[11]Yelp rounds the true average star rating to half-stars, so the idea that is tested for is that restaurants would be incentivized to leave fake reviews to push their average just above a rounding point (for example from 3.24 to 3.26). No evidence for this behavior was found.

[12]Such reviews are filtered and not displayed on the main listings, however they can still be accessed by solving a CAPTCHA puzzle.

[13]These characteristics are based on the following assumptions: a hotel with close neighbors is more likely to receive negative fake reviews; small owners, independent hotels and hotels with small management companies are more likely to engage in review manipulation than multi-unit owners, branded chain-hotels and hotels with a large management company.

views to total reviews between hotels with and without a close neighbor and then ask whether this difference differs between TripAdvisor and Expedia. They find that promotional reviews are more common on TripAdvisor than on Expedia and that organizational characteristics of hotels affect the extent of review manipulation.

Summarizing empirical evidence, I want to highlight the following points:[14]

1. Positive ratings have positive effects on prices and probability of sale.
2. The effect of negative ratings is less pronounced, the evidence points toward negative effects in general however.
3. There is a decreasing marginal benefit to ratings.
4. High reputation sellers are less likely to default.
5. Sellers are more likely to default shortly before exiting the market.
6. Rating manipulation is driven mainly by economic incentives and organizational characteristics of sellers.

In the following section I try to address all of these issues within the framework of my model.

3 The Model

3.1 Baseline Model

The goal of this basic approach is to model the price setting behavior of a new seller entering the market under the existence of an online rating system. I do not include a quality dimension in this basic setting, meaning that as long as a rating is given, the rating will be positive. The only effect a positive rating from a buyer has, is to increase the probability of the seller meeting another

[14]A complete overview over the mentioned studies can be found in Table 1 in the Appendix A.1.

buyer in the next period (i.e. the arrival rate). In a real world example, one could imagine this to be the probability with which a buyer would click on an offer, considering the rating that is displayed. The higher arrival rate can be interpreted as a premium for information; only a risk-loving buyer would choose a seller without any ratings, assuming the existence of alternatives and similar prices. Once a seller has (positive) ratings there should be more buyers willing to trade with him and thus a higher probability of being matched with a buyer within the model framework.

There is one seller and many buyers, interacting in 2 periods. There is one indivisible good which can be sold once in each period. $V_{i|j}$ describes the seller's value (or utility) for the rest of the game in period i, given his net reputation j.[15] I characterize the buyer's decision in two parts. Firstly, the probability with which a buyer considers to trade with a specific seller is the arrival rate $\alpha \in [0,1]$. Secondly, buyers have reservation utility $R \sim Uniform(0,1)$. A trade occurs only if the price P is below reservation utility R. Therefore, a buyer first decides whether to consider a seller based on the seller's reputation and his own risk-aversion (not modeled here). Then he decides whether to buy based on the price set by the seller and his own reservation utility. If a trade occurs in period 1, the buyer gives the seller a public rating on the online platform. This rating changes the arrival rate in period 2 from α to α', with $\alpha' > \alpha$. If no trade occurs, no ratings are given and the arrival rate does not change in period 2.

First, I consider the seller's value function in period 1. With probability α he meets a buyer. The first integral term in equation (5) describes the probability of the price being below the buyer's reservation utility. Only if the price is below the buyer's reservation utility there can be a trade. These two probabilities are multiplied with the price and the continuation value in the case that there is a positive rating in period 2. The second integral describes the probability of the price being above the buyer's reservation utility and thus no trade taking place despite the buyer and the seller meeting. With this probability, the continuation value of no rating in period 2 applies. The

[15] Net reputation is defined as the difference between positive and negative ratings. In this example where only positive ratings are possible, this is equal to the total amount of ratings received.

last term describes the case of the seller not meeting a buyer.

$$\begin{aligned} V_{1|0} = &\alpha \left[\int_{P_1}^{1} f(R)\,\mathrm{d}R\,[P_1 + V_{2|1}] + \int_{0}^{P_1} f(R)\,\mathrm{d}R V_{2|0} \right] \\ &+ (1-\alpha)\,V_{2|0} \end{aligned} \tag{5}$$

To summarize: With probability $\alpha \int_{P_1}^{1} f(R)\,\mathrm{d}R$ the seller meets a buyer and the price is below the buyer's reservation price. Thus a trade takes place, the seller receives the price and the continuation value from a higher rating in period 2. With probability $\alpha \int_{0}^{P_1} f(R)\,\mathrm{d}R$ the seller meets a buyer but the price is too high for the buyer and no trade takes place. The seller does not receive the price and does not receive a positive rating, thus receiving the continuation value from a lower reputation in period 2. Lastly, with probability $(1-\alpha)$ the seller does not meet a buyer at all and again does not receive a price or a positive rating.

In the second and last period, no more value is gained from a trade not taking place, so the only relevant component is the case of a trade taking place. There are two different possibilities: either the buyer didn't have a successful trade in period 1 and the arrival rate does not change or the buyer did have a successful trade in period one and the arrival rate increases from α to α'.

$$V_{2|0} = \alpha \left[P_2 \int_{P_2}^{1} f(R)\,\mathrm{d}R \right] \tag{6}$$

$$V_{2|1} = \alpha' \left[P_2' \int_{P_2'}^{1} f(R)\,\mathrm{d}R \right] \tag{7}$$

To find the optimal prices P_1 and P_2 we first need to solve the optimization problem in period 2:

$$\arg\max_{P_2} \alpha \left[P_2 \int_{P_2}^{1} f(R)\,\mathrm{d}R \right] \tag{8}$$

$$\arg\max_{P_2'} \alpha' \left[P_2' \int_{P_2'}^{1} f(R)\,\mathrm{d}R \right] \tag{9}$$

Solving this problem yields $P_2 = P_2' = \frac{1}{2}$. Next, we calculate the values in period 2 by plugging in this price in equations (6) and (7) and get $V_{2|0} = \frac{1}{4}\alpha$ and $V_{2|1} = \frac{1}{4}\alpha'$.

Using these values in equation (5) we get our new optimization problem in period one:

$$\arg\max_{P_1} \alpha \left[\int_{P_1}^{1} f(R)\,\mathrm{d}R\,[P_1 + \frac{1}{4}\alpha'] + \int_{0}^{P_1} f(R)\,\mathrm{d}R\,\frac{1}{4}\alpha \right] + (1-\alpha)\,\frac{1}{4}\alpha \tag{10}$$

Solving this problem yields:

$$P_1 = \frac{1 - \frac{1}{4}\alpha' + \frac{1}{4}\alpha}{2} \tag{11}$$

which is equivalent to:

$$P_1 = \frac{1 - V_{2|1} + V_{2|0}}{2} \tag{12}$$

Following our assumption that $\alpha' > \alpha$, it is easy to see that the numerator is smaller than 1 and thus the whole expression is smaller than $\frac{1}{2}$. Therefore we can say that $P_1 < P_2$ for any $\{\alpha, \alpha'\} > 0$ and $\alpha' > \alpha$.

This means that it is optimal for sellers to undercharge in period 1. This very basic and intuitive result confirms the findings of most other comparable models (see for example Friedman and Resnick (2000), Przepiorka (2013) or Jiang and Chen (2007)).

3.1.1 An additional time period

I now add a third time period to show that the model is not limited to 2 periods and still holds under multiple time periods. The expectation is that the price increases in each time period. The arrival rate has to be defined in a bit more detail now, as there is an additional time step to consider. Whereas before we simply defined it as increasing ($\alpha < \alpha'$), we now add the assumption of decreasing marginal benefit of ratings (i.e. a high reputation seller benefits less from a positive rating than a low repuation seller) or $f'(\alpha^*) > 0$ and $f''(\alpha^*) < 0$. This seems to be a sensible assumption and is confirmed by, for example, Livingston (2005). For our 3 period model this simply means that $\alpha < \alpha' < \alpha''$ and $(\alpha' - \alpha) > (\alpha'' - \alpha')$. The value function in period 1 still is the same as in equation (5).

$$\begin{aligned} V_{1|0} = &\alpha \left[\int_{P_1}^{1} f(R)\,\mathrm{d}R\,[P_1 + V_{2|1}] + \int_{0}^{P_1} f(R)\,\mathrm{d}R\,V_{2|0}) \right] \\ &+ (1-\alpha)\,V_{2|0} \end{aligned} \tag{13}$$

In period 2 however, we now have to factor in the continuation values from period 3:

$$\begin{aligned} V_{2|0} = &\alpha \left[\int_{P_2}^{1} f(R)\,\mathrm{d}R\,[P_2 + V_{3|1}] + \int_{0}^{P_2} f(R)\,\mathrm{d}R V_{3|0} \right] \\ &+ (1-\alpha)\,V_{3|0} \end{aligned} \tag{14}$$

$$\begin{aligned} V_{2|1} = &\alpha' \left[\int_{P_2'}^{1} f(R)\,\mathrm{d}R\,[P_2' + V_{3|2}] + \int_{0}^{P_2'} f(R)\,\mathrm{d}R V_{3|1} \right] \\ &+ (1-\alpha')\,V_{3|1} \end{aligned} \tag{15}$$

And in period 3 we have three possible states: 0 positive ratings received, 1 positive rating received or 2 positive ratings received:

$$V_{3|0} = \alpha \left[P_3 \int_{P_3}^{1} f(R)\,\mathrm{d}R \right] \tag{16}$$

$$V_{3|1} = \alpha' \left[P_3' \int_{P_3'}^{1} f(R)\,\mathrm{d}R \right] \tag{17}$$

$$V_{3|2} = \alpha'' \left[P_3'' \int_{P_3''}^{1} f(R)\,\mathrm{d}R \right] \tag{18}$$

Similar to the 2-period model, the price in the last period is constant at $\frac{1}{2}$.

Thus $P_3 = P_3' = P_3'' = \frac{1}{2}$. The value functions in period 3 solve to $V_{3|0} = \frac{1}{4}\alpha$, $V_{3|1} = \frac{1}{4}\alpha'$ and $V_{3|2} = \frac{1}{4}\alpha''$. We plug these values into equations (14) and (15) and get the maximization problems in period 2:

$$\begin{aligned} \arg\max_{P_2} \alpha \left[\int_{P_2}^{1} f(R)\,\mathrm{d}R \, [P_2 + \frac{1}{4}\alpha'] + \int_{0}^{P_2} f(R)\,\mathrm{d}R \, \frac{1}{4}\alpha \right] \\ + (1-\alpha)\,\frac{1}{4}\alpha \end{aligned} \tag{19}$$

$$\begin{aligned} \arg\max_{P_2'} \alpha' \left[\int_{P_2'}^{1} f(R)\,\mathrm{d}R \, [P_2' + \frac{1}{4}\alpha''] + \int_{0}^{P_2'} f(R)\,\mathrm{d}R \, \frac{1}{4}\alpha' \right] \\ + (1-\alpha')\,\frac{1}{4}\alpha' \end{aligned} \tag{20}$$

These problems solve to:

$$P_2 = \frac{1 - \frac{1}{4}\alpha' + \frac{1}{4}\alpha}{2} \tag{21}$$

$$P_2' = \frac{1 - \frac{1}{4}\alpha'' + \frac{1}{4}\alpha'}{2} \tag{22}$$

Slightly rewriting and applying our assumption of decreasing marginal benefits of ratings shows that:

$$P_2 = \frac{1 - \frac{1}{4}(\alpha' - \alpha)}{2} < P_2' = \frac{1 - \frac{1}{4}(\alpha'' - \alpha')}{2} < P_3 = \frac{1}{2} \tag{23}$$

This means that a seller who has already received a positive rating will charge a higher price in period 2 than a seller without a rating. This can easily be explained by the fact that with the first rating, the seller has already reaped the largest benefit of undercharging (or payed a part of his dues) and can therefore afford to charge a higher price from this point on (notice, however, that he still charges a lower price than in the following period). Next, I use the prices from (21) and (22) and plug them in to equations (14) and (15) to calculate the values. This yields:

$$V_{2|0} = \alpha \left[\frac{1}{2} + \frac{1}{64}\alpha'^{\,2} + \alpha' \left(\frac{1}{8} - \frac{1}{32}\alpha \right) - \frac{1}{8}\alpha + \frac{1}{64}\alpha^2 \right] \tag{24}$$

$$V_{2|1} = \alpha' \left[\frac{1}{2} + \frac{1}{64}\alpha''^{\,2} + \alpha'' \left(\frac{1}{8} - \frac{1}{32}\alpha' \right) - \frac{1}{8}\alpha' + \frac{1}{64}\alpha'^{\,2} \right] \tag{25}$$

Following equation (12), the optimal price in period 1 P_1 will be:

$$P_1 = \frac{1 - V_{2|1} + V_{2|0}}{2} \tag{26}$$

As this does not yield an analytically intuitive result, I plot the price P_1 depending on α' and α'', with the initial arrival rate α fixed at 0.25, 0.5 and 0.75 in Figure 2. Note that only the right side of the respective graphs

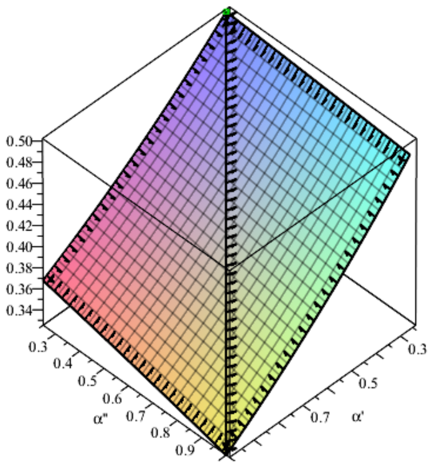

(a) $\alpha = 0.25$

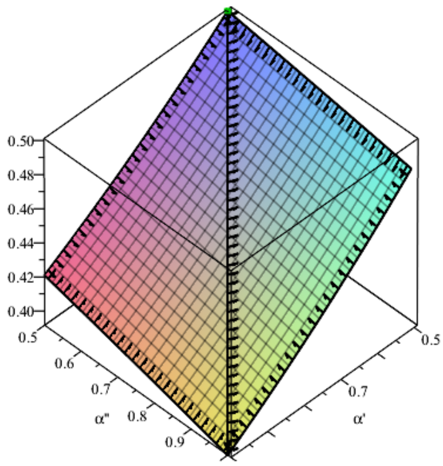

(b) $\alpha = 0.5$

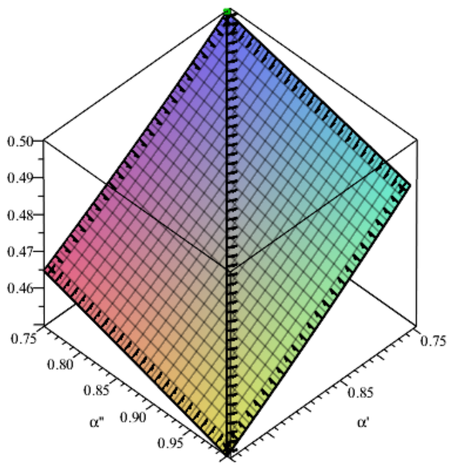

(c) $\alpha = 0.75$

Figure 2: P_1 with initial arrival rate α fixed at different values

is relevant (i.e. $\alpha'' \geq \alpha'$). The lowest possible prices are $P_1^{min} = 0.33$ for $\alpha = 0.25$, $P_1^{min} = 0.39$ for $\alpha = 0.5$ and $P_1^{min} = 0.45$ for $\alpha = 0.75$ at $\alpha' = \alpha'' = 1$; the highest possible price is $P_1^{max} = 0.5$ at $\alpha' = \alpha'' = \alpha$ in each case. Thus we can see that, as long as a positive rating leads to a higher arrival rate in the future, there will always be undercharging in period 1. Both, the arrival rate with one and with two positive ratings affect the price, however the effect of the immediate improvement α' is larger, as can be seen from the steeper curve along the α' axis, compared to the α'' axis.[16] We now know that both P_1 as well as P_2 are smaller than P_3. What we still have to confirm is that P_1 is smaller than P_2. To do this, I plot the difference between P_1 and P_2[17] in Figure 3, again for the same 3 initial arrival rates. It is clear that all differences are negative, thus $P_1 < P_2$. The difference decreases with increasing α which makes sense intuitively: If the initial arrival rate is very high already, relatively little can be gained from undercharging in each period and thus prices charged are similar.

Thus, I have shown that, in a three-period setting, sellers undercharge in every period and the extent of undercharging decreases in the second period, confirming my initial expectations.

3.2 Extension 1: Low-Quality Sellers and Bribing

Basic economic intuition would suggest that since a rating system adds a new utility - the utility of high reputation - sellers would be incentivized to share part of this utility with buyers in order to increase their total payoff. Practically speaking, sellers have an incentive to increase the rating they receive by paying the buyer either directly (monetary payment) or indirectly (additional service). Yet there is a surprising lack of literature on this specific subject. With the following extension I try to fill this gap.

So far I have assumed that the seller sells a high quality good and is therefore guaranteed a positive rating as long as a trade takes place. I now want to

[16] Note that even if one positive rating does not lead to a higher arrival rate ($\alpha' = \alpha$), an increase in the arrival rate after two positive ratings ($\alpha'' > \alpha$) still leads to some undercharging in period 1.

[17] We have seen that $P_2' > P_2$, so we can ignore P_2' for this exercise.

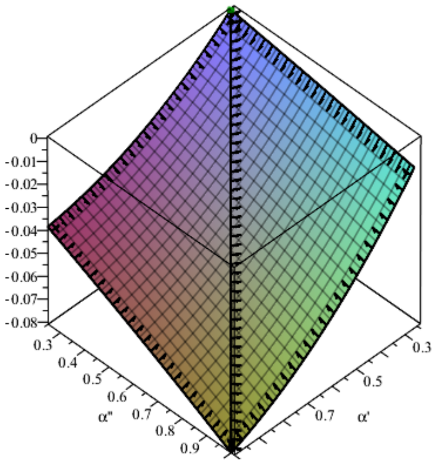

(a) $\alpha = 0.25$

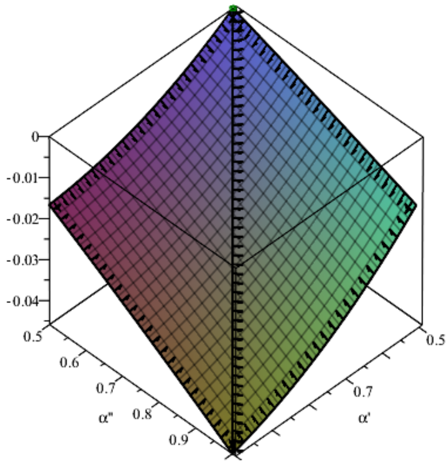

(b) $\alpha = 0.5$

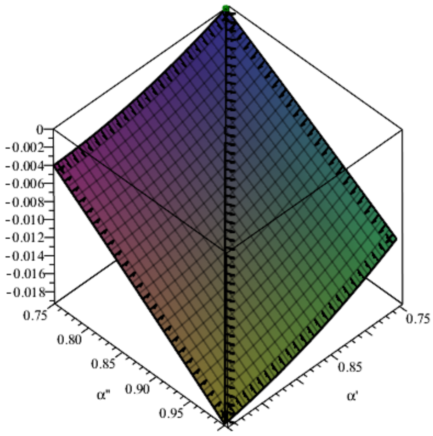

(c) $\alpha = 0.75$

Figure 3: $P_1 - P_2$ with initial arrival rate α fixed at different values

change this assumption and consider a low quality seller. Such a seller will receive no rating or a negative rating when a trade takes place. Re-entry into the market under a new identity is always possible without additional cost, so a negative rating can be treated equivalently to no rating being given (since a seller with a negative rating would clearly change his identity and start fresh).[18] As the continuation value doesn't change depending on the success

[18] An alternative interpretation of this assumption would be based on Resnick et al. (2006) who find that a small amount of negative ratings for a new seller does not affect the willingness-to-pay of buyers.

of the trade, the value function in period 1 reduces to:

$$V_{1|0}^{L} = \alpha \left[P_1 \int_{P_1}^{1} f(R) \, \mathrm{d}R \right] + V_{2|0}^{L} \tag{27}$$

and the single possible outcome in period 2 is described by:

$$V_{2|0}^{L} = \alpha \left[P_2 \int_{P_2}^{1} f(R) \, \mathrm{d}R \right] \tag{28}$$

It is easy to see that in this case $P_1 = P_2 = \frac{1}{2}$ which means that the existence of a rating system doesn't affect the seller and there is no undercharging in earlier periods. However, I now want to introduce a "Bribing" mechanism to the model: Once a trade is taking place, the seller can invest an additional amount of money to increase the rating he receives. This could be interpreted either as a direct bribe in terms of a monetary transfer to the buyer (the seller directly buys a better rating or offers a refund in return for a good rating) or an indirect bribe, for example in form of an additional service to the buyer (for example one could imagine a fruit basket in the hotel room or especially friendly and time consuming personal contact to the buyer). The goal of the model is to find out if and when such a bribe would be used by the seller and what this means for the pricing in period 1. First, I define the value function of a low quality seller in period 1 with a successful bribe B:

$$\begin{aligned} V_{1|0}^{b} = &\alpha \left[\int_{P_1^b}^{1} f(R) \, \mathrm{d}R \, [P_1^b + V_{2|1}^{b} - B] + \int_{0}^{P_1^b} f(R) \, \mathrm{d}R V_{2|0}^{L} \right] \\ &+ (1-\alpha) \, V_{2|0}^{L} \end{aligned} \tag{29}$$

The value functions in period 2 are equivalent to $V_{2|1}$ and $V_{2|0}$ from the baseline model (see equations (6) and (7)). Therefore the price and values in period 2 are still the same as well. The price in period 1 however, now changes to:

$$P_1^b = \frac{1 - \frac{1}{4}\alpha' + \frac{1}{4}\alpha + B}{2} \tag{30}$$

To evaluate this price we need to find out what the feasible values for B are. The highest bribe a seller would be willing to pay is the one that equalizes

$V^b_{1|0}$ and $V^L_{1|0}$, as any bribe higher than that would lead to a lower payoff than not bribing at all. This can be done mathematically or one can simplify the problem and just ask oneself: once a trade is taking place, how much would a seller be willing to pay to increase his rating from α to α'? The answer to this question is: the same amount that can be gained from a positive rating, which is the difference between $V_{2|0}$ and $V_{2|1}$. Therefore, we can say that:

$$B^{max} = V^b_{2|1} - V^L_{2|0} = \frac{1}{4}\alpha' - \frac{1}{4}\alpha = \frac{1}{4}(\alpha' - \alpha) \tag{31}$$

and

$$0 < B \leq \frac{1}{4}(\alpha' - \alpha) \tag{32}$$

I now plug in the maximum bribe to equation (29) to find:

$$P^b_1 = \frac{1 - \frac{1}{4}\alpha' + \frac{1}{4}\alpha + B}{2} = \frac{1 - \frac{1}{4}\alpha' + \frac{1}{4}\alpha + \frac{1}{4}\alpha' - \frac{1}{4}\alpha}{2} = \frac{1}{2} \tag{33}$$

This means that the price with a bribe will never exceed the price without a bribe, but still be higher than the price of a high quality seller:[19]

$$P_1 < P^b_1 \leq P^L_1 \tag{34}$$

With the presence of a bribing mechanism, a buyer in period 1 profits in two ways from a rating system: he pays a lower price for the good and, in case of meeting a low-quality seller, he receives the bribe transfer. Note that, even if the buyer is not satisfied with the product quality, there is no reason for him not to accept the bribe: The cost of the bad transaction is "sunk" but he can still profit by providing a positive rating.[20] The only drawback for a buyer offered such a deal is a potential moral conflict. Anecdotal evidence suggests that this rarely stops buyers from accepting the deal, however.[21] The loser in

[19]Notice that I disregard any kind of cost component, I only refer to prices relative to the optimal price without a rating system

[20]Compare this to the case of reciprocal ratings on eBay, where buyers would not give negative ratings for fear of retaliation. The same concept applies here.

[21]For example, Streitfeld (2012*b*) notes in his article for The New York Times concerning such a case that "310 out of 335 reviews [...] were five stars and nearly all the rest were four stars." and that when one customer accused the seller of scamming the review system, that customer was promptly chastised by another user, claiming this was a common practice and "not a scam but an incentive".

this system is the buyer in period 2: he is misled by the unjustified high rating and purchases a good he might not have purchased under complete information and doesn't receive the compensation the period-1-buyer received. As long as this interaction is repeated infinitely however, buyers as well as sellers are profiting from the rating system. Sellers share a part of their additional continuation value with the buyers and thus the rating system, while losing its function of distinguishing between high and low quality sellers, increases overall utility for the buyer. One problem that could arise in such a system over the long term is that the believability of the ratings will decrease once participants realize that the ratings are not truthful and therefore a high rating becomes meaningless, which would lead to the rating system losing its function of reducing adverse selection and a deterioration of overall quality.

3.3 Extension 2: Presence of Moral Hazard

I now describe seller behavior that would be relevant in a market where moral hazard is an issue. The game involves a decision for the seller to either follow up on his trade correctly (honest) or to defect (cheat). Should the seller cheat, he receives an additional payoff C,[22] in turn he also receives a negative rating. A negative rating will have a negative impact on the seller's arrival rate. However, obtaining a new identity is always possible without additional cost, so a seller with a net negative rating will just restart with a new identity. The game lasts for 3 periods, after the third period the seller exits the market.

Once again, there are diminishing marginal effects of ratings. I model this here by assuming that the second positive rating does not have an impact on the arrival rate anymore, i.e. buyers do not (or cannot) distinguish between a seller with 1 or 2 positive ratings. Figure 4 illustrates the possible arrival rates a seller can attain. Sellers start with α and decide to cheat or to be honest. If they cheat, they receive a negative rating and restart with a new

[22]For simplicity, the model is set up in a way such that when the seller decides on cheating, he will always receive the payoff C, even if he does not meet a trading partner. In a practical sense, this would mean that he is saving in fixed costs (say by producing a lower quality product), rather than saving in variable costs (for example by saving on packaging or not delivering at all).

identity in period 2, thus keeping α. If they are honest they improve to α'. In period 2, an honest seller can be honest again to keep α' or decide to cheat and fall back to α. A cheating player can again decide to improve to α' or to restart and stay at α. In the last period it is clear that everyone will cheat as there is no future period to consider. Additionally, there is always the possibility that no trade takes place, either due to the seller not finding a buyer or due to the price being too high. In such cases, the arrival rate does not change.[23]

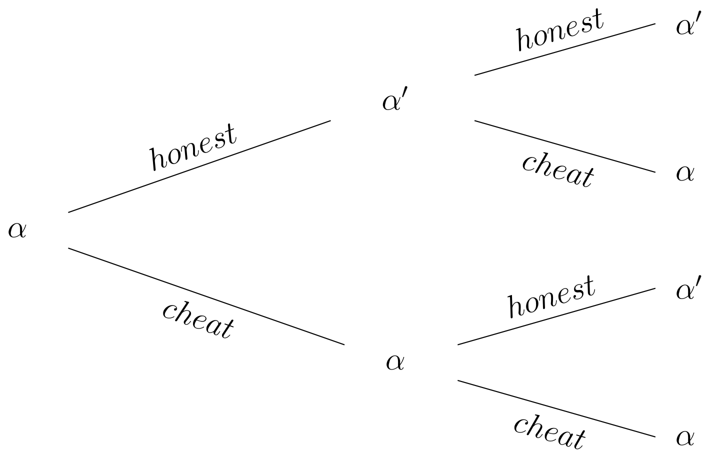

Figure 4: attainable arrival rates

In period 1, everyone starts with a score of 0. The choice is to either cheat, resulting in $V^C_{1|0}$ or to be honest, receiving $V^H_{1|0}$. For now, I assume that a seller with a positive score will be honest in period 2 and a seller with a neutral score will cheat in period 2. I will show later on, that this assumption holds in most scenarios.

$$V^H_{1|0} = \alpha \left[\int_{P^H_{1|0}}^{1} f(R)\, \mathrm{d}R\, [P^H_{1|0} + V^H_{2|1}] + \int_0^{P^H_{1|0}} f(R)\, \mathrm{d}R V^C_{2|0} \right] + (1-\alpha)\, V^C_{2|0} \tag{35}$$

[23]Thus, the label "cheat" in Figure 4 also includes "no trade taking place".

$$V^C_{1|0} = \alpha \left[\int_{P^C_{1|0}}^1 f(R)\,\mathrm{d}R\,[P^C_{1|0} + V^C_{2|0}] + \int_0^{P^C_{1|0}} f(R)\,\mathrm{d}R V^C_{2|0} \right] + (1-\alpha)\,V^C_{2|0} + C \tag{36}$$

As in the case of cheating a seller will always receive continuation value $V^C_{2|0}$, the second equation simplifies to:

$$V^C_{1|0} = V^C_{2|0} + C + P^C_{1|0}\alpha \int_{P^C_{1|0}}^1 f(R)\,\mathrm{d}R \tag{37}$$

In period 2 there can be 2 states: either the seller received a positive rating in period 1, or he didn't. In each period he can choose to either cheat or be honest.

$$V^H_{2|1} = \alpha' \left[\int_{P^H_{2|1}}^1 f(R)\,\mathrm{d}R\,[P^H_{2|1} + V^C_{3|2}] + \int_0^{P^H_{2|1}} f(R)\,\mathrm{d}R V^C_{3|1} \right] + (1-\alpha')\,V^C_{3|1} \tag{38}$$

$$V^C_{2|1} = \alpha' \left[\int_{P^C_{2|1}}^1 f(R)\,\mathrm{d}R\,[P^C_{2|1} + V^C_{3|0}] + \int_0^{P^C_{2|1}} f(R)\,\mathrm{d}R V^C_{3|1} \right] + (1-\alpha')\,V^C_{3|1} + C \tag{39}$$

$$V^H_{2|0} = \alpha \left[\int_{P^H_{2|0}}^1 f(R)\,\mathrm{d}R\,[P^H_{2|0} + V^C_{3|1}] + \int_0^{P^H_{2|0}} f(R)\,\mathrm{d}R V^C_{3|0} \right] + (1-\alpha)\,V^C_{3|0} \tag{40}$$

$$V^C_{2|0} = \alpha \left[\int_{P^C_{2|0}}^1 f(R)\,\mathrm{d}R\,[P^C_{2|0} + V^C_{3|0}] + \int_0^{P^C_{2|0}} f(R)\,\mathrm{d}R V^C_{3|0} \right] + (1-\alpha)\,V^C_{3|0} + C \tag{41}$$

Again, the last equation simplifies to:

$$V^C_{2|0} = V^C_{3|0} + C + P^C_{2|0}\alpha \int_{P^C_{2|0}}^{1} f(R)\,\mathrm{d}R \tag{42}$$

Additionally, because by assumption $V^C_{3|2} = V^C_{3|1}$, (38) simplifies to:

$$V^H_{2|1} = V^C_{3|1} + P^H_{2|1}\alpha' \int_{P^H_{2|1}}^{1} f(R)\,\mathrm{d}R \tag{43}$$

Lastly, in period 3 the seller will always cheat and, as in the other models, he will always demand the same price P_3. There are 3 different states, however the outcome is the same in 2 of them because the buyers cannot differentiate between a net rating score of 1 and 2.

$$V^C_{3|2} = P_3\,\alpha' \int_{P_3}^{1} f(R)\,\mathrm{d}R + C \tag{44}$$

$$V^C_{3|1} = P_3\,\alpha' \int_{P_3}^{1} f(R)\,\mathrm{d}R + C \tag{45}$$

$$V^C_{3|0} = P_3\,\alpha \int_{P_3}^{1} f(R)\,\mathrm{d}R + C \tag{46}$$

As mentioned above, price P_3 is the same in all cases and it is easy to show that it is $\frac{1}{2}$. Thus the values in period 3 are:

$$V^C_{3|2} = \frac{1}{4}\alpha' + C \tag{47}$$

$$V^C_{3|1} = \frac{1}{4}\alpha' + C \tag{48}$$

$$V^C_{3|0} = \frac{1}{4}\alpha + C \tag{49}$$

Plugging in the result from equation (48) to equation (43) yields:

$$V^H_{2|1} = \frac{1}{4}\alpha' + C + P^H_{2|1}\alpha' \int_{P^H_{2|1}}^{1} f(R)\,\mathrm{d}R \tag{50}$$

The price $P^H_{2|1}$ again is $\frac{1}{2}$, thus this solves to:

$$V^H_{2|1} = \frac{1}{2}\alpha' + C \tag{51}$$

In the case of cheating, the optimal price is:[24]

$$P^C_{2|1} = \frac{1 - (\frac{1}{4}\alpha + C) + (\frac{1}{4}\alpha' + C)}{2} = \frac{1 - \frac{1}{4}\alpha + \frac{1}{4}\alpha'}{2} \tag{52}$$

Plugging in this price, as well as the continuation values from equations (48) and (49) into equation (39) yields the following value:

$$V^C_{2|1} = \frac{1}{2}\alpha' + \underbrace{\frac{1}{8}\alpha\alpha' + \frac{1}{64}\alpha'\alpha^2 + \frac{1}{64}\alpha'^3 - \frac{1}{8}\alpha'^2 - \frac{1}{32}\alpha\alpha'^2}_{\text{cost of cheating}} + \underbrace{2C}_{\text{gain of cheating}} \tag{53}$$

It can be shown that the middle part of the equation is always < 0 (as long as $\alpha' > \alpha$). Comparing this to the Honest version (equation (51)), it becomes clear that the middle term describes the cost of cheating, while, by definition, the extra C is the gain of cheating. Therefore the optimal decision depends on whether the middle term is, in absolute terms, smaller or greater than C. Next, I repeat this process for the case of having received no positive rating in period 1. For being honest, I get a price of $P^H_{2|0} = \frac{1-\frac{1}{4}\alpha'+\frac{1}{4}\alpha}{2}$. This yields the following value:

$$V^H_{2|0} = \frac{1}{2}\alpha + \underbrace{\frac{1}{8}\alpha\alpha' + \frac{1}{64}\alpha^3 + \frac{1}{64}\alpha\alpha'^2 - \frac{1}{8}\alpha^2 - \frac{1}{32}\alpha'\alpha^2}_{\text{gain of being honest}} + \underbrace{C}_{\text{(opportunity) cost of being honest}} \tag{54}$$

The parallels to equation (53) are obvious. This time, the middle term is always > 0, thus this is the gain of being honest. The cost of being honest is indicated by the lack of C, compared to equation (55).

[24] Notice that finding a trading partner actually reduces the continuation value in this case, thus it is ideal to overcharge.

The price in the case of cheating with 0 positive ratings is, once again, $P^C_{2|0} = \frac{1}{2}$. The value is, accordingly:

$$V^C_{2|0} = \frac{1}{2}\alpha + 2C \tag{55}$$

In the first period I find the following value function and price for being honest:

$$\begin{aligned} V^H_{1|0} = \alpha & \left[\int_{P^H_{1|0}}^{1} f(R)\, \mathrm{d}R \, [P^H_{1|0} + (\frac{1}{2}\alpha' + C)] + \int_{0}^{P^H_{1|0}} f(R)\, \mathrm{d}R (\frac{1}{2}\alpha + 2C) \right] \\ & + (1-\alpha)\, (\frac{1}{2}\alpha + 2C) \end{aligned} \tag{56}$$

$$P^H_{1|0} = \frac{1 - (\frac{1}{2}\alpha' + C) + \frac{1}{2}\alpha + 2C}{2} = \frac{1 - \frac{1}{2}\alpha' + \frac{1}{2}\alpha + C}{2} \tag{57}$$

Applying this price yields the following value for being honest in period 1:

$$\begin{aligned} V^H_{1|0} = & \frac{1}{16}\alpha^3 + \frac{1}{4}\alpha^2 C - \frac{1}{4}\alpha^2 - \frac{1}{8}\alpha'\alpha^2 - \frac{1}{2}\alpha C - \frac{1}{4}\alpha\alpha' C \\ & + \frac{3}{4}\alpha + \frac{1}{16}\alpha\alpha'^2 + \frac{1}{4}\alpha\alpha' + 2C \end{aligned} \tag{58}$$

The case of cheating is a lot simpler: Because we assume that a seller who cheats in period 1 will cheat in every period, he never cares about his continuation value and always sets his price as $\frac{1}{2}$. His value in period one is 3 times $\frac{1}{4}\alpha + C$ or:

$$V^C_{1|0} = \frac{3}{4}\alpha + 3C \tag{59}$$

The next step is now to equalize the payoff for cheating and being honest in period 1, in order to make sure that a seller in period 1 is indifferent between cheating and being honest. Solving for C, this yields the following cheat-payoff:

$$C = -\frac{\frac{1}{4}\alpha^3 - \frac{1}{2}\alpha'\alpha^2 + \frac{1}{4}\alpha\alpha'^2 - \alpha^2 + \alpha\alpha'}{\alpha^2 - \alpha\alpha' - 2\alpha - 4} \tag{60}$$

Note that this is a positive number under the given assumptions: The numerator is strictly positive, the denominator is strictly negative.

Now I show that, using this value for C, a seller with a positive score will be honest in period 2 and a seller with a neutral score will cheat in period 2. To do this, I first plot the difference between $V_{2|1}^H$ and $V_{2|1}^C$, which means equation (51) - equation (53). Applying the usual constraints, I receive the plot depicted in Figure 5.[25] Interestingly, the result is slightly negative for

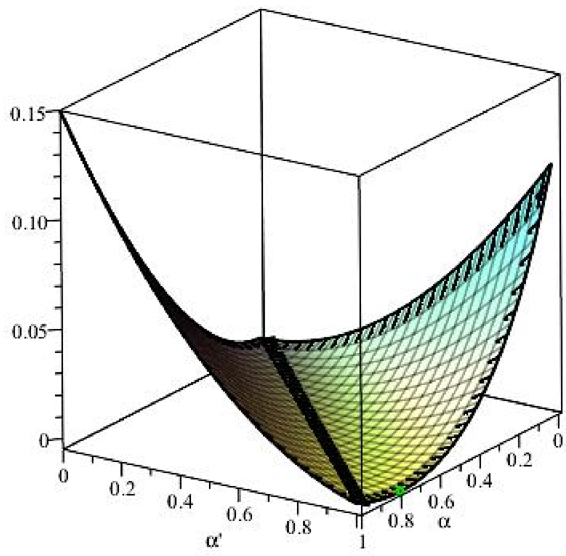

Figure 5: Difference between being honest and cheating in period 2: positive reputation

some values, specifically it is negative for relatively high values of α. This makes sense intuitively: if the probability of meeting a trading partner is very high even for newcomers, then the seller does not have much to lose and might thus be more likely to cheat. However, for most reasonable arrival rates[26] the result is positive, thus a seller with an existing reputation will most likely be honest again to keep his reputation. However, it is still interesting to note that under some circumstances there are incentives for a previously honest

[25]Note that again only the right half of the plot is relevant

[26]The result is actually marginally negative for some reasonable arrival rates as well, specifically for $\alpha \geq 0.5$. However, the differences are so small that this could be considered to be indifference. To give an example: for $\alpha = 0.5$ and $\alpha' = 0.75$, the difference would be -0.0032. For a comparison, C would be 0.026 in this scenario. I would argue that a seller who decided to be honest when indifferent before is likely to be honest again when faced with such a decision.

seller to cheat in the second period. Turning to the case of not having a positive rating, either because the seller cheated and started fresh or because he didn't complete a trade at all, I take the difference between $V^H_{2|0}$ and $V^C_{2|0}$ or equation (54) - equation (55) and receive the plot depicted in Figure 6. The maximum value of this plot is 0 (which only occurs for either $\alpha = \alpha'$

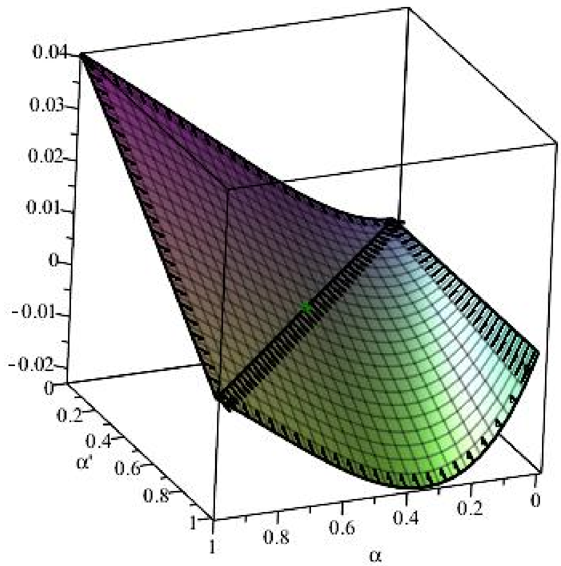

Figure 6: Difference between being honest and cheating in period 2: neutral reputation

or $\alpha = 0$, both of which are special cases that can be disregarded), all other values are negative, which means that if a seller has not managed to build up a reputation by period 2, he will always cheat for the rest of the game. As opposed to an honest seller, a cheating seller never has an incentive to change his behavior.

Therefore, I have shown that (under most reasonable assumptions) it is optimal for a cheating seller (and a seller who did not complete a trade, despite potentially good intentions) to continue cheating and for an honest seller to continue being honest. To illustrate this, I update the tree of attainable arrival rates from Figure 4 to a tree of feasible arrival rates in Figure 7.

As to the price setting, it is interesting to note that $P^C_{1|0} = P^H_{2|1} = P^C_{2|0} =$

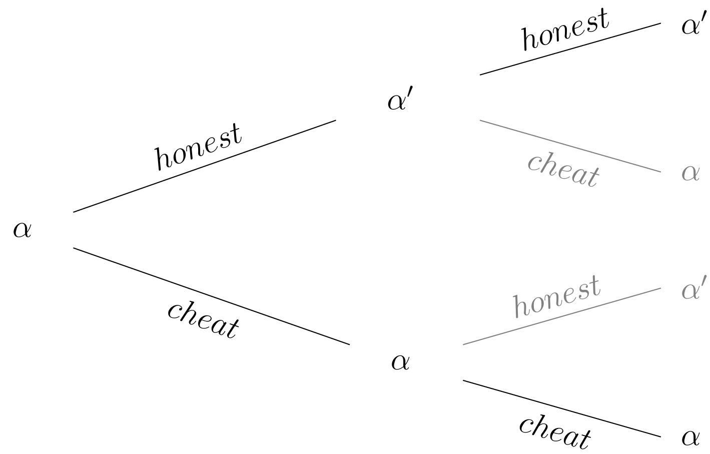

Figure 7: feasible arrival rates

$P_3 = \frac{1}{2}$.[27] Thus, the only time where undercharging occurs is in the case of an honest seller in the first period: $P^H_{1|0} = \frac{1-\frac{1}{2}\alpha'+\frac{1}{2}\alpha+C}{2}$.[28]

I illustrate this model with a numerical example. Assume the arrival rate of a seller with a neutral rating is $\alpha = 0.25$ and the arrival rate of a seller with at least one positive rating is $\alpha' = 0.75$. The reward for cheating is then $C = -\frac{\frac{1}{4}\alpha^3-\frac{1}{2}\alpha'\alpha^2+\frac{1}{4}\alpha\alpha'^2-\alpha^2+\alpha\alpha'}{\alpha^2-\alpha\alpha'-2\alpha-4} = 0.030$. Plugging in the values for α, α' and C in the respective equations, I summarize the payoffs in every state in Figure 8.

This model illustrates, how in the early stage of their career, sellers decide on either honestly trading or cheating and repeatedly creating new identities. The latter type of sellers will keep this behavior forever while the former faces another decision once they are established with a high reputation. They can keep their high reputation and continue behaving honestly. Or alternatively, they can cash in on their high reputation and start cheating. I show that honest sellers will most likely stay honest, the rating system thus successfully fulfils its role as a sanctioning device (at least for previously honest sellers). This result does not hold for all combinations of arrival rates however, in some

[27] This is because in each of these cases, completing a trade is not rewarded with a higher future payoff (i.e. the arrival rate does not change). Therefore, there is no incentive for additional investments into reputation.

[28] It can be shown that $C \leq (\alpha' - \alpha)$. Thus there is undercharging, but not as much undercharging as in the baseline model.

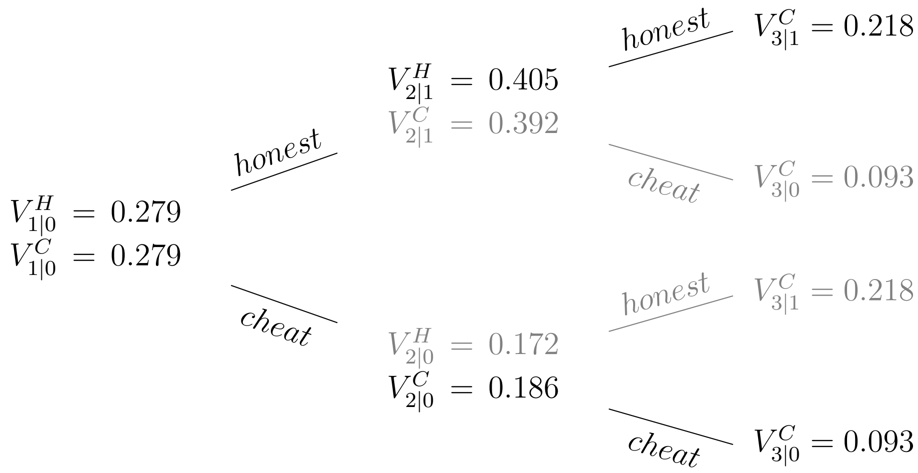

Figure 8: payoffs for $\alpha = 0.25$ and $\alpha' = 0.75$

markets sellers are indifferent or might even choose to cheat in the second period. Additionally I show that, unsurprisingly, sellers will always cheat at the end of their careers, a fact that has been shown empirically as well (see for example Cabral and Hortacsu (2010)). Lastly, it should be noted that this model is consistent from the buyer-side as well. A low reputation seller can be honest (period 1 honest seller) or cheat (any period cheating seller), but a high reputation seller is more likely to be honest (period 2 honest seller). So a buyer is more likely to choose a high reputation seller (thus the higher arrival rate), but choosing a low reputation seller is feasible as well (with a lower arrival rate).

4 Conclusion

Reviewing existing literature on online ratings, I found that basic features of ratings, for example their effect on (auction) prices and likelihood of sale have been fairly well researched empirically. Current rating systems do fulfill their most basic functions of reducing adverse selection (i.e. signaling high quality) and moral hazard (i.e. punishing bad behavior). However, there are several unique features of online markets that have to be considered when designing rating systems. For example, new identities are cheap to attain which reduces the extent to which misbehavior can be punished and the anonymous nature of the internet enables sellers to post fake reviews or encourage buyers to rate higher than they usually would, thus diluting the quality signal. That this is not only a theoretical concern but also an issue in practice has been shown by several examples. eBays previous reciprocal rating system has shown that both parties rate predominantly strategically. Sellers change their behavior shortly before exiting a market and several media reports have shown that some sellers do try to increase their reputation by posting fake reviews or influencing buyer ratings and that buyers in general seem to accept such bribes. Empirical evidence on this behavior is still scarce, but the literature is growing.

In the main part of this thesis I tried to model price setting behavior under an online rating system in general and then extended the model with two examples of seller behavior. My model confirms the common assumption of newcomers having to "pay their dues" by setting lower prices in earlier periods. In a first extension I then show that lower quality sellers have an incentive to pay an additional amount directly or indirectly to their buyers in order to increase their reputation in the second period. When designing rating systems it is important to understand that such incentives exist and are a very real issue. If the practice of "buying" a higher reputation becomes widespread, buyers lose faith in the signal quality and the rating system loses its function of mitigating adverse selection. In a second extension I model moral hazard, i.e. the possibility of sellers cheating on buyers. I show that under some circumstances it is possible that sellers are indifferent between cheating and not cheating in the first period but then usually keep their strategy in the second period. Thus, in my model, a high rating is a signal for a trustworthy buyer. This result is not very strong however. While a seller cheating in period 1 will always cheat in period 2, an honest seller is roughly indifferent between cheating and being honest in period 2. This is due to my assumption of decreasing marginal benefits of ratings: The seller can only lose his current reputation, but cannot gain additional reputation. So there might be an incentive to cheat, despite a high reputation. It is therefore important for the designer of the reputation system (especially when moral hazard is present), that buyers can differentiate between "high" and "very high" reputation. This has been implemented for example by eBay in the shape of a multicolored star-ranking, with additional potential ranking gains for even the highest reputation users.

All in all, the design of the ideal rating system remains an interdisciplinary challenge. It should set correct incentives for buyers to rate honestly and for sellers to behave correctly, while eliciting sufficient ratings and attracting enough users on the very competitive and continuously growing online market. There are many possibilities for further research, both theoretically as well as empirically, especially for markets in which adverse selection is a dominant issue. The economic importance of sellers influencing their own ratings in various ways needs to be further explored and ways to detect or prevent such practices might need to be incorporated into rating systems.

References

Akerlof, G. A. (1970), ‘The market for ’lemons’: Quality uncertainty and the market mechanism’, *The Quarterly Journal of Economics* **84**(3), 488–500.

Anderson, M. and Magruder, J. (2012), ‘Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database’, *Economic Journal* **122**(563), 957–989.

Andrews, T. and Benzing, C. (2007), ‘The determinants of price in internet auctions of used cars’, *Atlantic Economic Journal* **35**(1), 43–57.

Avery, C., Resnick, P. and Zeckhauser, R. (1999), ‘The market for evaluations’, *American Economic Review* **89**(3), 564–584.

Ba, S. and Pavlou, P. A. (2002), ‘Evidence of the effect of trust building technology in electronic markets: Price premiums and buyer behavior’, *MIS quarterly* pp. 243–268.

Bolton, G., Greiner, B. and Ockenfels, A. (2013), ‘Engineering trust: reciprocity in the production of reputation information’, *Management Science* **59**(2), 265–285.

Bruce, N., Haruvy, E. and Rao, R. (2004), ‘Seller rating, price, and default in online auctions’, *Journal of Interactive Marketing* **18**(4), 37–50.

Cabral, L. and Hortacsu, A. (2010), ‘The dynamics of seller reputation: Evidence from ebay’, *The Journal of Industrial Economics* **58**(1), 54–78.

Chevalier, J. A. and Mayzlin, D. (2006), ‘The effect of word of mouth on sales: Online book reviews’, *Journal of marketing research* **43**(3), 345–354.

Chwelos, P. and Dhar, T. (2007), ‘Differences in “truthiness” across online reputation mechanisms’, *Sauder School of Business* .

Dellarocas, C. (2006), ‘Reputation mechanisms’, *Handbook on Economics and Information Systems* pp. 629–660.

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

Dewally, M. and Ederington, L. (2006), ‘Reputation, certification, warranties, and information as remedies for seller-buyer information asymmetries: Lessons from the online comic book market’, *The Journal of Business* **79**(2), 693–729.

Dewan, S. and Hsu, V. (2004), ‘Adverse selection in electronic markets: Evidence from online stamp auctions’, *The Journal of Industrial Economics* **52**(4), 497–516.

Eaton, D. H. (2005), ‘Valuing information: Evidence from guitar auctions on ebay.’, *Journal of Applied Economics & Policy* **24**(1), 1 – 19.

Friedman, E. J. and Resnick, P. (2000), ‘The social cost of cheap pseudonyms’, *Journal of Economics and Management Strategy* **10**, 173–199.

Friedman, E., Resnick, P. and Sami, R. (2007), ‘Manipulation-resistant reputation systems’, *Algorithmic Game Theory* pp. 677–697.

Houser, D. and Wooders, J. (2006), ‘Reputation in auctions: Theory, and evidence from ebay’, *Journal of Economics & Management Strategy* **15**(2), 353–369.

Jian, L., MacKie-Mason, J. K. and Resnick, P. (2010), ‘I scratched yours: The prevalence of reciprocation in feedback provision on ebay’, *The BE Journal of Economic Analysis & Policy* **10**(1).

Jiang, B.-J. and Chen, P.-Y. (2007), ‘An economic analysis of online product reviews and ratings’, *SSRN eLibrary* .

Jin, G. Z. and Kato, A. (2006), ‘Price, quality, and reputation: Evidence from an online field experiment’, *The RAND Journal of Economics* **37**(4), 983–1005.

Jolivet, G., Jullien, B. and Postel-Vinay, F. (2013), ‘Reputation and prices on the e-market: Evidence from a major french platform’, *Institut d’Économie Industrielle (IDEI), Toulouse* .

Klein, T. J., Lambertz, C., Spagnolo, G. and Stahl, K. O. (2006), 'Last minute feedback', *SFB/TR 15 Discussion Paper No. 62* .

Klein, T. J., Lambertz, C. and Stahl, K. O. (2013), 'Market transparency, adverse selection, and moral hazard', *CESifo Working Paper Series* (4552).

Livingston, J. A. (2005), 'How valuable is a good reputation? a sample selection model of internet auctions', *Review of Economics and Statistics* **87**(3).

Luca, M. (2011), 'Reviews, reputation, and revenue: The case of yelp. com', *Harvard Business School* .

Luca, M. and Zervas, G. (2013), 'Fake it till you make it: Reputation, competition, and yelp review fraud', *Harvard Business School NOM Unit Working Paper* (14-006).

Lucking-Reiley, D., Bryan, D., Prasad, N. and Reeves, D. (2007), 'Pennies from ebay: The determinants of price in online auctions', *The Journal of Industrial Economics* **55**(2), 223–233.

Mayzlin, D., Dover, Y. and Chevalier, J. A. (2012), 'Promotional reviews: An empirical investigation of online review manipulation', *National Bureau of Economic Research Working Paper* (w18340).

McDonald, C. G. and Slawson, V. C. (2002), 'Reputation in an internet auction market', *Economic Inquiry* **40**(4), 633–650.

Melnik, M. I. and Alm, J. (2002), 'Does a seller's ecommerce reputation matter? evidence from ebay auctions', *The journal of industrial economics* **50**(3), 337–349.

Mickey, R. (2010), 'The impact of a seller's ebay reputation on price.', *American Economist* **55**(2).

Miller, N., Resnick, P. and Zeckhauser, R. (2002), Eliciting Honest Feedback in Electronic Markets, Working Paper Series rwp02-039, Harvard University, John F. Kennedy School of Government.

Miller, N., Resnick, P. and Zeckhauser, R. (2005), 'Eliciting informative feedback: The peer-prediction method', *Management Science* **51**(9), 1359–1373.

Przepiorka, W. (2013), 'Buyers pay for and sellers invest in a good reputation: More evidence from ebay', *The Journal of Socio-Economics* **42**(0), 31 – 42.

Resnick, P. and Zeckhauser, R. (2002), 'Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system', *Advances in applied microeconomics* **11**, 127–157.

Resnick, P., Zeckhauser, R., Swanson, J. and Lockwood, K. (2006), 'The value of reputation on ebay: A controlled experiment', *Experimental Economics* **9**(2), 79–101.

Segal, D. (2011), 'A rave, a pan, or just a fake?', *The New York Times, May 22, 2011* .
URL: *http://www.nytimes.com/2011/05/22/your-money/22haggler.html*

Shapiro, C. (1983), 'Premiums for high quality products as returns to reputations', *The Quarterly Journal of Economics* **98**(4), 659–79.

Smith, D. (2004), 'Amazon reviewers brought to book', *The Observer, February 15, 2004* .
URL: *http://www.theguardian.com/technology/2004/feb/15/books.booksnews*

Streitfeld, D. (2012*a*), 'The best book reviews money can buy', *The New York Times, August 26, 2012* .
URL: *http://www.nytimes.com/2012/08/26/business/book-reviewers-for-hire-meet-a-demand-for-online-raves.html?smid=pl-share*

Streitfeld, D. (2012*b*), 'For $2 a star, an online retailer gets 5-star product reviews', *The New York Times, January 27, 2012* .
URL: *http://www.nytimes.com/2012/01/27/technology/for-2-a-star-a-retailer-gets-5-star-reviews.html*

A Appendix

A.1 Summary of empirical evidence on the effectiveness of online rating systems

Study	Platform	Item	Observation Period	Methodology	Main Results
Anderson and Magruder (2012)	Yelp	Restaurants	July 2010 - February 2011	Regression discontinuity design	An extra half-star causes restaurants to sell out 49% more frequently
Andrews and Benzing (2007)	eBay	Used cars	Summer 2004	OLS, Logit regression	Reputation increases price and likelihood of sale
Ba and Pavlou (2002)	eBay	Music, Software, Electronics	January - March 2000	Multivariate Regression on Field Data and Field Experiment	Data: Positive ratings increase price, negative ratings show no significant impact; Experiment: Negative ratings have stronger impact than positive ratings
Bruce et al. (2004)	eBay	Laptops, PCs, DVDs, Books	May 2002	Maximum likelihood regression	Rating increases amount of bids and prices for cheap and used items, prior default is indicative of future default and terminal sellers are more likely to default
Cabral and Hortacsu (2010)	eBay	Notebooks, collectible coins, Beanie Babies	October 2002 - March 2003	Cross-section and panel regression	Negative ratings decrease sales and increase the rate of new negative ratings, sellers receive more negative ratings just before exiting

Chevalier and Mayzlin (2006)	Amazon, BarnesAndNoble	Books	May 2003	Regression on the difference between sales rank of the same book on each page	Higher rating on a site leads to higher sales on that site
Dewally and Ederington (2006)	eBay	Comic Books	January - June 2001	Tobit and Heckman's least squares	Both, amount of previous ratings as well as quality of those ratings increase price
Dewan and Hsu (2004)	eBay	Collectible stamps	July - December 2001	OLS, Tobit and Probit Regression	Economically modest but statistically significant positive effect of reputation on prices and probability of sale
Eaton (2005)	eBay	Electric guitars	January - April 2001	Logit Regression, OLS	Negative feedback reduces probability of sale but increases price
Houser and Wooders (2006)	eBay	Pentium Processors	September - December 1999	Two-Step GLS	Seller (but not buyer) reputation has statistically significant effect on prices
Jin and Kato (2006)	eBay	Collectible Baseball Cards	April - December 2001	Field experiment, Probit regression	High reputation is effective at identifying good-faith sellers but does not infer higher quality
Jolivet et al. (2013)	PriceMinister	Books, CDs, video games, DVDs	January 2001 - December 2008	OLS, GMM	Seller reputation has a statistically significant positive effect on prices

Livingston (2005)	eBay	Golf clubs	October 2000 - August 2001	Probit Regression	Price and probability of sale increase with reputation, marginal effects are severely decreasing with increasing reputation
Luca (2011)	Yelp	Restaurants	January 2003 - October 2009	Regression discontinuity desgin	A one star increase in the restaurants rating leads to a 5-9% increase in revenue
Luca and Zervas (2013)	Yelp	Restaurants	2004 - 2012	Fixed Effects	16% of reviews are identified as fake, review fraud is a response to economic incentives rather than a small number of unethical businesses
Lucking-Reiley et al. (2007)	eBay	Collectible US 1-cent-coins	July - August 1999	Censored Tobit maximum likelihood regression	Negative ratings have larger effect on prices than positive ratings
Mayzlin et al. (2012)	TripAdvisor, Expedia	Hotels	October 2011	Difference in Differences	There are more signs for manipulated reviews on TripAdvisor than on Expedia, independent hotels are more likely to engage in review manipulation
McDonald and Slawson (2002)	eBay	Limited Edition Barbie Doll	January - July 1998	Multivariate regression	Seller reputation increases both price and number of bids
Melnik and Alm (2002)	eBay	Mint condition 1999 US $5 coins	May - June 2000	Censored Tobit maximum likelihood regression	Quality of reputation has statistically significant positive effect on prices

Mickey (2010)	eBay	iPod Touch	April 2008	OLS	Price decreases with the share of negative ratings
Przepiorka (2013)	eBay	SD memory cards	October - December 2006	Logit and OLS regressions	Positive ratings increase and negative ratings decrease probability of sale and price in auctions as well as fixed price offers
Resnick and Zeckhauser (2002)	eBay	MP3 players, Beanie Babies	February - June 1999	Logit regression	Reputation doesn't affect prices but increases probability of sale, buyer and seller feedback are highly correlated
Resnick et al. (2006)	eBay	Vintage postcards.	January 2002	Controlled field experiment	Price increases by 8.1% if seller is reputable

Table 1: Empirical evidence on the effectiveness of online rating systems

Plagiatserklärung

Ich bezeuge mit meiner Unterschrift, dass meine Angaben über die bei der Abfassung meiner Arbeit benutzten Hilfsmittel sowie über die mir zuteil gewordene Hilfe in jeder Hinsicht der Wahrheit entsprechen und vollständig sind. Ich habe das Merkblatt zu Plagiat und Betrug vom 22. Februar 2011 gelesen und bin mir der Konsequenzen eines solchen Handelns bewusst.

Ort, Datum	Unterschrift