THREE ESSAYS ON ONLINE RATINGS AND AUCTION THEORY

by

Jonathan M. Lafky

B.S. Mathematics, University of Oregon, 2004

Submitted to the Graduate Faculty of

the Department of Economics in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2010

UNIVERSITY OF PITTSBURGH

ECONOMICS DEPARTMENT

This dissertation was presented

by

Jonathan M. Lafky

It was defended on

June 24th 2010

and approved by

John Duffy, Economics, University of Pittsburgh

Lise Vesterlund, Economics, University of Pittsburgh

Andreas Blume, Economics, University of Pittsburgh

Roberto Weber, Social and Decision Sciences, Carnegie Mellon University

Dissertation Director: John Duffy, Economics, University of Pittsburgh

THREE ESSAYS ON ONLINE RATINGS AND AUCTION THEORY

Jonathan M. Lafky, PhD

University of Pittsburgh, 2010

This work includes three papers, focusing on online ratings and auction theory. The first chapter, *Why Do People Write Reviews? Theory and Evidence on Online Ratings* examines what motivates consumers to provide uncompensated ratings for products on the internet. It finds that both a desire to inform other buyers and to punish or reward sellers are at play.

The second chapter *Optimal Availability of Online Ratings* looks at what role consumer ratings have in determining market outcomes. It shows that more information for buyers does not always lead to increased buyer welfare, and that in some cases it may be preferable to keep some buyers uninformed of existing ratings.

The third chapter *All Equilibria of the Multi-Unit Vickrey Auction* characterizes all Nash Equilibria for the Vickrey auction with three or more bidders and any number of units. It shows that all equilibria fall into one of two basic families, and that there cannot be equilibria of any other type.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

PREFACE

This work would not have been possible by the support and encouragement of many people. I greatly appreciate the expertise and assistance of my dissertation committee members: John Duffy, Lise Vesterlund, Andreas Blume and Roberto Weber. I am especially grateful to my committee chair and advisor, John Duffy, for his guidance, endless patience, humor, enthusiasm and encouragement.

In addition to my committee, I would like to gratefully acknowledge professors George Loewenstein, Jack Ochs and Ted Temzelides for perspectives and advice that have made me a better economist.

This work benefited greatly from the comments of Sourav Bhattacharya, Oliver Board, Tim Hister, Wooyoung Lim, Soiliou Namoro, and seminar and conference participants at the University of Pittsburgh, 2009 International and North American Economic Science Association meetings, and 2009 Western Economic Association.

Lastly, I am most grateful to my parents, John and Linda for their support and friendship and for giving me the best of all possible educations.

1.0 WHY DO PEOPLE WRITE REVIEWS? THEORY AND EVIDENCE ON ONLINE RATINGS.

1.1 INTRODUCTION

1.1.1 OVERVIEW

Internet commerce is a large and rapidly growing component of the economy. Internet retail is projected to be greater than \$156.1 billion for 2009, up 11% from \$141.3 billion in 2008 (Mulpuru, 2009). Typical growth over the past decade has been even higher, averaging approximately 20% annually.

The rapid growth and popularity of internet retail is not surprising. Virtually any good can be purchased on the internet, in every model, style, or color produced. The enormous selection offered to consumers means that they must often choose between several goods with similar observable characteristics but potentially different levels of quality. Without firsthand experience, it may be difficult or impossible for consumers to tell which of several similar-looking products is of the highest quality.

In an effort to alleviate this problem and to encourage sales, many internet retailers provide customer-based rating and review systems for their products. In these systems, consumers (sometimes restricted only to previous buyers) are allowed to leave written reviews as well as numerical scores for products. These ratings are then made available to future buyers to inform them of the product's qualities, allowing them to make more informed purchases.

The average review score can vary considerably for products that have otherwise similar characteristics, and may be the only insight consumers have into a product's unobservable qualities before they buy. As Chevalier and Mayzlin (2003) demonstrate, ratings can significantly influence buyers' behavior, and as a result have a substantial impact on the success or failure of a product. But why are ratings given in the first place? Are people taking time to give these ratings in order to help their anonymous fellow shoppers, or are they writing out of gratitude or anger that they feel towards online merchants? Are raters equally likely to evaluate all products, or do they speak up only if they have a strong opinion? This paper examines some possible motivations for the provision of ratings in a theoretical framework, and then isolates those motivations in an experimental setting.

To preview the results, I find evidence that consumers are motivated by concern for both buyers and sellers when they decide to rate products. Ratings given to affect buyers are relative to the average quality found in the market, while ratings given to affect sellers are relative to the what raters consider to be fair behavior. Making rating less attractive through the introduction of a small cost has a large effect on the volume and distribution of ratings. Ratings in the presence of a cost take on a U-shaped distribution, which can lead to average ratings that are not representative of true quality. A possible solution to this problem is to provide small discounts to consumers who provide ratings, thereby compensating for any inconveniences or opportunity costs associated with rating products.

The remainder of the paper is organized as follows. Section 1.1.2 surveys relevant past research. Section 1.1.3 provides motivating data and poses the basic questions to be addressed. Section 1.2 introduces a theoretical framework for analyzing rating behavior and isolating concern for sellers from concern for buyers. Section 1.3 lays out the experimental design and hypotheses. Section 2.3.1 presents results from the experiments while section 1.5 discusses implications of those findings. Section 1.6 concludes.

1.1.2 RELATED LITERATURE

There is a small but growing literature on online ratings, with most existing work focusing on how consumers are influenced by ratings and how well those ratings can predict market outcomes. The first paper to demonstrate that consumer-generated ratings significantly impact consumer behavior is Chevalier and Mayzlin (2003), which examines how sales ranks of books at amazon.com and bn.com vary based on customer ratings. They find that ratings have a significant influence on sales, with the addition of just a few user-generated ratings significantly improving the amazon.com sales rank of previously unrated books.

Duan et al. (2005) and Dellarocas et al. (2004) examine consumer-generated ratings on movie review websites, showing that ratings given by consumers are significantly better than expert reviews at predicting movies' box office success. These papers also suggest that user ratings can be seen as a gauge of underlying word-of-mouth communication. For the domain of movies, their findings suggest that ratings are predictors rather than drivers of success and failure.

A simple but elegant theoretical framework for examining reputation mechanisms is introduced in Bolton et al. (2004). They consider an interaction in which a buyer must pay a seller in advance, after which the seller may choose to fulfill their commitment, or to cheat the buyer by not sending a product. They examine the environment with and without feedback, and find that the presence of feedback significantly increases trustworthiness among sellers.

The paper perhaps most similar to the current project is Li and Hitt (2007), who consider possible distortions across time in ratings for newly released products on amazon.com. They find that average ratings immediately after release begin high, drop rapidly, and then gradually rise to an intermediate level. They attribute this phenomenon to "avid fans" who rush to buy a product immediately after its release, and a later backlash by typical consumers against unrealistically high initial ratings.

Chen et al. (2008) use social comparisons to encourage users of movieLens, a movie recommendation website, to rate more movies. By providing users with a brief summary of how their rating output compares to others, the authors are able to substantially increase the volume of ratings. They also find some evidence that a user's altruism, as determined in a post-experimental survey, predicts their likelihood of rating movies that have few existing ratings. This strongly suggests that at least some users are motivated by altruism when providing ratings.

It is important to distinguish the current line of research from several papers that have been written on two-sided reputation systems. Houser and Wooders (2005), for example, examine the impact of reputations in eBay auctions, in which buyers and sellers rate one another. As evidenced by eBay's change in 2008 to a one-sided rating system (buyers may rate sellers, sellers cannot rate buyers), two sided systems can introduce the undesirable possibility of strategic rating behavior. In contrast, consumers in the one-sided system considered in this paper need not worry about being punished or rewarded for their ratings, and can rate products based solely on their own opinions.

1.1.3 MOTIVATING DATA

To give a picture of real-world ratings, data was taken from the amazon.com website in November 2008. The distribution of ratings for more than 400 products were collected, encompassing more than 17,500 separate ratings. The products evaluated were from the "Home Improvement" section of the website, which includes products such as lawn mowers, flashlights and electric chainsaws. The Home Improvement section was chosen in an effort to find products which have relatively objective quality. Unlike previous research which has looked primarily at books and music, the data is restricted to products with more objective quality in an effort to simplify the task of interpreting ratings.

Figure 1 shows the distribution of individual ratings for products with different average ratings. For reference, customers at amazon.com can give ratings from one star to five stars, in one-star increments. Only one item (less than 0.25% of all products) had an average rating of less than two stars, and thus is not included.

Looking at the distributions one feature is particularly striking: Middling ratings, especially ratings of two stars and three stars are uncommon, even among products with *average* ratings of two or three stars. The reasons behind these distributions are not clear, however, as several distinct mechanisms could generate the same pattern. The simplest explanation would be that quality itself tends towards extremes, where products realize binary qualities of "success" and "failure," with little else in between. It is also possible, however, that the pattern comes not from the underlying distribution of quality, but from the motivations that drive consumers to provide ratings.

If people view the act of rating a product to be intrinsically burdensome but nonetheless

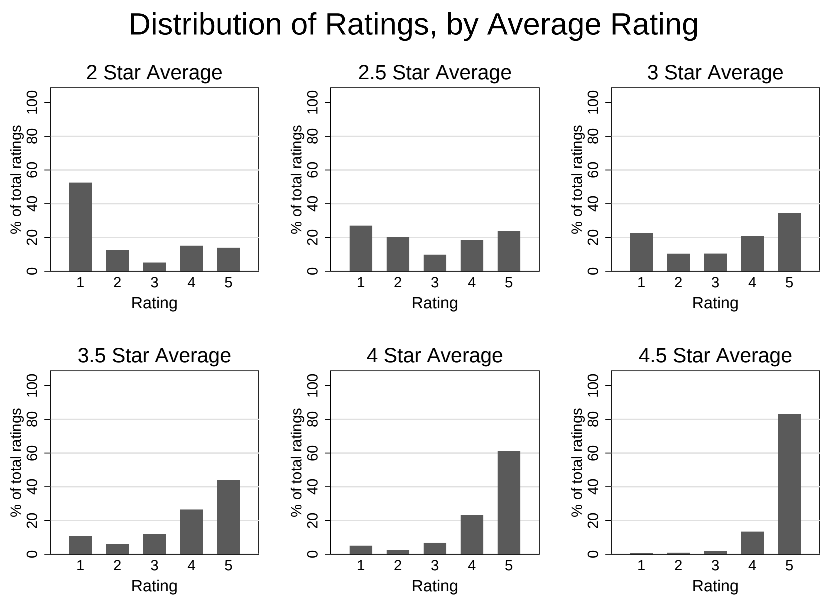

Figure 1: Ratings distributions by average rating.

want to help other buyers, the same pattern could emerge. Acceptable but unremarkable products would not be rated because the benefit to the rater from informing others would be smaller than the cost of providing the rating. High or low quality products would be rated because the rater could have a large impact on other buyers' welfare.

Alternatively, raters may take the time to rate in an attempt to punish or reward sellers for their quality. A buyer who receives a defective product may seek retribution against the good's seller by damaging their reputation with a negative rating. Likewise, a buyer who is pleased with a recently purchased good may give the seller a positive rating as a reward or encouragement for their high quality. In both cases the reaction elicited from the buyer is intense enough to outweigh any costs of rating. Products of moderate quality, however, would elicit neither reward nor punishment.

1.2 THEORY

This section develops a theoretical model to formalize the insights described above. It characterizes behavior for buyers and sellers interacting in a simple stylized market and generates predictions that can be tested in the laboratory.

1.2.1 THE MODEL

In this model there are n buyers, $B_1, B_2 \ldots B_n$, each deciding which of 2 sellers, S_1 and S_2 to buy from.[1] Buyer B_1 will be referred to as the "first buyer" and $B_2, \ldots B_n$ will be known as "second buyers." At the start of the game, each seller chooses a quality level, $q_i \in [0, q_{max}]$.[2] The choice of q_i is a seller's private information. After sellers choose their qualities, buyer B_1 selects one of the sellers. Since B_1 has no information to distinguish one seller from the other, for ease of notation assume that B_1 chooses S_1. B_1 learns S_1's quality, q_1 and is then given the opportunity to pay a cost of c in order to provide a rating $r \in [\underline{r}, \overline{r}] \subset \mathbb{R}$ for S_1.[3] After B_1 makes his rating decision, all other buyers learn what rating, if any, B_1 gave. If B_1 did not give a rating, the other buyers cannot tell which seller B_1 selected. Finally, $B_2 \ldots B_n$ simultaneously each select one of the sellers.

Sellers' payoffs are $U_{S_i}(q_i, n_i) = n_i \cdot (\overline{u} - aq_i)$ where $\overline{u}$ is the utility from setting $q = 0$, a is the marginal cost of quality, and n_i is the number of buyers who selected S_i. B_1's payoffs are given by $U_{B_1}(q) = bq - I_r c$, where b is the marginal benefit of quality and I_r is an indicator function for whether B_1 rated or not. The payoffs for all other buyers are simply $U_{B_i}(q) = bq$.

Given this framework the unique equilibrium is for sellers to set the minimum quality of $q = 0$, and for B_1 to never provide a rating so long as $c > 0$. This model does not reflect observed behavior, however, in that there are tens of millions of buyer-generated ratings on the internet. In order to explain this discrepancy I extend the model to include regard for

[1] Two is the minimum number of sellers that prevents unrealistic signalling behavior. With a single seller, not rating has the potential to convey as much information as rating.

[2] For simplicity prices are normalized to zero. This is done to remove the possibility that prices would be used as a signal for quality, which would complicate the task of inferring rater's motivations. One interpretation of this model is an analysis of products at a given price, meaning that quality can be thought of as value for money.

[3] This can be interpreted as both the opportunity cost of rating as well as any effort a consumer expends from the act of rating.

others.[4] The timing of the game remains the same, however B_1's payoffs are rewritten as

$$U_{B_1}(q_1) = bq_1 + \alpha \cdot (q_1 - R)U_{S_1}(q_1, n_1) + \beta \sum_{i=2}^{n} U_{B_i}(q_{j(i)}) - I_r c \tag{1.1}$$

α and β measure B_1's concern for sellers and buyers, respectively, $R \in \mathbb{R}^+$ is a quality level that B_1 considers to be fair treatment, and $q_{j(i)}$ is the quality corresponding to whichever seller j is selected by buyer i. If $q_1 < R$, S_1 was selfish in choosing quality, we have $q_1 - R < 0$ and B_1 has spiteful concern for his seller. When $q_1 > R$, S_1 was generous in choosing quality, and B_1 has altruistic concern for his seller. When $q_1 = R$, then $q_1 - R = 0$ and B_1 is unconcerned by S_1's utility. Note that the further S_1's choice is from B_1's opinion of fair quality, the more intense becomes B_1's altruism or spite towards his seller.

Sellers' preferences remain the same, and for simplicity do not include other-regarding components. Introducing other-regarding preferences for sellers would alter the quality levels provided, but would not qualitatively affect the analysis of first or second buyer behavior. Buyer behavior is characterized below for the full range of possible qualities, and thus all variations in quality levels are already accounted for.

The addition of other-regarding preferences to the model gives a reasonable starting point for describing behavior, although it is not ideal for an experimental analysis. As was the case with the motivating data, concern for sellers and second buyers still cannot be separately identified from the first buyer's actions. Isolating these motivations requires that one additional feature be added to the model.

Prior to the beginning of the game nature randomly determines if it will be a *buyer-fixed* or *seller-fixed* game. The type of game is known only to B_1, although sellers and second buyers know that it will be either buyer-fixed or seller-fixed with equal probability. A buyer-fixed game has the same structure as the previously described model, except that $B_2, \ldots, B_n$ receive fixed payoffs of $f \in \mathbb{R}$, independent of the actions taken by any player. Similarly, in a seller-fixed game all buyers receive their normal payoffs, while the sellers receive payoffs of f, independent of any player's action. In this way it is possible to "deactivate" either sellers or second buyers from B_1's decision to rate, as B_1 is affected only by his concern for sellers

[4] An alternate explanation for voluntary rating, especially in the realm of non-durable goods, is a repeat-customer motive. A consumer may rate to improve their own future interactions with a merchant. Such a motivation may drive some online ratings, although it is outside the scope of this paper.

or second buyers in each role's respective game type. This means that in a buyer-fixed game B_1's rating decision is determined entirely by his value of α and the quality he receives. Likewise in a seller-fixed game B_1's rating decision is affected only by his value of β and his received quality. Note that, because they cannot tell which type of game is being played, all players other than B_1 behave the same in both types of games.

1.2.2 SECOND BUYER BEHAVIOR

Before characterizing B_1's rating behavior it is first necessary to understand how second buyers respond to B_1's ratings. Second buyers condition their decision upon it being a seller-fixed game, as their decision in a buyer-fixed game is irrelevant to their payoffs. If B_1 rates in a seller-fixed game his preferences are perfectly aligned with those of second buyers.[5] Second buyers know that B_1's preferences are in line with their own and can rely upon the first buyer to provide ratings in their best interest. Thus they will buy from the rated seller when they observe a high rating and switch to the unrated seller when they observe a low rating. Notice that because later buyers' decisions are essentially binary (either choose the rated seller or choose an unrated seller) there is no need for more than two ratings: high (buy) and low (don't buy). The exact choice of messages is irrelevant, so long as first and second buyers share a common convention. It may be convenient to think of the "buy" message as $\overline{r}$ and "don't buy" as $\underline{r}$. Although only two messages are necessary, this paper allows, both theoretically and later experimentally, for a wider range of possible ratings for consistency with commonly used online rating systems.

1.2.3 FIRST BUYER BEHAVIOR

To understand B_1's behavior in the buyer-fixed game, two cases must be considered: $q < R$ and $q \geq R$. If $q < R$, meaning that the observed quality is less than B_1's threshold to trigger spitefulness, then B_1 will provide a negative rating if:

[5]Note that this implicitly assumes that B_1 does not exhibit spite towards second buyers. Allowing for the possibility of spiteful raters complicates the analysis significantly, as ratings can no longer be trusted by second buyers.

$$-c > \alpha \cdot (q - R)\frac{(n-1)}{2}(\overline{u} - aq)$$

Solving for q gives the cutoff value

$$\underline{q}_S = \frac{\alpha \cdot (n-1)(aR + \overline{u}) - \sqrt{\alpha \cdot (n-1)(-8ac + (n-1)(a^2R^2\alpha - 2aR\alpha\overline{u} + \alpha\overline{u}^2))}}{2\alpha \cdot (n-1)} \tag{1.2}$$

If S_1 is generous with quality, and chooses $q \geq R$, B_1 provides a positive rating if:

$$\alpha \cdot (q - R)(n-1)(\overline{u} - aq) - c > \alpha \cdot (q - R)\frac{(n-1)}{2}(\overline{u} - aq)$$

Solving for q gives the cutoff value

$$\overline{q}_S = \frac{\alpha \cdot (n-1)(aR + \overline{u}) - \sqrt{\alpha \cdot (n-1)(8ac + (n-1)(a^2R^2\alpha - 2aR\alpha\overline{u} + \alpha\overline{u}^2))}}{2\alpha \cdot (n-1)} \tag{1.3}$$

B_1 thus rates if:

$$q \in [0, \underline{q}_S) \cup [\overline{q}_S, q_{max}]$$

Where B_1 gives a negative rating in the first interval and a positive rating in the second interval. This leads to the following proposition:

Proposition 1. *In the buyer-fixed game, the first buyer employs a double-cutoff strategy for rating. He gives a negative rating if $q \in [0, \underline{q}_S)$, no rating if $q \in [\underline{q}_S, \overline{q}_S)$, and a positive rating if $q \in [\overline{q}_S, q_{max}]$.*

Figure 2 shows how the cutoff values vary with α, with parameters $n = 3, c = .25, a = .36, \overline{u} = 6$ and $R = 5.5$.

Note also that $\frac{d\overline{q}_S}{dc} > 0$ and $\frac{d\underline{q}_S}{dc} < 0$, meaning that the range of values that will be unrated by first buyers is increasing in the cost of rating. This is a key insight in explaining the U-shaped distribution of ratings. If $c = 0$, all quality levels will be rated, while if $c > 0$, a "blind spot" of unrated qualities emerges centered around R.

Behavior in the seller-fixed game is similar, with one important difference. Because each second buyer can potentially select any of the sellers, B_1's utility in a seller-fixed round can

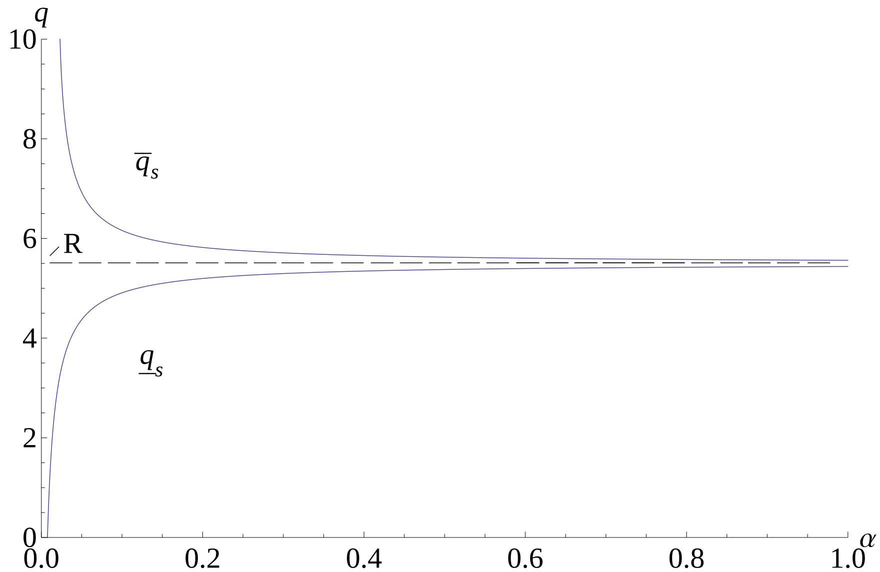

Figure 2: Cutoffs for buyer-fixed rating, as a function of alpha.

be influenced not only by S_1's quality, but also by the quality of sellers who have not yet been selected. B_1 thus needs to have beliefs about the quality buyers will receive if they switch from S_1 to another seller. Denote the average quality that B_1 believes to be offered by other sellers by q'. Note that no assumptions are made about the source of q', allowing for the possibility that it corresponds to the actual quality of the other seller, but not requiring it to do so.

To understand B_1's decision in a seller-fixed round we have to consider the cases $q < q'$, when the other seller is expected to be better, and $q \geq q'$, when the other seller is expected to be (weakly) worse. B_1 will provide a negative rating when $q < q'$ if:

$$\beta(n-1)bq' - c > \beta(n-1)\frac{bq + bq'}{2}$$

Which gives a cutoff value of

$$\underline{q}_B = q' - \frac{2c}{\beta b(n-1)} \tag{1.4}$$

B_1 will rate positively when $q \geq q'$ if:

$$\beta(n-1)bq - c > \beta(n-1)\frac{bq + bq'}{2}$$

Which gives a cutoff value of

$$\overline{q}_B = q' + \frac{2c}{\beta b(n-1)} \tag{1.5}$$

Proposition 2. *In the seller-fixed game, the first buyer employs a double-cutoff strategy for rating. He gives a negative rating if* $q \in [0, \underline{q}_B)$*, no rating if* $q \in [\underline{q}_B, \overline{q}_B)$*, and a positive rating if* $q \in [\overline{q}_B, q_{max}]$.

Figure 3 shows the cutoff values for different values of β, with parameters $n = 3$, $c = .25$, $b = .92$ and $q' = 3.5$.

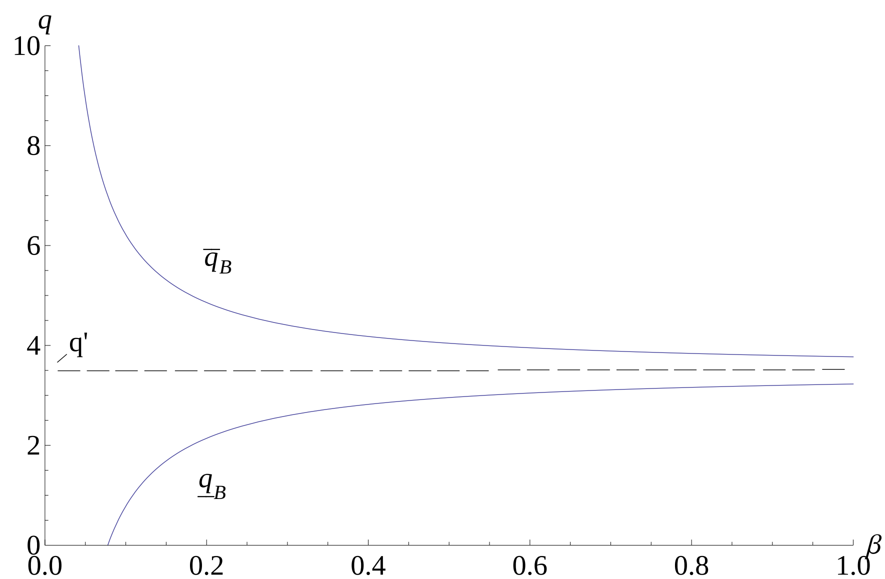

Figure 3: Cutoff seller-fixed rating, as a function of beta.

As in the buyer-fixed game, $\frac{d\overline{q}_B}{dc} > 0$ and $\frac{d\underline{q}_B}{dc} < 0$, meaning that the range of unrated qualities is again increasing in the cost of rating.

1.2.4 SELLER BEHAVIOR

Similar to later buyers, sellers need to condition their behavior only on it being a buyer-fixed game, and thus have three potentially payoff maximizing actions: $q = 0, q = \underline{q}_S$ and $q = \overline{q}_S$. Any other quality levels are dominated by one of these three. To simplify analysis, several parameters will be held constant to allow for an investigation of the variables of interest. Setting $a = .36, R = 5, \overline{u} = 6$ and $n = 3$ permits a simple analysis of c and α. Behavior is characterized for two values of c: $c = 0$ when rating is free and $c = .25$, when rating is costly.[6]

When $c = 0$, there is no cost to rating and all qualities are rated, meaning that $\overline{q}_S = \underline{q}_S$. Assuming that B_1 rates when he is indifferent (e.g. when $\alpha = 0$), sellers are always rated when $c = 0$, regardless of B_1's level of α. In this case $q = \overline{q}_S$ gives strictly higher payoffs than $q = 0$ and the unique equilibrium is both sellers offering high quality by setting $q = \overline{q}_S$.

If $c = .25$ and $\alpha = 0$, B_1 never rates and the unique optimal action for both sellers is setting $q = 0$. Comparing behavior when $\alpha = 0$ at $c = 0$ and $c = .25$ gives an important insight. When rating is costly sellers can get away with providing $q = 0$, knowing that they will not be punished for selfish behavior. When rating is free, however, sellers know that they will be rated poorly for low quality and respond by setting high quality of $q = \overline{q}_S$. An increase in the cost of rating thus leads to a decrease in the quality offered by sellers.

If $c = .25$ and $\alpha > 0$ the choice of quality level by both sellers can be described by the following symmetric game matrix:

	0	$\underline{q}_S$	$\overline{q}_S$
0	$U_S(0, \frac{3}{2}), U_S(0, \frac{3}{2})$	$U_S(0, 1), U_S(\underline{q}_S, 2)$	$U_S(0, \frac{1}{2}), U_S(\overline{q}_S, \frac{5}{2})$
$\underline{q}_S$	$U_S(\underline{q}_S, 2), U_S(0, 1)$	$U_S(\underline{q}_S, \frac{3}{2}), U_S(\underline{q}_S, \frac{3}{2})$	$U_S(\underline{q}_S, 1), U_S(\overline{q}_S, 2)$
$\overline{q}_S$	$U_S(\overline{q}_S, \frac{5}{2}), U_S(0, \frac{1}{2})$	$U_S(\overline{q}_S, 2), U_S(\underline{q}_S, 1)$	$U_S(\overline{q}_S, \frac{3}{2}), U_S(\overline{q}_S, \frac{3}{2})$

Figure 4: Game matrix for sellers.

[6] This parameterization is examined as it is used below in the laboratory experiment. Other parameterizations yield similar predictions.

The equilibria in this game depend on the value of α faced by the sellers.[7] For a sufficiently unconcerned first buyer, ($\alpha < .038$), both sellers providing low quality, $(\underline{q}_S, \underline{q}_S)$ is an equilibrium. Alternatively, for a sufficiently concerned first buyer ($\alpha > .0287$), both sellers providing high quality, $(\overline{q}_S, \overline{q}_S)$ is an equilibrium. Offering zero quality, $(0, 0)$ is never an equilibrium of this game since $\alpha = 0$.

Thus when $c = .25$ there are four regions of behavior, depending on the value of α. When $\alpha = 0$ the unique equilibrium is (0,0). For $0 < \alpha < .0287$, the unique equilibrium is $(\underline{q}_S, \underline{q}_S)$. When $.0287 < \alpha < .038$ either equilibrium is possible, and when $\alpha > .038$ the unique equilibrium is $(\overline{q}_S, \overline{q}_S)$.

1.3 EXPERIMENTAL DESIGN

Despite having theoretical predictions for behavior, it is difficult to test these predications against real online ratings. In analyzing data from ratings websites it is not possible to control for product quality or cost of rating, two variables essential to identifying behavior. These problems can be overcome by moving to the laboratory, where it is possible to perfectly control for both quality and the cost of rating.

This paper uses a novel experimental design intended to isolate subjects' motivations for giving ratings. The experiment was conducted using Fischbacher's (2007) z-Tree software over networked computers in the Pittsburgh Experimental Economics Laboratory. A total of 200 subjects were recruited from the student populations of the University of Pittsburgh and Carnegie Mellon University. Each session consisted of 20 subjects with no prior knowledge of the experiment. Each session began with the distribution of written instructions which were then read aloud to all subjects. A brief comprehension quiz was administered, subjects played 20 rounds of the experiment, and then completed a brief questionnaire. The experimental materials are included at the end of the paper. The instructions used in the Costly and Free treatments were identical, with a single additional sentence added to the Costly treatment.

[7] The equilibria depend on α since sellers' payoffs depend on the cutoffs $\overline{q}$ and $\underline{q}$, which are in turn determined by α.

Sessions lasted one hour or less and average earnings were approximately \$11.00, including a \$5.00 show-up fee. At the beginning of each round subjects were randomly assigned into four groups of five players. Within each group subjects were randomly assigned roles, with two subjects taking the role of sellers, one subject in the role of first buyer and two subjects in the role of second buyers. Sellers moved first by choosing an integer quality level from 0 to 10, inclusive. Without knowing the qualities selected, the first buyer then chose one of the sellers to "purchase" from. After making his choice, the first buyer learned the quality of his seller and was given the option to provide that seller with a rating. The cost of giving a rating varied by treatment, and was either \$0.25 (Cost treatment) or \$0.00 (Free treatment). A rating consisted of an integer score from 1 to 5, inclusive. This score was shown to sellers in buyer-fixed rounds and to second buyers in both types of rounds. Ratings were not made visible to sellers in seller-fixed rounds to exclude the possibility that first buyers would rate negatively to express their displeasure to sellers, as demonstrated in Xiao and Houser (2008).[8]

Ratings *did not* persist between rounds. When subjects were randomly assigned to new groups at the beginning of each round any ratings they received in previous rounds were not visible to the new group. This is essential to understanding the experiment, as it means that ratings were not accumulated throughout the course of each session, but existed only during the round in which they were given. Additionally, because roles were switched between rounds the incentive to rate to influence a future partner were minimized.

After the first buyer decided whether to give a rating, the second buyers were informed of what, if any rating was given. If no rating was given the second buyers could not tell which seller the first buyer picked. After seeing what, if any, rating was given the second buyers each simultaneously selected a seller for themselves. Sellers with quality level q received payoffs of $\$6.10 - \$0.34q$ each time a buyer picked them. First and second buyers who chose a buyer with quality q received payoffs of $\$0.92q$. To ensure that first buyers in the Costly treatment would never lose money, each subject was also given a \$1.00 "round completion fee" at the end of each round.

[8]Xiao and Houser show that responders in an ultimatum game accept lower offers when they are provided with the ability to send payoff-irrelevant messages after the proposers have made their offers. This suggests that subjects may simply wish to express their displeasure, even if it is not relevant to their earnings.

Each round was selected with equal probability by the computer to be either seller-fixed or (second) buyer-fixed. In a seller-fixed round all sellers received \$6.00, regardless of what decisions were made. Likewise, in a second buyer-fixed round all second buyers received \$6.00 total, independent of all subjects' decisions.

All subjects knew that each round would be either seller-fixed or second buyer-fixed, but only the first buyer knew the round's type while they made their decisions. Sellers and second buyers learned the round type only at the *end* of the round, after their decisions had been made. Sellers and second buyers were faced with the same decision and incentives in each type of round, even though their actions would only affect their payoffs 50% of the time. By implementing this payoff and information structure I have effectively "deactivated" either sellers or second buyers as targets for the first buyer's concern. For example, in a seller-fixed round the first buyer cannot influence his seller's payoffs in any way, since the seller will only receive the fixed payment of \$6.00.

1.3.1 EXPERIMENTAL HYPOTHESES

The first and most straightforward prediction to be tested is that a higher cost of rating will decrease the number of ratings, regardless of round type.

Hypothesis 1 (Ratings Volume)**.** *The frequency of rating will be significantly higher when rating is free than when it is costly.*

Based on the theoretical predictions that $\frac{d\overline{q}_B}{dc}, \frac{d\overline{q}_S}{dc} > 0$ and $\frac{d\underline{q}_B}{dc}, \frac{d\underline{q}_S}{dc} < 0$, first buyers faced with a cost of rating should be more inclined to provide high and low ratings than moderate ones.

Hypothesis 2 (Polarization of ratings)**.** *Ratings will be more polarized in the Costly treatment than in the Free treatment. High and low quality sellers will receive a larger percentage of all ratings when rating is costly than when it is free.*

If first buyers provide ratings in buyer-fixed rounds their actions must be an attempt to affect sellers in some way. This may be done either as a reward or a punishment for sellers, leading to the next two hypotheses.

Hypothesis 3 (Altruism toward sellers)**.** *In buyer-fixed rounds first buyers will give high quality sellers positive ratings, even when it is costly to do so.*

Hypothesis 4 (Spite toward sellers)**.** *In buyer-fixed rounds first buyers will give low quality sellers negative ratings, even when it is costly to do so.*

Similarly, if first buyers rate sellers in seller-fixed rounds they must be attempting to affect second buyers. In this case a rating serves as an informative signal to second buyers, and can be viewed as an altruistic act. [9]

Hypothesis 5 (Altruism toward buyers)**.** *In seller-fixed rounds first buyers will provide truthful ratings in order to aid other buyers, even when it is costly to do so.*

1.4 RESULTS

Table 1 lists summary statistics for the experiment. Non-parametric tests show that quality, ratings and the probability of rating are significantly higher in the Free treatment than in the Costly one ($p < .01$, Mann-Whitney U-test).

Table 1: Summary statistics. Standard deviations in parentheses.

	Free Buyer-Fixed	**Free Seller-Fixed**	**Costly Buyer-Fixed**	**Costly Seller-Fixed**
Rating	3.24	3.28	2.23	2.55
	(1.23)	(1.15)	(1.60)	(1.59)
Prob. Rate	.88	.90	.37	.35
	(.28)	(.24)	(.38)	(.39)
Quality	5.35	5.22	3.29	3.42
	(2.07)	(2.03)	(2.45)	(2.57)

[9] It is possible that first buyers could provide intentionally misleading ratings specifically to harm second buyers. Indeed, there are a handful (< 1%) of observations in the data that appear to be spiteful behavior toward second buyers.

As indicated in Figure 5 below, there is no significant difference in the accuracy of ratings between buyer-fixed and seller-fixed rounds within each of the Costly and Free treatments. This means that, contingent upon giving a rating, subjects provide the same rating for the same quality in both round types. However, there is a significant difference in ratings between the Costly and Free treatments. Notice in Figure 5 that there is a substantial jump in average ratings as quality increases from the 3-4 bin to the 5-6 bin in the Costly treatment, but a much smaller difference in the Free treatment. In the costly treatment first buyers are essentially giving binary recommendations, whereas in the free treatment we can see a broader range of ratings.

Figure 5: Average rating per quality bin, by cost and round type.

Behavior is even more interesting when we examine the frequency of rating, rather than the ratings themselves. Pooling both types of rounds in Figure 6 shows that the probability of rating in the Free treatment (88.9%) is more than twice that of the Costly treatment (35.9%), and the difference is significant at the 1-percent level (Mann-Whitney).

Finding 1. *The volume of ratings is significantly higher in the Free treatment than in the Costly treatment.*

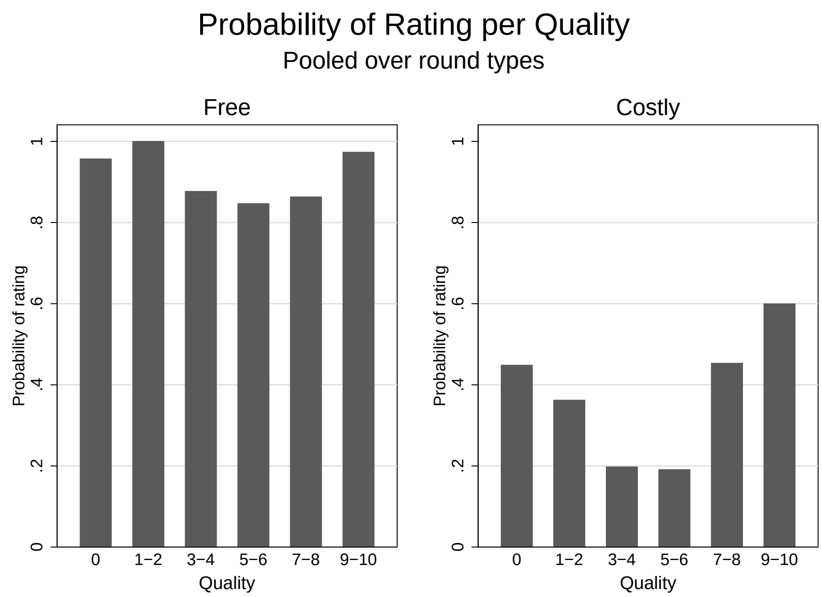

Figure 6: Probability first buyer rated, by quality and cost.

While it is not surprising that fewer ratings are given in the Costly treatment, the magnitude of the difference is striking. Removing a cost of only \$0.25, or approximately 2.3% of the average subject payment, leads to a 248% increase in the volume of ratings. Additionally, the decrease in ratings is not uniform across qualities ($p < .01$, Kruskal-Wallis test). Regressing rating choice on quality and quality^2 shows that there is a positive and significant correlation with the quadratic term in both treatments, although much more so in the costly case. The quadratic coefficient is .018 ($p < .01$) when rating is costly but only .005 ($p < .05$) when it is free. Introducing a cost of rating clearly causes subjects to decrease the probability with which they rate middling qualities, relative to extreme ones.

Finding 2. *The ratio of ratings for extreme qualities to ratings for moderate qualities is significantly greater in the Costly treatment than in the Free treatment.*

Finding 2 gives support to the polarization hypothesis, and provides a first glimpse into what may be causing the U-shaped distributions observed in online rating data. It shows that, in the face of a small cost of rating, people are more willing to rate when they have either a very positive or very negative experience relative to a more moderate one.

The next question is whether subjects are more likely to rate in either buyer-fixed or seller-fixed rounds. Figure 7 shows the probabilities of rating different qualities for each round type.

Figure 7: Probability first buyer rated, by quality, cost and round type.

Ratings are given for high and low quality sellers in all instances. This gives support to hypotheses 3-5, showing that raters are motivated by both buyers and sellers. Since ratings are given for both high and low quality sellers in each round type, this means that raters are driven to rate by altruism toward buyers and sellers, as well as spite toward sellers. This finding is a simple but powerful insight into the workings of rating systems, shedding some light on the basic motives that drive people to give ratings.

Notice that there is no significant difference between buyer-fixed and seller-fixed rounds in the Free treatment. In both round types the rating probabilities are approximately 90%. Probabilities vary slightly by quality, though none of the differences are statistically significant.

Looking at the Costly treatment shows a different picture. As already noted, rating probabilities are much lower across the board than in the Free treatment, and probabilities in the Costly treatment exhibit a U-shaped distribution. The U-shape is more pronounced in buyer-fixed rounds than in seller-fixed ones. The minimum of the distribution is also different between rounds, with the 5-to-6 bin being least likely to be rated in buyer-fixed rounds and the 3-to-4 bin least likely in seller-fixed rounds. This behavior can be explained theoretically by values of $q' = 3.36$ and $R = 5.14$, as discussed below.

Finding 3. *When ratings are costly, the least frequently rated quality is lower in seller-fixed rounds than buyer-fixed rounds.*

This difference is important, as it shows support for the prediction that seller-centric and buyer-centric ratings are made relative to different reference points. The theory predicts that ratings in a buyer-fixed round should be based on deviations from what first buyers perceive to be fair treatment, while in seller-fixed rounds they should be based on deviations from what the first buyer believes to be the quality of the untried seller.

Choosing a quality of 5 gives the smallest difference between seller and buyer payoffs and thus provides the most equitable division of earnings. Post-experimental questionnaires also showed the mean quality subjects believed to be "fair" was 5.14, due to the equitable payoffs generated for buyers and sellers. Finding 3 thus supports the theoretical prediction that in seller-fixed rounds, first buyers are rating based on the seller's deviation from a fair quality level.

The average quality offered by sellers in the Costly treatment is 3.36, which lies squarely in the 3-4 bin. If first buyers' beliefs about sellers' quality equaled the observed average quality, the theory predicts that first buyers would be least likely to rate qualities in the 3-4 bin, exactly as observed in the data. Finding 3 thus supports the theoretical prediction that first buyers are rating relative to their expectation of sellers' quality.

The data on the probability of rating also describes the level of concern first buyers show for sellers and second buyers. For example, if a first buyer rates a seller with quality $q < R$ in a second buyer-fixed round, we can infer that $q < \underline{q}_S$. Finding the α corresponding to that cutoff then gives a lower bound on the first buyer's α.

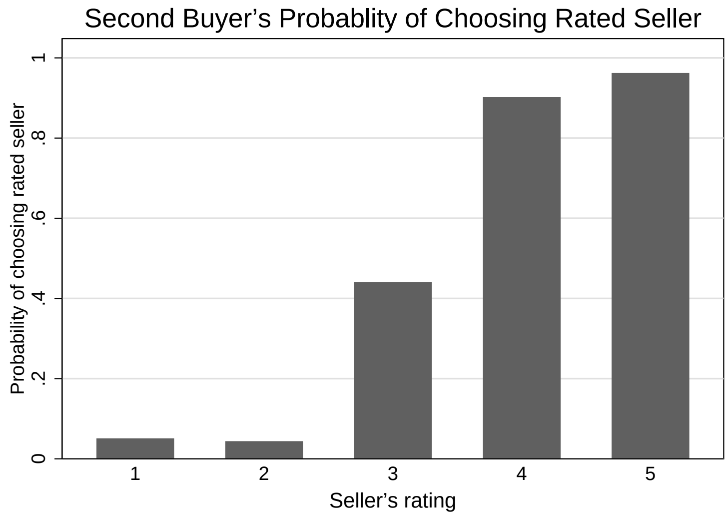

Figure 8: Probability second buyer chooses the rated seller, by rating

Using the corresponding cutoff values from the theory on costly second buyer-fixed rounds, 4% of subjects exhibit behavior implying an α of at least .425, meaning that they rate sellers regardless of their quality. An additional 42.3% have an α between .425 and .008, indicating that their decision to rate depends upon the quality they receive. A majority of first buyers, 53.7%, have an α of less than .008, showing very little concern for sellers. These first buyers never rate, regardless of quality. In costly seller-fixed rounds, behavior implies that 6.7% of first buyers always rate, having a β of at least .425. A majority of first buyers (66%), are sensitive to the quality they observe, having β between .425 and .041. Just over a quarter (27.3%) never rate in seller-fixed rounds, implying a β less than .041. All of these numbers are summarized in Table 2.

In addition to examining why ratings are given, it is also important to know the effect of ratings on market outcomes. Since ratings are assumed to be provided in an attempt to affect sellers and second buyers, it is important to check what impact ratings have on behavior for each of those roles.

Figure 11 shows that second buyers are heavily affected by the recommendations of

Table 2: Types of Raters.

	Never Rate	Sometimes Rate	Always Rate
Buyer-Fixed	4%	42.3%	53.7%
Seller-Fixed	6.7%	66%	27.3%

first buyers. More than 95.2% stay away from sellers with ratings of 1 or 2, and a similar percentage (93.4%) choose sellers with ratings of 4 or 5. Slightly less than half (44%) of second buyers choose a seller with a rating of 3, suggesting that buyers are essentially indifferent when faced with a middling rating. There is no significant difference in behavior between the Costly and Free treatments.

What impact does the cost of rating have on seller behavior? Figure 9 shows the distribution of qualities offered by sellers in the Costly and Free treatments.

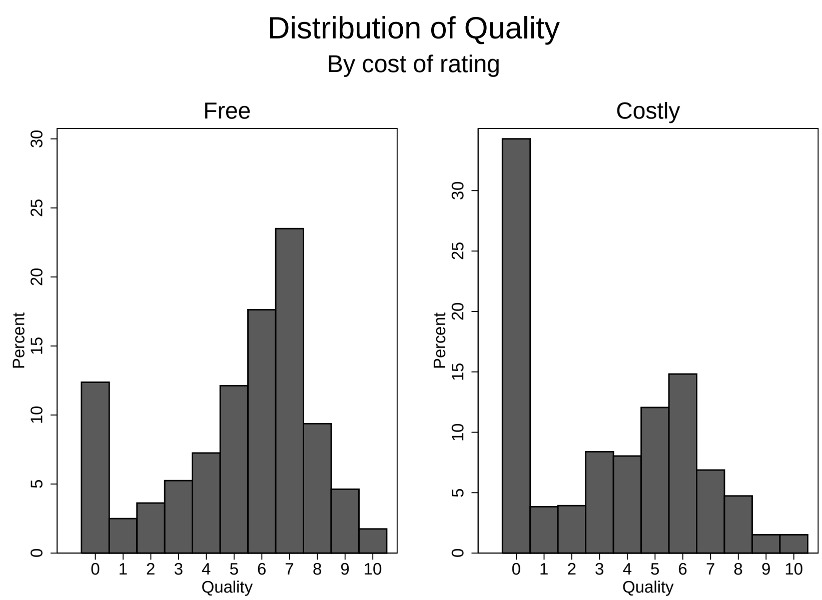

Figure 9: Distribution of qualities offered in the Free and Costly treatments.

Finding 4. *Sellers offer significantly higher quality levels in the Free treatment than in the*

Costly treatment.

Finding 4 indicates that the cost of rating is a significant factor in a seller's decision of what quality to provide to their buyers. There is a large and significant difference in the quality levels offered by sellers between the Free and Costly treatments. In the Free treatment the average quality level is 5.21, 65% higher than the average of 3.16 for the Costly treatment ($p < 0.01$, Mann-Whitney). This difference is especially striking when comparing the distributions of quality in each of the treatments. While 34.3% of sellers offer a quality of 0 in the Costly treatment, about one third as many (12.4%) do so in the Free treatment.

This difference can be explained by sellers correctly anticipating the frequency of rating in the Costly and Free treatments. When rating is free, sellers anticipate that they are relatively likely to be rated when they offer low qualities and offer higher qualities to avoid a negative rating. When rating is costly they know that it is relatively more likely that they will be able to offer very low qualities and escape without a rating. Higher cost of rating thus results in lower quality being offered by sellers.

It is important to understand how the cost of rating affects the welfare of buyers and sellers. Figure 10 shows the adjusted average earnings for each type of subject in free and costly rounds.[10] Sellers earnings decrease by 11.8% in the Free treatment, while first and second buyers' earnings increase by 49.1% and 39.3%, respectively ($p < .01$ for each difference, Mann-Whitney). As would be expected, sellers' earnings are lower in the more heavily rated Free environment, while first and second buyers' earnings are increased. Even if we exclude the direct benefit of not having to pay for ratings, both types of buyers are significantly better off when ratings are free.

[10]Earnings reported are based on amounts subject would receive if neither side of the market was fixed. Including the actual fixed payments skews the average earnings of sellers and second buyers towards the fixed earnings of $6.00.

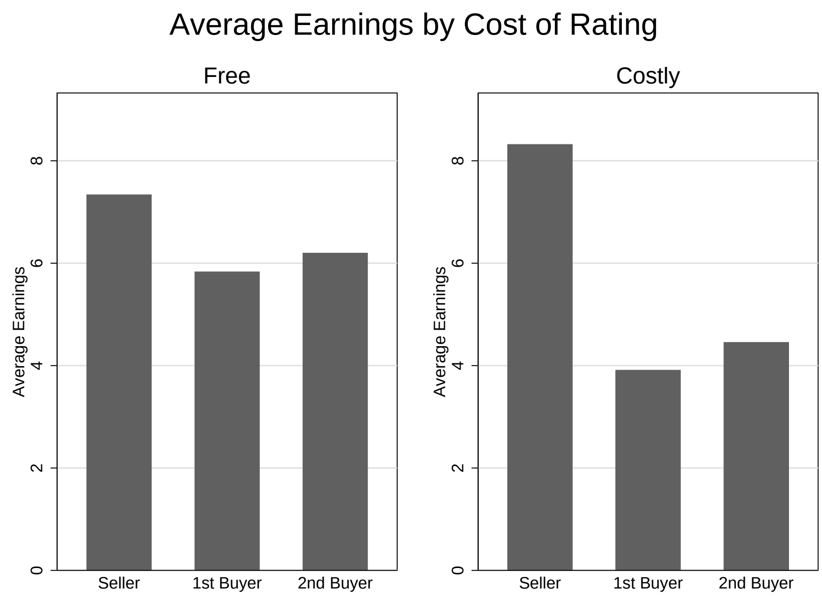

Figure 10: Average earnings in the Free and Costly treatments.

1.5 DISCUSSION

What can we learn from these findings, and how can they be applied? First, understanding why people write reviews may help to improve the design of future rating systems. Encouraging customers to rate products, especially those which have not yet been rated, is in the best interests of merchants and consumers.

Second, different systems may affect ratings' ability to accurately reflect product quality. For example, consider a product which is of acceptable quality, but has some small probability of failure. In the face of even a small cost, consumers who receive a functional, though unremarkable product would be unlikely to provide a rating online. However, the small number of consumers who do have a negative experience would be very likely to provide negative ratings for the product. This can lead the product to have inaccurately negative ratings.

As a simple example, consider a product which is generally of moderate quality, but occasionally fails utterly. If 80% of consumers receive a product of quality 5 and 20% receive quality 0, the average rating in the presence of cost will be 3.07, compared with 3.6 when rating is free. Designers of rating systems should pay special attention to minimizing any costs which might discourage consumers from providing ratings. Given that the very act of providing a rating may be burdensome, designers may want to provide small incentives to buyers for rating products. A small discount off future purchases, for example, could be all that is necessary to offset the cost of rating[11].

Designers should also be mindful of how they frame their requests for users to give ratings. They may receive different ratings if they focus consumers' attention on the seller or on future buyers. Requests that emphasize helping other buyers can be expected to produce ratings driven more by comparisons with other products or sellers, whereas requests that focus on the sellers will have a greater focus on fairness. It is not clear if a rating based on perceived fairness or relative quality is more desirable in general, as there are likely scenarios in which each bias is preferred.

1.6 CONCLUSION

This paper examines the factors that influence consumers' decisions to rate products online. Using a laboratory experiment I show that consumers are motivated to rate both by a concern for punishing or rewarding sellers and by a desire to inform future buyers. Introducing even a small cost of rating has a large effect on rating behavior, leading to fewer and more polarized ratings. One implication of this finding is that any cost of rating, even a small, and implicit one, may cause a "blind spot" for moderate quality products. This can cause inaccurate average ratings for products of variable quality, especially those whose quality distributions are asymmetric.

There is some evidence that raters rate relative to different reference points when motivated by sellers than when they are motivated by buyers. When concern for sellers is driving

[11]Note that this approach has the potential to create new incentive problems, as consumers may be motivated solely by the reward, which could decrease the accuracy of ratings.

ratings, raters evaluate sellers relative to what they perceive to be fair behavior. That is, they punish sellers offering quality too far below fair and reward sellers whose quality is significantly greater than fair.

When concern for other buyers is the motivation, raters evaluate their sellers relative to the alternative they expect to be available from other sellers. If their seller's quality is substantially below the alternative, they warn other buyers away with a negative rating. If the seller's quality is substantially above the alternative, they signal the high quality of their seller with a positive rating.

Sellers are responsive to buyers' cost of providing ratings, and adjust their quality accordingly. A small decrease in the cost of rating causes a large and statistically significant increase in the level of quality offered by sellers. This contributes to existing evidence, such as Bolton et al. (2004) that suggests consumer generated ratings systems may significantly increase consumer welfare. This finding further demonstrates that the cost of rating should be a major concern for designers of rating systems.

This paper suggests several directions for future research. One important feature of online ratings that has been intentionally removed from this setting is the accumulation of ratings over time. Since only one person can be the first rater for each product, most ratings are given in the shadow of many previous buyers' opinions. Given a series of pre-existing ratings by other raters, do consumers provide their honest opinion of a product, or do they attempt to adjust the mean rating toward what they feel is the correct value?

It would also be valuable to see portions of this experiment replicated in a more naturalistic environment. While the laboratory gives us unrivaled control over experimental conditions, it would be useful to document behavior described in this paper "in the wild." In particular, it would be interesting to see how the distribution of ratings for a real product varies with the cost of rating.

2.0 OPTIMAL AVAILABILITY OF ONLINE RATINGS

2.1 INTRODUCTION

How do online ratings affect market outcomes? It may seem that so long as ratings are accurate, greater access to ratings and more informed buyers will enhance market efficiency. Buyers will know which sellers are of high quality, and the best sellers will be the most popular. However, this paper shows that providing buyers with more ratings information does not necessarily improve market outcomes. In fact, increasing the number of informed buyers can actually lead to worse sellers dominating the market, and lower buyer welfare.

The basic intuition behind this result is that uninformed buyers carry out an essential role experimenting with new sellers. Even though it is in each buyer's interest to be as informed as possible, social welfare is maximized when some buyers are forced to buy without knowing what experiences previous buyers have had in the market. This leads to a tradeoff between informing buyers, which allows them to choose the best known seller, or not informing them, leading to a richer set of options for later buyers.

Previous literature has shown that online ratings are used extensively by consumers, and that they significantly influence product sales (Chevalier and Mayzlin, 2006). While it is has been established that ratings influence market behavior, it is not clear how much they improve it.

Existing work in the literature on herding and information cascades has studied how consumers aggregate information from those who have moved before them to make more informed decisions. A common thread throughout this extensive literature is the focus on the tension for buyers between following private and public information. Individuals can rationally ignore their own valuable private information in the face of sufficiently many other people taking actions that suggest that information is incorrect. The canonical model of information cascades that first captured these insights was simultaneously introduced in Banerjee (1992) and Bikhchandani et al. (1992).

The classic information cascade model describes agents with private, noisy signals of the state of the world. Each agent in turn makes a decision based on their signal. An agent's action is visible to all subsequent agents, although their signal itself is never observed. It is possible for a string of misleading signals to cause a sequence of suboptimal decisions to be made. Later agents observe that all previous agents have chosen the same action, and even if their own information suggests that they take a different action, they rationally ignore their signal and follow the behavior of those who came before.

There are two important factors in the cascade model. First, each agent has some private information that is never publicly observed. Second, and more crucially, later agents are influenced by previous ones because they know that earlier agents likely had access to more information than they do. In other words, they are trusting that the collective decision making of previous agents contains better information than their own single, noisy signal.

The information cascade model is used to explain inefficient herding, in which groups of people crowd into a single suboptimal choice, despite a better option being available. This paper presents a simple complimentary explanation for suboptimal herding. I show that when there is *no* private information among agents similar inefficient herding can occur. For some contexts this is a significant a step towards a more realistic model. More generally, it is a simplification of the traditional information cascade model while still yielding similar results.

There are also several lines of research outside of the information cascade literature that are relevant to the current paper. The literature on optimal experimentation examines the tradeoff between experimentation and exploitation when choosing between a known safe option, and an unknown option. The most relevant paper is Bolton and Harris (1999) who study the classic "n-armed bandit" problem, extended to many agents. In their paper each agent may choose the amount and timing of their experimentation, and all information they acquire is visible to all others. This leads to a strategic interaction in which there is an incentive for each agent to free-ride on other's experimentation. They find that agents exhibit a small amount of initial experimentation that gradually increases before rapidly dropping to zero, as the optimal action is discovered.

Bergemann and Välimäki (2000) examine the issue of consumers performing searches to determine whether a company newly entered into the market is of higher quality than an incumbent of known quality. They find excess experimentation among consumers, driven by the firms ability to change prices in response to consumers observed outcomes, thereby extracting some of the benefit from the consumers experiments. The firms ability to benefit from consumers experimentation modifies the quantity and efficiency of consumer search, distorting the view of its effectiveness. This is an especially important result, as many papers examining consumer experimentation ignore prices.

Another work showing the potential for ratings to have undesirable effects on markets is Satterthwaite (1979). He explores perverse price effects from reputations and the supply of doctors in local markets, finding that an increase in the supply of doctors may lead to an *increase* in the cost of medical care. This counterintuitive result stems from the reduction in consumer information sharing about doctor quality due to fewer patients frequenting each doctor. This is a product of the fixed number of consumers and increased number of doctors. Since there are fewer patients per doctor, the probability of a new patient meeting one of the doctors current patients shrinks as the number of doctors grows.

King (1995) addresses free-riding in a model with publicly observable search. He finds equilibrium behavior characterized by a single individual conducting extensive search while all others follow along after the search is completed. He finds there to be serious inefficiency in both the collection and distribution of information among consumers. Like Bergemann and Välimäki, King's model focuses on the role of prices in consumer search.

This paper's main contribution is showing that market outcomes vary non-monotonically in the amount of information available to consumers. This paper also illustrates the potential for suboptimal outcomes with consumer search in the absence of many of the assumptions made in previous works. The environment has been simplified in several ways. First, I assume that each agent has no private knowledge about the state of nature. Second, I assume that the outcome resulting from a given action is revealed as soon as an agent has experienced it. This means that the usual information-cascade story no longer holds, as there are no inferences to be made about other's information or behavior. Third, I make no assumptions on the number of sellers in the market, nor do I require a known outside option or "incumbent" seller.

The remainder of the paper is organized as follows. Section 2.2 address the model with deterministic quality. Sections 2.2.1 and 2.2.2 examine the cases with perfect, and imperfect usage of rating, respectively. Section 2.3 examines the model when quality varies stochastically. Section 2.4 looks at the potential for suppressing some buyers' access to ratings to improve market outcomes. Section 2.5 concludes.

2.2 DETERMINISTIC QUALITY

The first and simplest environment I study assumes that each of a number of sellers offers an ex-ante identical experience good. A sequence of buyers select sellers, each trying to select a seller of the highest possible quality. Each buyer rates the seller she interacts with, informing future buyers of that seller's quality.

I first consider a model with perfect rating, in which buyers always learn the ratings that have been given by other consumers. I then examine a model with imperfect rating, when buyers are only probabilistically informed.

2.2.1 PERFECTLY INFORMED BUYERS

There are m sellers, called $S_1, \ldots, S_m$ and n buyers, $B_1, \ldots, B_n$. Each seller i has a quality q_i that is an i.i.d. draw from the uniform distribution over the unit interval $[0, 1]$, and is not known to buyers. Beginning with B_1, buyers move sequentially, each choosing a seller and receiving a payoff equal to that seller's quality. After choosing a seller the buyer "rates" that seller and reveals their quality to all other buyers. Subsequent buyers thus observe the outcomes of all earlier buyers' purchases prior to selecting a seller of their own.

As a simple example, consider the case with three buyers and three sellers. Without loss of generality, assume that B_1 chooses S_1. If $q_1 > 1/2$ then both the second and third buyers will also select S_1. B_2 will select S_1 since $q_1 > E(q_2)$ and $q_1 > E(q_3)$. B_2 will not reveal any new information, since she chose the same seller as B_1, and thus B_3 will face the same decision as B_2 and also select S_1. Notice that while all buyers will choose the same seller if that seller's quality is greater than 1/2, the expectation of the highest quality among S_2 and S_3 is 2/3. Thus if $1/2 < q_1 < 2/3$, the buyers most likely chose seller of suboptimal quality.

This simple example can easily be extended to any number of buyers and sellers. Regardless of how large the market is, once a seller with $q > 1/2$ is found all subsequent buyers will rationally choose that seller rather than trying a new seller with expected quality equal to 1/2. Since the first seller with $q > 1/2$ will receive all later sales, the most popular seller will have expected quality $E(q|q > 1/2) = 3/4$. With four sellers the expected quality of the best seller is 3/4, and the optimal seller is just as likely to be selected as one of the suboptimal ones. In general, as the number of sellers grows large the probability that the selected seller is of optimal quality goes to zero.

While this situation is undoubtedly an improvement over a market without ratings, it is surprising that providing buyers with the complete history of transactions does not lead to a fully efficient market.

2.2.2 IMPERFECTLY INFORMED BUYERS

Consider now an environment in which consumers are not always informed of previously sampled sellers. This can be modeled by consumers being informed, or "reading" previous ratings with probability r. When a consumer has read ratings (with probability r) she knows the quality of all sellers who have been sampled before her. When she has not read (with probability $1 - r$) she is totally uninformed of seller quality.

For ease of analysis I consider the case with a continuum of sellers available, when $m \to \infty$. This assumption avoids having to account for the possibility that an uninformed buyer will choose a previously sampled seller. Reducing the number of sellers and allowing for this complication leads to a qualitatively similar result, but with a lower rate of experimentation, since some uninformed buyers will "waste" their choice on a previously tried seller.

Since the probability of a buyer having $q > 1/2$ is itself 1/2, the expected number of sellers tried before such a high quality seller is found is 2. The expected number of sellers who have been tried after k buyers have purchased is thus $(1-r)(k-2)+2$. The expected best known quality is the highest order statistic for the number of sellers sampled, which is[1]

$$\frac{(1-r)(k-2)+2}{(1-r)(k-2)+3} = \frac{k+2r-kr}{1+k+2r-kr}$$

and the expected utility of the $k+1$th buyer is

$$E(u_{k+1}) = r\frac{k+2r-kr}{1+k+2r-kr} + (1-r)\frac{1}{2} = \frac{1+k+r+2r^2-kr^2}{2+2k+4r-2kr}.$$

Figure 11 shows the best observed quality for varying numbers of buyers. The best observed quality is decreasing in the probability of reading and increasing in the number of buyers.

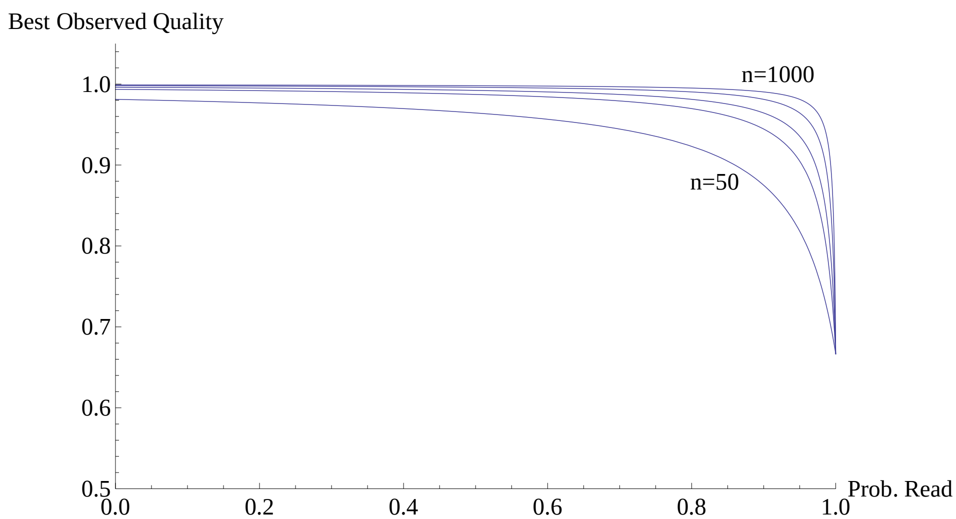

Figure 11: Expected best quality observed after k sellers have been sampled, for $k = 50, 150, 250, 500, 1000$.

The average welfare for all buyers can then be calculated as

$$\frac{1}{n}\sum_{k=1}^{n} \frac{1+k+r+2r^2-kr^2}{2+2k+4r-2kr} = \frac{(1+n)(r^2-1)+2r\cdot\Psi(\frac{2+n+r-nr}{1-r})-2r\cdot\Psi(\frac{1+2r}{1-r})}{2n(r-1)}$$

[1]It may be helpful to remember that the nth order statistic for n draws from the uniform distribution is $n/(n+1)$

Where $\Psi(z)$ is the digamma function defined as

$$\Psi(z) = \int_0^\infty \frac{e^{-t}}{t} - \frac{e^{-zt}}{1 - e^{-t}} dt$$

Unfortunately, due to the presence of the digamma function the expression for average welfare cannot be solved analytically for an optimal r. Since an analytical solution for r is not possible, I instead report here numerical solutions for several values of n. The welfare maximizing value of r increases from .83 with 50 buyers to .89 with 250 buyers and .93 with $1,000$ buyers. The average utility over the range of r and for a variety of n is shown in Figure 12 below. While the optimal r is increasing in the number of buyers, even with $n = 1,000$, a large number in real-world terms, it remains optimal to have a significant fraction (7%) of buyers uninformed. As the number of buyers grows the benefit of added information also grows, necessitating additional uninformed buyers to provide information through experimentation.

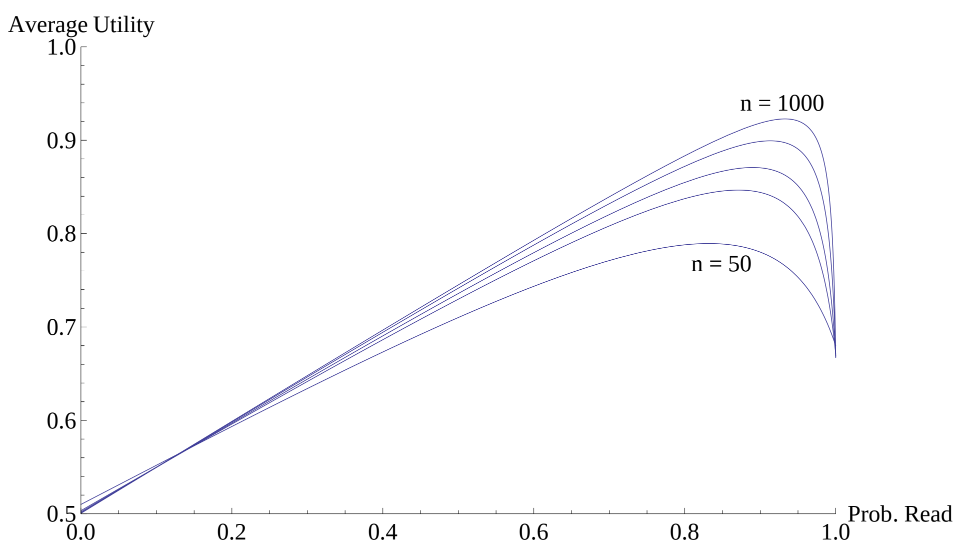

Figure 12: Average utility for n total buyers, $n = 50, 150, 250, 500, 1000$

It is possible to obtain an analytical solution for average welfare if we assume an infinite sequence of buyers. As the number of buyers tends to infinity we have

$$\lim_{n\to\infty}\frac{1}{n}\sum_{k=1}^{n}\frac{1+k+r+2r^2-kr^2}{2+2k+4r-2kr}=\frac{1+r}{2}$$

meaning that with sufficiently many buyers, average payoffs are strictly increasing in r and full reading is optimal.

It should be noted that while this paper focuses on how well-informed buyers should be, it is equally natural to ask what effect the frequency of rating has on market outcomes. The problem of encouraging more consumers to leave ratings has been studied before (Chen et al. 2008) since more information about sellers and products is generally considered to be beneficial for consumers. This paper's model shares that prediction, however for this model the dynamics of a change in the probability of rating are straightforward.

The effect of changing the rating frequency is simple: If each consumer rates with probability s, the market behaves as if the number of buyers is decreased by $(1-s)m$. Each buyer who does not rate has no impact on future buyers, since he does not reveal the quality of his seller. One way to think of this is that each buyer has a probability $1-s$ of leaving the market exactly as he found it, leading the next buyer to be faced with an identical decision.

It is easy to imagine alternative models of buyer behavior in which a change in the probability of rating has a more complex effect on buyer behavior. For example, if buyers choose whether or not to participate in the market an increase in the number of ratings may lead to more buyers participating. More buyers participating would in turn lead to an increase in the number of ratings, thus making it attractive for even more buyers to enter the market. This is an interesting and appealing model although it greatly complicates the focus of this paper: the analysis of how reading probabilities affect the market.

2.3 STOCHASTIC QUALITY

In practice, product quality tends not to be deterministic. Quality may be objectively variable due to manufacturing inconsistencies introducing defects into products, or it may vary subjectively due to consumers having heterogeneous tastes. Both cases can be captured by allowing quality to vary stochastically.

Allowing for stochastic quality is an important step to increase the validity and robustness of the model's predictions. In the deterministic model, a seller's quality is perfectly and permanently revealed after it is sampled. In the stochastic case a seller's quality is never known perfectly, and learning about that quality is a gradual process involving multiple buyers. Because of this randomness, the exact sequence of outcomes observed from a seller of a given quality is also random. This means that a high-quality seller could generate a bad outcome the first time he is selected by a buyer. This would cause future buyers' expectation of that seller's quality to be below that of untried sellers, leading him to have no further sales. This added source of uncertainty means that even sampling all sellers is not enough to identify the highest quality seller.

Because of the difficulty in solving the environment with stochastic quality analytically, I use a small computational simulation of buyer behavior to examine the market with stochastic quality. This model is similar to before, except that each buyer's payoff is no longer equal to her seller's quality, but is randomly determined by it. Qualities are still drawn from the uniform distribution over $[0, 1]$, but a seller with quality q now provides a buyer with a good outcome with probability q and a bad outcome with probability $1 - q$. Buyers receive a payoff of one from a good outcome and zero from a bad outcome. The expected payoff to a buyer from choosing a seller i is just q_i, as in the deterministic case.

The simulations were programmed in Python and run with 50 sellers and between 50 and $1,000$ buyers, in 50 buyer increments. The probability of reading was varied between zero and one in increments of $1/60$. Each combination of reading probability and number of buyers was then run for $2,000$ iterations, leading to a total of $2,400,000$ observations.

2.3.1 SIMULATION RESULTS

The behavior observed in the stochastic case is very similar to the analytical results in the deterministic case. As shown in Figure 13, the best discovered quality is decreasing in the probability of reading, with more buyers leading to a higher quality seller being found. Even for relatively large numbers of buyers (1,000), the quality of the best seller is significantly lower when all buyers read ($r = 1$) than for any other value ($p < .01$ Mann-Whitney). With full reading the best observed quality is .822 (95% confidence interval $[.821, .823]$) for any number of buyers. This is substantially higher than the prediction for the deterministic case, in which full reading leads to values between .66 and .68, depending on the number of buyers.

This difference in best discovered quality is likely due to the fact that above average sellers with quality near the mean are relatively likely to have a bad initial outcome and be passed over by future buyers. These are sellers, who, by virtue of their above-average quality, would continue to be chosen in the deterministic case. The reverse situation, a below average seller having an initially good outcome, is more likely to be corrected since additional buyers are likely to reveal the sellers true, below-average quality. The combination of these two effects leads to more experimentation, and an increase in the best observed quality.

Average welfare also shows a similar pattern to the deterministic case. As Figure 14 shows, welfare is strictly increasing across much of the range of r, peaks at a high value of r, and then decreases as r approaches one. Sequential Mann-Whitney tests show the welfare maximizing level of r varies based on the number of buyers, ranging from .93 with 100 buyers to .95 with $1,000$ buyers ($p < .01$ in each case). Average welfare is also increasing in the number of buyers, resulting both from the improved quality found and from fact that the worse-informed early buyers are a smaller fraction of the total number of buyers.

A significant difference between the stochastic and deterministic cases is the outcome for a given number of buyers. With stochastic quality more buyers are necessary to obtain similar results to the deterministic model. This is not surprising, considering that a seller's true quality is known immediately after selection in the deterministic case, but is only learned gradually in the stochastic case. Since real-world products often have variable or subjective quality, this suggests that the deterministic case may give overly conservative estimates of the optimal reading level for a given number of buyers.

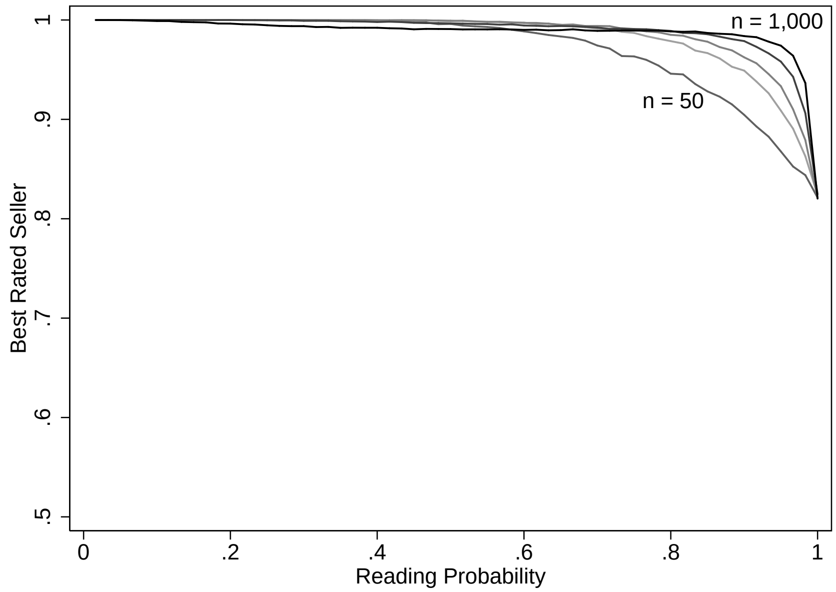

Figure 13: Best observed quality, $n = 50, 150, 250, 500, 1000$.

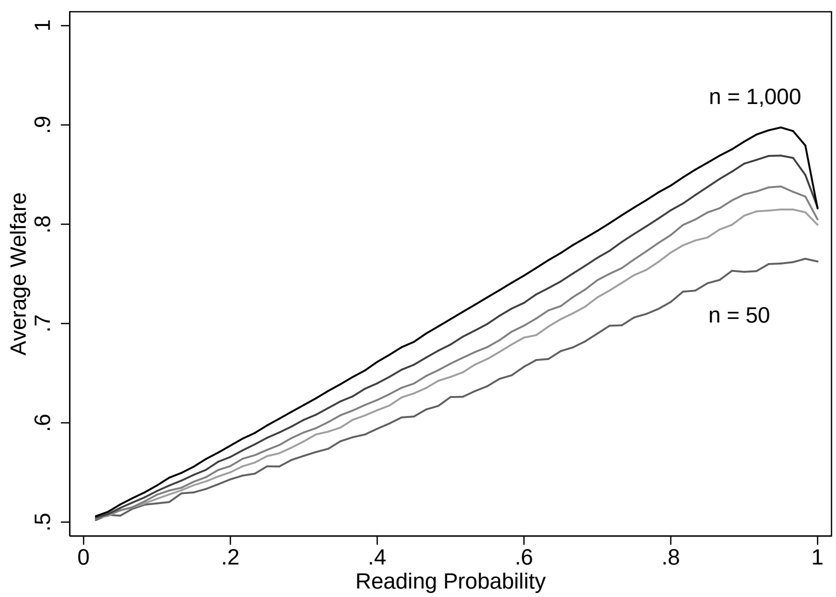

Figure 14: Average buyer welfare, $n = 50, 150, 250, 500, 1000$.

2.4 SUPPRESSED RATINGS

Despite the differences between the deterministic and stochastic cases, the first best outcome in both instances is still significantly worse than optimal. Even using the ideal reading frequency leads to a seller of significantly less then maximal quality becoming most popular, meaning that the market remains inefficient.

To improve this situation we can allow for the frequency of reading to change across time. Rather than setting a single reading frequency once and for all, there are instead two stages to informing buyers. In the first stage, early buyers are kept uninformed, leading to high levels of experimentation and a rapid improvement in the best known quality. In the second stage, the remaining buyers are perfectly informed of previous outcomes, allowing them to fully exploit all existing information. This suppression of early ratings eliminates the long-term inefficiency of always having some fraction of buyers uninformed, while still allowing for the experimentation necessary to find high-quality sellers.

Consider again the deterministic case, where buyers' payoffs are equal to their seller's quality. If $r = 0$, meaning that buyers are completely uninformed, the best expected quality observed after k buyers have purchased is $\frac{k}{k+1}$. Thus a relatively small number of uninformed buyers leads to a high quality seller being discovered. For example, if the first 10 buyers are uninformed, the expected best seller has quality of $10/11 = .91$. This is the same quality discovered by 50 buyers with a reading probability of .85. For the case with 50 buyers however, the exploration is a very gradual process, since 85% of the time buyers will choose a previously selected seller.[2] Early buyers will have lower quality sellers available to them than later buyers.

While keeping all buyers uninformed does lead to much more rapid experimentation, is also comes with the obvious downside that buyers are never able to exploit the improved information. Suppressing reading in the first stage, then allowing full reading in the second stage attempts to take the best from both approaches.

The social welfare from having t uninformed buyers and $n - t$ informed buyers is

[2] This assumes a seller with $q > 1/2$ has been found. Early buyers will perform more experimentation, since they are more likely to not know of an acceptable seller. This is a minor concern, however, since on average an acceptable seller will be found after only 2 buyers have moved

$$t \cdot \frac{1}{2} + (n - t)\frac{t}{t+1} = \frac{t(1 + 2n - t)}{2(1+t)}$$

and the optimal number of informed buyers is thus

$$-1 + \sqrt{2}\sqrt{1+n}$$

As Figure 15 shows, for $n = 50$ buyers it is optimal to have 9 (18%) uninformed buyers, and the remainder fully informed. For $n = 500$ the number of uninformed only increases to 27 (5.4%).

Compared to the standard rating system with $r = 1$, suppressed ratings provide a large improvement in average welfare, especially as the number of buyers increases. As Figure 16 illustrates, the welfare from the suppressed system goes from being equal to the standard system with 14 buyers, to being 40% greater with with 500 buyers.

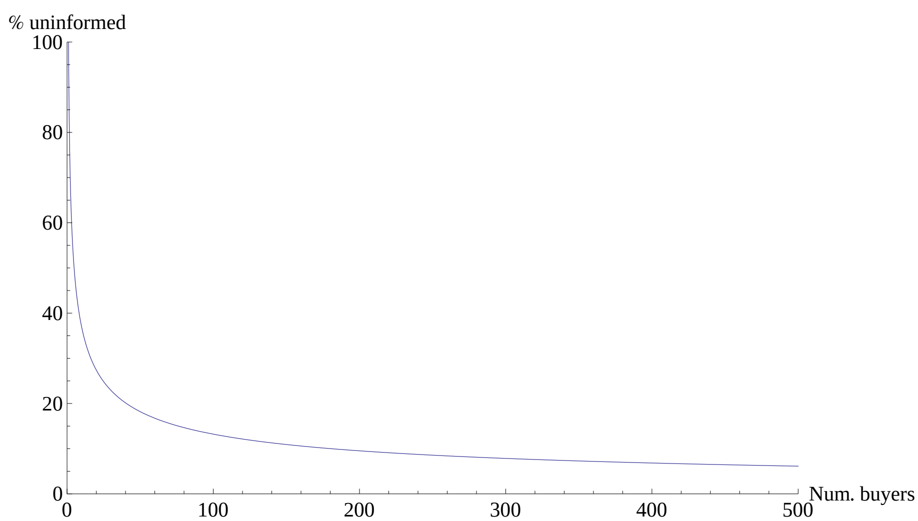

Figure 15: Optimal percentage of uninformed buyers.

Suppressing ratings does harm the earliest buyers by withholding information from them. However given the increase in average buyer welfare, these early buyers could be compensated by either a market-making retailer or directly from later buyers.

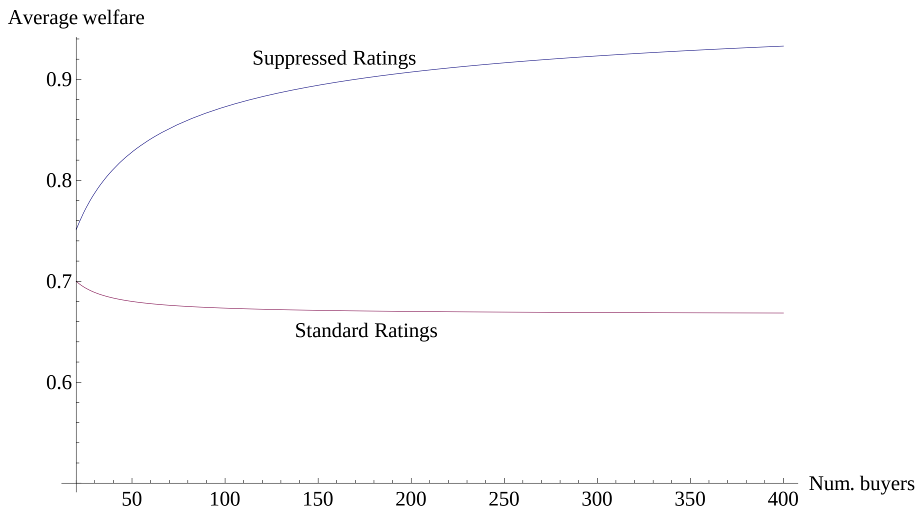

Figure 16: Average welfare comparison with standard and suppressed ratings.

A system with suppressed ratings has already been implemented by eBay.com, although for different reasons. Rather than encouraging experimentation, eBay delays the release of the first several ratings in order to encourage ratings. By waiting until several ratings have been given they make it difficult for sellers to connect ratings to buyers, thereby preventing retribution by buyers giving negative feedback. The observation that suppressing early ratings may aid buyers in finding the best seller recommends use and perhaps expansion of this existing system.

2.5 CONCLUSION

This paper examines a model of consumer search with a variable supply of information to buyers. This is simple alternative to the canonical information cascade model of herding, presented in the context of online ratings. It shows that individually rational but socially suboptimal herding can occur without any private signals guiding agents.

Findings do not differ significantly between the simplest model, with deterministic quality, and the richer, stochastic model. The highest quality that is found by any buyer is decreasing in the probability that buyers read ratings. This is driven by higher reading probabilities meaning that fewer buyers make uninformed purchases. Informed buyers will generally buy from a seller that was already selected by a previous buyer, whereas uninformed buyers may choose a new seller, increasing the number of rated sellers for future buyers. Having more sellers to choose from increases the expected quality of the best of those sellers, leading to higher qualities from lower reading probabilities.

A policy implication of these findings is that online retailers may actually benefit from showing ratings to few buyers, at least initially. This requires some additional incentive to uninformed buyers to offset the loss of information, but can ultimately lead to better information for consumers and thus to higher sales for retailers.

There are many possible extensions to this work. Studying the effect of alternative models of stochastic quality is a very natural next step. The model in this paper describes stochastic quality as being binary, either "good" or "bad." Obviously quality can be more nuanced, taking on many different levels and even multiple dimensions.

It would greatly enrich the model to endogenize the number of buyers entering the market based on observed ratings. This means that buyers could be drawn into the market by more and better ratings, or that the market could potentially collapse with enough poor ratings.

Related to the idea of endogenizing entry to the market, it is reasonable to endogenize a buyer's decision to rate, depending on the outcome they have observed. As Lafky (2010) has shown, buyers are more likely to rate products of extreme quality, leading to a U-shaped distribution of ratings. Including this behavior in the model could actually lead to an *improved* market, as good though sub-optimal sellers may remain unrated even after being selected. This would decrease the likelihood that buyers would get stuck in the type of undesirable outcome described in this paper.

3.0 ALL EQUILIBRIA OF THE MULTI-UNIT VICKREY AUCTION

Andreas Blume[a]
Paul Heidhues[b]
Jonathan Lafky[c]
Johannes Münster[d]
and
Meixia Zhang[e]

[*] Blume thanks the WZB (Berlin) and ISER (Osaka) for their hospitality and financial support. Heidhues and Münster gratefully acknowledge financial support from the Deutsche Forschungsgemeinschaft through SFB/TR 15. We benefitted from comments by seminar audiences at the University of Bonn, LMU (Munich), and WZB as well as Dan Kovenock and Benny Moldovanu. [a]Department of Economics, 230 South Bouquet Street, University of Pittsburgh, Pittsburgh, PA 15260, USA, *E-mail address:* `ablume@pitt.edu`; [b]Department of Economics, University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany and CEPR, *E-mail address:* `heidhues@uni-bonn.de`; [c]Department of Economics, 230 South Bouquet Street, University of Pittsburgh, Pittsburgh, PA 15260, USA, *E-mail address:* `jml44@pitt.edu`; [d]Wissenschaftszentrum Berlin für Sozialforschung (WZB), Reichpietschufer 50, 10785 Berlin, Germany and Free University of Berlin, *E-mail address:* `muenster@wzb.eu`; [e] Department of Economics, Princeton University, Fisher Hall, Princeton, NJ 08544, USA, *E-mail address:* `meixiaz@Princeton.edu`

3.1 INTRODUCTION

This paper completely characterizes the set of Bayesian Nash equilibria of the Vickrey auction for multiple identical units when buyers have non-increasing marginal valuations and there are at least three potential buyers. Equilibria fall into two classes:[1] In the first class, there is positive probability that there are positive bids below the maximum valuation. In this class, there is a threshold for valuations such that all bidders bid truthfully on any unit for which they have a valuation exceeding the threshold. Furthermore, there is a distinct bidder who bids the threshold value on any unit for which his valuation is below the threshold. The remaining bidders bid zero on any unit for which their valuation is below the threshold. In the second class of equilibria, there is zero probability of positive bids below the highest valuation. In this class, each bidder bids at or above the highest valuation on some number of units and bids zero on the remaining units in such a manner that the total number of positive bids across all bidders equals the number of units that are for sale. In any equilibrium, except the conventional equilibrium in dominant strategies, there is positive probability that a bidder wins a unit at a price of zero. In this sense all of these equilibria are collusive.

We also observe that all equilibria of the Vickrey auction are *ex-post* equilibria, i.e. bidders have no incentive to change their behavior even after all private information is revealed and therefore suffer no regret. Indeed, the entire set of equilibria within the first class remain equilibria for *any* change of the distribution function of bidders' valuations, including changes that affect the support of the distribution of bidders' valuations.

With any positive reserve price equilibrium becomes unique: Bidders bid truthfully on all units for which their valuation exceeds the reserve price. From this perspective, our result can be interpreted as providing an alternative foundation for the focus on the truthful-bidding equilibrium. Finally, we establish that only the truthful-bidding equilibrium ensures that the final allocation is in the core for all realizations of valuations.

We are interested in a full characterization of the equilibrium set of the Vickrey auction

[1] In the introduction we ignore inessential variations of behavior on measure zero sets of valuations. These are explicitly taken into account in Section 4.

for the following reasons: First, it is often argued that the single- or multi-unit Vickrey auction is susceptible to collusion. We show that a continuum of collusive equilibria exists, thereby providing some support for this view. (Although as noted above, we also establish that with any effective reserve price and at least three bidders, equilibrium is unique.) Second, for many distributions of valuations, there is a tension between selecting weakly undominated and payoff dominant equilibria in the multi-unit Vickrey auction. The truthful-bidding equilibrium results in comparatively high prices. In contrast, in the second class of equilibria the bidders obtain the units for free. In this class of equilibria, however, the units are allocated inefficiently across bidders, which reduces the bidders' total expected payoffs. Nevertheless, for many distribution functions bidders prefer equilibria in the second class to truthful bidding.[2] Furthermore, for some distributions of valuations, equilibria in the first class with a strictly positive threshold payoff dominate both truthful bidding and equilibria in the second class. Such equilibria improve the allocation relative to equilibria in the second class while still suppressing prices relative to truthful bidding (see Blume and Heidhues (2001) for an example in the single-unit case).[3] Third, if the auction is repeated (with an arbitrarily small positive probability), the one-shot equilibria can be used to construct collusive equilibria in which no players uses a weakly dominated strategy for *any* arbitrarily small positive discount factor. Fourth, if collusion is the result of a non-binding agreement, then it seems reasonable to assume that such communication has a small effect on the preferences in the subsequent auction. Any arbitrarily small (psychological) cost of breaking one's word and deviating from the informal agreement, however, transforms all collusive equilibria into strict equilibria.[4] Finally, the issue arouses our intellectual curiosity: We think it is interesting to know the implication of the most fundamental solution concept in game theory for

[2]Proposition 2 in Blume and Heidhues (2008) shows for the single-unit case that the equilibria in class two payoff dominate the truthful equilibrium for the uniform distribution and any distribution that first-order stochastically dominates it.

[3]Collusion via equilibria in the first class may also be more difficult to detect than collusion via equilibria in the second class because prices vary in this class of equilibria, which provides another reason to study these equilibria.

[4]Think of the dinners held by Judge Gary, chairman of U.S. Steel's board of directors, at the beginning of the 20th century that according to him lead to such mutual respect among steel industry leaders that all industry leaders found the obligation to cooperate "more binding [...] than any written or verbal contract". For a brief discussion of this case and, more generally, the role of "gentlemen agreements" in collusive arrangements see Scherer and Ross (1990, p. 235-236).

the multi-unit Vickrey auction—especially, given the simple structure of the set of Bayesian Nash equilibria.

Vickrey (1961) introduced the second-price sealed-bid auction for both the single- and the multi-object case. With private values, there is a unique equilibrium in undominated strategies: Bidders bid their valuations. Milgrom (1981) notes the existence of other (asymmetric) equilibria in the single-unit case. For two bidders, Plum (1992) describes yet more equilibria in the single-unit case. Blume and Heidhues (2004) characterize all equilibria of the single-unit Vickrey auction with independent private values and three or more bidders. Blume and Heidhues (2001) also cover the two-bidder case. Tan and Yilankaya (2006) show the existence of asymmetric equilibria with participation costs that are undominated. In contrast to these papers, here we consider the more complex multi-unit case.

3.2 SETUP

There are n bidders indexed by $i = 1, \ldots, n$, and m identical units $j = 1, \ldots, m$. Bidder i's vector of valuations is denoted $\mathbf{v}^i = (v_1^i, \ldots, v_m^i)$, where v_j^i represents the marginal valuation of the j-th unit. Thus the value of obtaining j units is $\sum_{k=1}^{j} v_k^i$. We assume marginal valuations to be non-increasing, i.e. $v_j^i \geq v_{j+1}^i$, for all $j = 1, \ldots, m-1$. Each bidder i's valuation vector is independently drawn from some distribution F_i that has positive density everywhere on the set $V := \{\mathbf{v}^i \in [0, v^h]^m | v_j^i \geq v_{j+1}^i, \forall j = 1, \ldots, m-1\}$ and no mass points.

The auction rules are as follows: Each bidder i submits a bid vector $\mathbf{b}^i = (b_1^i, \ldots, b_m^i) \in B := \{\mathbf{b}^i \in \mathbb{R}_+^m | b_j^i \geq b_{j+1}^i, \forall j = 1, \ldots, m-1\}$ independently from and simultaneously with the other bidders. Restricting bid vectors to belong to the set B is without loss of generality. It simply expresses that bids in any bid vector are automatically ranked from highest to lowest, and permits us to talk about "a bidder's bid on his first, second, ... unit." The auctioneer collects all bidders' bids and ranks them from the highest to the lowest bid, breaking ties by choosing with equal probability among all possible rankings among tying bids. Each bidder receives a unit for each of his bids that is among the m highest ranked bids. If bidder i wins k^i units, then he pays the k^i highest losing bids among his rivals. Formally, define $\mathbf{c}^{-i}$ as the vector consisting of the m highest bids submitted by bidders other than bidder i, ordered so that $c_1^{-i} \geq \ldots \geq c_m^{-i}$. A bidder who gains k^i units pays $\sum_{k=1}^{k^i} c_{m-k^i+k}^{-i}$.

We are interested in the Bayesian Nash equilibria of this game. Define $\mathcal{B}$ as the collection of Borel subsets of B. Then, a (behaviorally) mixed strategy for bidder i is a function $\beta^i : \mathcal{B} \times V \to [0,1]$ that satisfies (1) for all $W \in \mathcal{B}$, the function $\beta^i(W, \cdot) : V \to [0,1]$ is measurable, and (2) for every $\mathbf{v}^i$, the function $\beta^i(\cdot, \mathbf{v}^i) : \mathcal{B} \to [0,1]$ is a probability measure.[5] A profile of mixed strategies for all bidders has the form $\beta = (\beta^1, \ldots, \beta^n) = (\beta^i, \beta^{-i})$. Let $U_i(\beta, \mathbf{v}^i)$ denote the expected payoff of type $\mathbf{v}^i$ of bidder i from the mixed strategy profile β. Then the profile $\hat{\beta}$ is a Bayesian Nash equilibrium, henceforth simply equilibrium, if

$$U_i(\hat{\beta}, \mathbf{v}^i) \geq U_i(\beta^i, \hat{\beta}^{-i}, \mathbf{v}^i) \ \forall i, \ \forall \mathbf{v}^i, \forall \beta^i.$$

We say that the strategy β^i is pure if there exists a bid function $\mathbf{b}^i : V \to B$ such that $\beta^i(\{\mathbf{b}^i(\mathbf{v}^i)\}, \mathbf{v}^i) = 1$ for all $\mathbf{v}^i$. For simplicity, we refer to the bid function $\mathbf{b}^i$ as a pure strategy below. We denote a profile of bid functions by $\mathbf{b}(\cdot)$.

3.3 EXAMPLES

In order to illustrate the panoply of equilibria that are possible with only two bidders, to understand the role of limiting attention to the case with three or more bidders, and to get intuition for the proof of our characterization result, here we briefly discuss a few simple examples.

In FIGURE 1, we represent the essential aspects of one (type of) equilibrium in an auction with two bidders and two items for sale. Some features of this example survive when we restrict attention to three or more bidders, others do not. The two panels on top represent bidder one's bid function. Bidder two's bid function is shown in the bottom two panels. In this equilibrium bidder i's bid on his first (second) unit depends only on the higher (lower) of his two valuations. This feature, that a bidder's bid on his jth unit is independent of his

[5] See Milgrom and Weber (1985).

valuation for his other units will be a general characteristic of all equilibria in which there is a positive probability of positive bids below the maximum valuation.

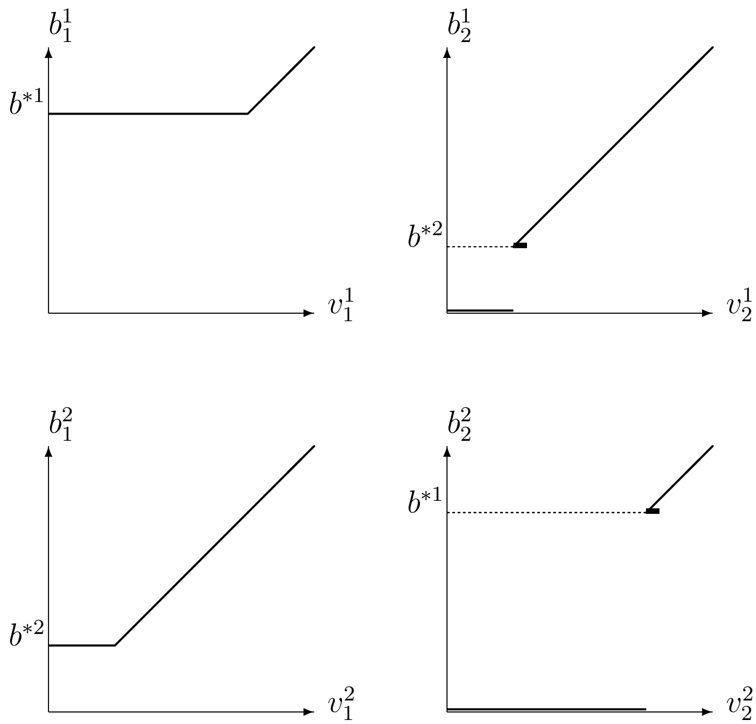

FIGURE 1

Bidder 1 bids b^{*1} on his first unit provided his high valuation satisfies $v_1^1 \leq b^{*1}$. Otherwise, he bids truthfully on his first unit. Since bidder 1 bids above his valuation on his first unit whenever he does not bid truthfully on that unit, we refer to him as a *high bidder* on his first unit. Bidder 1 bids zero on his second unit whenever his low valuation satisfies $v_2^1 \leq b^{*2}$. Otherwise, he bids truthfully on his second unit. Since bidder 1 bids below his valuation on his second unit whenever he does not bid truthfully on that unit, we refer to him as a *low bidder* on his second unit.

In the example, each bidder is a high bidder on his first unit and a low bidder on his second unit. We will find that having multiple high bidders is possible with more than two bidders only if the probability of positive bids below the maximum valuation is zero.

Observe also that in the example the critical value b^{*i} at which bidder i switches from bidding above his valuation to bidding truthfully on his first unit, differs across the two bidders. Again, this cannot occur with three or more bidders. With three or more bidders, if there is positive probability of positive bids below the maximum valuation, an equilibrium with distinct threshold values (b^{*1} and b^{*2}) as in FIGURE 1 is ruled out. To understand this, notice that a third bidder facing bidding behavior by the other two bidders as in FIGURE 1 would want to bid truthfully on his first unit for valuations above b^{*2}. But with two bidders bidding truthfully on their first unit above b^{*2}, it is no longer optimal for bidder 1 to maintain his postulated bidding strategy on the first unit.

This, in a nutshell, is the contagion effect that drives much of our result: If some bidder in a putative equilibrium puts in positive bids below the maximum valuation on some unit, e.g. bidder 1 bids near b^{*2} on his second unit, then for any other bidder, say bidder 2, who competes for that unit, those bids become potential prices. This disciplines this bidder's bidding behavior on that unit (viz. bidder 2 does not overbid on his first unit for valuations above b^{*2}). With three or more bidders in the auction, this discipline extends to at least two bidders. As a consequence the discipline extends to other units. In the example, with three bidders, it is no longer optimal for bidder 1 to overbid on his first unit for valuations for that unit in $[b^{*2}, b^{*1}]$.

A further possibility for equilibria in the two-bidder case, which generalizes the example of FIGURE 1, is that bidders have multiple gaps in their bid function, bidding truthfully outside the gaps, and adopting complementary roles of high and low bidders over the gaps. Here bidder one's gaps in his bid function for his first unit match bidder two's gaps in his bid function for his second unit, and vice versa. Any number and (matching) placement of gaps is possible. With three or more bidders all equilibria of this form disappear. The reason is that if two bidders have bid functions with gaps of this form, the third bidder has an incentive to bid inside these gaps. As a consequence, the bid functions with gaps are no longer optimal.

3.4 RESULTS

We have two principal results that jointly characterize the entire set of equilibria of the multi-unit Vickrey auction with three or more bidders. Our first result describes all equilibria in which there is positive probability of positive bids below the maximum valuation; i.e., there exists at least one bidder i for whom $b_j^i \in (0, v^h)$ with positive probability for at least one item j. We show that for each equilibrium in this class, there is a critical value b^* such that every bidder i bids truthfully on any unit j for which $v_j^i \geq b^*$. Furthermore, there is a single *high bidder* $\hat{\imath}$ who bids b^* on any unit j for which $v_j^{\hat{\imath}} \leq b^*$. The remaining bidders will be referred to as *low bidders.* Any low bidder i bids zero on any unit j for which $v_j^i \leq b^*$. It is important to emphasize that in each equilibrium the high bidder is unique and that the critical value b^* is the same for all units. All proofs are in the Appendix.

Proposition 3. *Let the number of bidders satisfy $n \geq 3$. Consider the class S^1 of profiles of bid functions $\mathbf{b}(\cdot)$ for which there is a bidder $\hat{\imath}$ and some $b^* \in [0, v^h)$ such that*

$$b_j^{\hat{\imath}}\left(\mathbf{v}^{\hat{\imath}}\right) = \begin{cases} v_j^{\hat{\imath}} & \textit{if } v_j^{\hat{\imath}} \geq b^* \\ b^* & \textit{otherwise} \end{cases}$$

for all $j = 1, ..., m$ and

$$b_j^{i}\left(\mathbf{v}^{i}\right) = \begin{cases} v_j^{i} & \textit{if } v_j^{i} > b^* \\ 0 & \textit{otherwise} \end{cases}$$

for all $i \neq \hat{\imath}$ and all $j = 1, ..., m$. Any profile in the class S^1 forms an equilibrium. Conversely for any equilibrium β in which $b_j^i \in (0, v^h)$ with positive probability for some bidder i and unit j, there is a profile of bid functions in the class S^1 that describes the behavior of each bidder for almost all valuations.

To check that these strategies are equilibria it suffices to verify that no bidder can gain by deviating and playing his dominant strategy. First, consider the low bidders. For valuations above b^*, these bidders bid their valuation anyhow. In case the valuation for some unit j is below b^*, these bidders bid zero for this unit and do not obtain it. Raising such bids to their true valuation $v_j^i < b^*$, does not increase the probability of obtaining the unit, since the mth highest bid is at b^* or above. Thus low bidders are playing a best response to the strategy of the high bidder. Similarly, suppose the high bidder has a valuation for unit j below b^*. In equilibrium, with probability one, he only obtains the unit if the $m-j+1$-th highest rival bid is zero, and thus pays a price of zero. Thus, lowering his bid to his valuation neither affects the probability with which he obtains the unit nor the price. Again therefore, he plays a best response.

The converse statement of Proposition 1 requires that we allow for variants on sets of measure zero of valuations. To see why, observe, for example, that if a bidder i bids above v^h— say $2v^h$—for unit j whenever his valuation for that unit, v_j^i, equals v^h, then the resulting strategy profile is still an equilibrium. Similarly, the high bidder can bid in the interval $(0, b^*)$ on a given unit for a set of valuations that has measure zero without affecting the equilibrium payoffs and incentives of his rivals or her own payoff. The converse statement proves that all differences from the strategy profiles in S^1 are inessential in the sense that they are restricted to sets of measure zero of valuations.

For the converse result, the key observation is that if a bidder bids at or near some interior value b^* with positive probability, this induces a contagion process with the result that all bidders bid their true value above b^* for all units. Suppose bidder 1 bids at or near b^*. Then this bid must sometimes win as otherwise there would be a bidder who would obtain the unit at a price weakly above b^* also when he has valuations below b^* for all units. This bidder would gain by switching to bidding his valuation on all units. Now consider bidder 2 and hold the behavior of all bidders other than bidder 1 and 2 fixed. Since the bid b^* sometimes wins, there exists at least one unit for which bidder 2 competes directly with bidder 1's bid b^*—in the sense that by bidding slightly above b^* rather than below, bidder 2 increases the probability of obtaining that unit. As this is true for all bidders other than 1, all bidders sometimes bid at or above b^*. Hence, with $n \geq 3$, there are both potentially many bids above as well as below b^*, which induces bidder 1 to bid at or above b^* for many units, which in turn induces other bidders to bid at or above b^* for many units. This contagion process continues until all bidders bid with positive probability at or above b^* for all units. Furthermore, as argued in Section 3, there can be no gap in the bid functions above b^* with $n \geq 3$ bidders as otherwise some bidder has an incentive to bid in the gap. Hence, for any unit all bidders bid their valuation above b^*. Finally, it is impossible for two bidders to have a mass point at b^*.

We are left to consider the case in which no bidder bids in the interval $(0, v^h)$. In an equilibrium for this case, there are exactly m high bids at or above v^h and all other bids are zero. Each high bid wins a unit at a price of zero. The allocation of units to bidders is arbitrary but independent of the realized valuations.

Proposition 4. *Let the number of bidders satisfy $n \geq 3$. Consider the class S^2 of strategy profiles β in which for each bidder i there exists a $k^i \in \{0, \dots, m\}$ such that for all $\mathbf{v}^i$, $\beta^i([v^h, \infty)^{k^i} \times \{0\}^{m-k^i}, \mathbf{v}^i) = 1$ and $\sum_{i=1}^n k^i = m$. Any strategy profile in the class S^2 forms an equilibrium. Conversely, suppose that β is an equilibrium strategy profile in which $b^i_j \in \big(0, v^h\big)$ with probability zero for all bidders i and units j. Then β prescribes the same behavior as one of the equilibrium profiles in S^2 for all i and almost all $\mathbf{v}^i$.*

To see that the above strategy profiles are equilibria, observe that any bidder who submits a positive bid for some unit, obtains that unit for free. Thus, submitting any positive bid on these units is part of a best response. Furthermore, the only way a bidder could increase the number of units he obtains with positive probability is to bid at or above the highest possible valuation for some additional unit(s). For each unit he would obtain over and above the ones he gets in equilibrium, his payment increases by at least the highest possible valuation v^h. Thus deviating is unprofitable. The converse statement is established in the Appendix.

Again note that—similar to Proposition 1—the converse statement of Proposition 2 requires that we allow for variants on sets of measure zero of valuations. For example, if a bidder who is meant to bid zero for unit j for all of her valuations bids instead in the interior $(0, v^h)$ for a set of measure zero of her valuations then this effects neither her nor her rivals' payoffs and is consistent with equilibrium.

Next we consider the robustness of the equilibria of the Vickrey auction. Four types of robustness are considered, robustness against varying the type distribution on a fixed payoff-type space, robustness against removing bidders, robustness against adding bidders with a larger set of payoff types, and robustness against introducing a positive reserve price. We find that the last of the four tests is the most stringent.

Remark 1 *All equilibria in $S^1 \cup S^2$ are ex-post equilibria.*

A Bayesian Nash equilibrium in a Bayesian game is an *ex-post* equilibrium if players' strategies remain optimal even if all private information is made public. This condition clearly holds for all equilibria in $S^1 \cup S^2$. In an *ex-post* equilibrium agents will never have to face regret. Furthermore, these equilibria are robust in the sense that they are invariant to changes of the distribution of players' private information (on a given type space).

The asymmetric equilibria are not, however, *ex post* in the sense of Holzman and Monderer (2004). They require what they refer to as "*ex-post* equilibria in Vickrey-Clarke Groves mechanisms" to remain equilibria when an arbitrary subset of players is excluded from playing. In the asymmetric equilibria of Proposition 1, if the high bidder is excluded, the remaining low bidders who bid zero would have a positive probability of obtaining the unit when bidding zero. Thus, once the high bidder is taken out, low bidders gain from bidding their true valuations. As Holzman and Monderer show, their notion of "*ex post* equilibria in Vickrey-Clarke Groves mechanisms" requires players to use symmetric strategies. Since the equilibria of Proposition 2 also do not satisfy their requirement, their notion selects the unique weakly dominant strategy profile. If, however, the high bidder is known to be active, our asymmetric equilibria of Proposition 1 are robust to adding or excluding low bidders.

Remark 2 *All equilibria in S^1 are robust to enlarging the type space by including bidders with higher valuations than v^h.*

The equilibria of Proposition 1 are not only robust to changing the distribution over a given payoff-type space but also allow arbitrary extensions of the type space to bidders with possibly higher valuations.[6] This is not the case with the equilibria of Proposition 2. Indeed, only if there is a single high bidder who always submits the same bid b^* at or above v^h for all units is it possible to prescribe equilibrium bidding behavior for the new types (i.e. bid their valuation above b^* and behave as before below b^*) without changing the bidding behavior of existing types.

[6] With free disposal, zero is a natural lower bound on possible types.

Suppose the auctioneer sets a positive reserve price $r > 0$, such that for any unit a bidder obtains, his bid has to be at least as high as the reserve price and the reserve price is the minimum price for any unit. Without loss of generality, we can identify bids below the reserve price, or not bidding, with bidding zero. If m' is the number of bids at or above r, the auctioneer hands out $\mu = \min\{m', m\}$ units to the bidders with the μ highest bids. A bidder who gains k^i units pays $\sum_{k=1}^{k^i} \max\left\{c_{m-k^i+k}^{-i}, r\right\}$.

A positive reserve price below the maximum valuation has the same contagious effect as having bids with positive probability in that range. Bidders with valuations below the reserve price will refrain from bidding above the reserve price, for fear of winning a unit at a price above their valuation for that unit. As a consequence, any bidder will put in a bid above the reserve price for any unit for which his valuation exceeds the reserve price. Bidders with valuations between those bids and the reserve price will want to bid in that range. With three or more bidders, this eliminates any potential gaps in the bid function above the reserve price, and therefore bidders bid truthfully for any unit with a valuation above the reserve price. The details of the proof are virtually identical to that of Proposition 1 and are therefore omitted.

Corollary 1. *With a positive reserve price $r > 0$, equilibrium is (essentially) unique: Bidders refrain from bidding on any unit for which their valuation is less than the reserve price. Otherwise they bid their valuation for each unit (except possibly at r or v^h).*

We conclude by analyzing the relationship between the equilibria of the Vickrey auction and the core since it is often argued that the desirability of a trading institution requires that its equilibria are in the core. Let k^i be the number of units that bidder i obtains and $k = (k^1, ..., k^n)$. The set of feasible allocations of the units is

$$X = \left\{ k \in \mathbb{Z}_+^n : \sum_{i=1}^{n} k^i \leq m \right\}.$$

Following Ausubel and Milgrom (2002), the coalitional game associated with the Vickrey auction consists of the set of players $N = \{0, 1, ..., n\}$, where 0 denotes the seller, and the coalitional value function

$$w(C) = \begin{cases} \max_{k \in X} \sum_{i \in C} \left(\sum_{j=1}^{k^i} v_j^i \right), & \text{if } 0 \in C \\ 0, & \text{if } 0 \notin C \end{cases}, \quad \forall C \subseteq N.$$

Denote the payoff of bidder i when paying p^i for k^i units by $\pi^i = \left(\sum_{j=1}^{k^i} v_j^i\right) - p^i$ and the payoff of the seller by $\pi^0 = \sum_{i=1}^{n} p^i$. The set of core payoffs is given by (e.g. Milgrom (2004, p. 303))

$$Core\,(w, N) = \left\{ (\pi^0, \ldots, \pi^n) \left| \sum_{i \in N} \pi^i = w\,(N) \text{ and } w\,(C) \leq \sum_{i \in C} \pi^i \text{ for all } C \subset N \right. \right\}.$$

The following three remarks relate the equilibria of the multi-unit Vickrey auction to the core.

Remark 3 *For any profile of valuations, the profile of payoffs resulting from the truthful equilibrium is in the core.*

Consider any coalition including the seller, which does not obtain all units in the truthful equilibrium. If the members of this coalition trade only among themselves, they obtain all units, but lose the payments from the bidders who are not part of the coalition. Since each bidder pays the externality he imposes, these payments are at least as high as the value of the additional units for the coalition. Therefore, no coalition can block the truthful equilibrium allocation.[7]

Remark 4 *For any equilibrium that differs from the truthful equilibrium (except possibly at v^h), there is a positive measure set of valuation profiles such that the equilibrium payoffs are not in the core.*

For any equilibrium except the truthful equilibrium, there is positive probability that all units are sold at a price of zero and there is a bidder who has a positive willingness to pay for an additional unit. Similarly, if there is a positive reserve price, with positive probability not all units are sold even though bidders have positive valuations for all units. In this case the outcome is obviously inefficient and, thus, not in the core.

[7] A simple proof based on the above argument is contained in the appendix. The result also follows from the fact that the units are substitutes in the sense of Ausubel and Milgrom (2002) and Theorem 8.5 in Milgrom (2004).

Remark 5 *If there is a positive reserve price, there is a positive measure set of valuation profiles such that the equilibrium payoffs are not in the core.*

3.5 PROOFS

Lemma 1 establishes the contagion property that is central to the proof of Proposition 1. According to Lemma 1, if one bidder, say bidder i, submits a bid c in the interior of the support of the valuation distribution for any unit, then all bidders bid truthfully for every valuation of every unit that is above c. The (somewhat lengthy) proof is broken down into steps, for which we first briefly sketch the intuition:

1. In step 1, we show that whenever bidder i bids c on his jth unit, then he wins at least j units with positive probability. The reason is that otherwise at least one other bidder with positive probability would win at least one unit at a price above the valuation for that unit.
2. In step 2 we conclude from step 1 that bidders other than i have an incentive to bid above c on $m - j + 1$ units if they have that many or more valuations above c.
3. In step 3 we show that with positive probability all bidders bid below v^h on their first unit, and therefore with positive probability all bidders bid above zero on their last unit. Intuitively, by step 2, bidder i bidding at c leads to a contagion process that leads other bidders to put in increasingly high bids. As a consequence, bids above v^h would win with positive probability at prices near v^h. Bidders with low enough valuations would want to avoid such bids.
4. In step 4, we first strengthen the conclusion from step 3 to infer that indeed the probability of anyone ever bidding at or above v^h is zero. Hence, there are bids below v^h that win all units and for high enough valuations a bidder will want to make such bids. This implies that every bidder bids with positive probability in the interval $(0, v^h)$ on all of his units.

5. In step 5, using step 4, and iterating on bidders, it is not hard to see that for high enough valuations all bidders must bid near v^h.
6. Step 6 is central to the argument and establishes that there are no gaps in first-unit bids above c : in every interval there is some bidder who, with positive probability, bids on his first unit in this interval. Essentially, if there were a gap any bidder whose lowest valuation falls into the gap would want to have his lowest bid in the gap. But this would induce other bidders to bid in the gap on their first unit, where their first-unit bid is kept below the upper bound of the putative gap because of step 5.
7. Step 7 strengthens the conclusion from step 6 by first showing that there is no bidder who has gaps in his first-unit bid function above c. This implies that every bidder bids truthfully on his last unit above c. Hence every bidder bids truthfully on his first unit for valuations above c.
8. Step 8 uses step 7 to argue that for any unit and for any valuation on that unit there is positive probability that all rival bids are either in a small interval immediately above or below that valuation. Therefore it is uniquely optimal to bid truthfully on that unit.

Lemma 1. *Suppose that β is an equilibrium in which for some bidder i, some unit j, some $c \in (0, v^h)$ and all $\epsilon > 0$ there is positive probability that $b^i_j \in [c, c+\epsilon)$, i.e. that $\int_V \beta^i(\{\mathbf{b}^i \in B | b^i_j \in [c, c+\epsilon)\}, \mathbf{v}^i) dF_i(\mathbf{v}^i) > 0$. Then $\beta^l(\{\mathbf{b}^l \in B | b^l_k = v^l_k\}, \mathbf{v}^l) = 1$ for all $\mathbf{v}^l$ with $v^l_k \in (c, v^h)$, all $l = 1, \ldots, n$, and all $k = 1, \ldots, m$; i.e., every bidder bids truthfully on every unit for which his valuation belongs to the interval (c, v^h).*

Proof. Since the argument is somewhat lengthy and for readability the proof will be presented in a series of steps. The property that is established in each step will be highlighted by italics.

Step 1. Whenever $b^i_j \geq c$, bidder i wins at least j units with positive probability. Otherwise, there is probability one that at least $m-j+1$ bids made by bidders other than i are above c. But this cannot be because then one of these bidders, say bidder l, with positive probability would win $k_l > 0$ units and pay more than his value for his k_lth unit.

Step 2. If bidder $l_1 \neq i$ has $m-j+1$ or more valuations above c, then he has at least one bid above c, by Step 1. In the case of $m = j$ this means that he has $m - j + 1$ bids above c. The remainder of the argument in this step generalizes this observation to the case where $m > j$. If conditional on bidder l_1 having one bid above c there is positive probability that all bids $b^i_j \geq c$ remain winning bids, and bidder l_1 has $m - j + 1$ or more valuations above c, then bidder l_1 has at least two bids above c. If on the other hand conditional on bidder l_1 having one bid above c, for some $\eta > 0$ all bids $b^i_j \in [c, c + \eta)$ become losing bids with probability one, then there is probability one that there are at least $m - j$ other bids above c. Suppose these $m - j$ bids are made by bidders other than l_1. This cannot be because then one bidder other than bidder l_1, say bidder l, with positive probability would win k_l units and pay more than his value for his k_lth unit. Thus, again if bidder l_1 has $m - j + 1$ or more valuations above c, he will have at least two bids above c. Iterating this argument, we conclude that *if bidder l_1 has $m - j + 1$ or more valuations above c, he will have at least $m - j + 1$ bids above c. Furthermore, from Step 1 it follows that if a bidder $l_1 \neq i$ has less than $m - j + 1$ valuations above c, he will have no more than $m - j$ bids above c.*

Step 3. One implication of *Step 2* is that there is positive probability that there are at least $m + 1$ bids at or above c made by bidders i and l_1.

A bidder $l_2 \neq i, l_1$ cannot bid with probability one at or above v^h for his first unit. If such a bid wins in the event of $m + 1$ or more positive bids by i and l_1, then with positive probability l_2 pays more than his value for his first unit. If such a bid loses, then a bidder other than bidder l_2 would pay more than his value for at least one of his units.

Exchanging the roles of bidders in the above argument, it follows that with positive probability all bidders other than bidder i bid below v^h on their first unit. Hence, with positive probability bidder i's bid on his last unit exceeds the highest bid of bidders other than i on their first unit.

The probability of bidder i's bid on his mth unit being equal to or exceeding v^h, however, equals zero. Otherwise the support of the distribution of bids of bidders other than bidder i on their first unit has to be bounded away from v^h. In that case, denote the supremum of this support by $\bar{b}$. Note (using *Step 2*) that $v^h > \bar{b} > c > 0$, and hence it follows from *Step 1* that a bid at or above $\bar{b}$ wins at least one unit with positive probability. Then, if there is a bidder $l_1 \neq i$ whose distribution of first-unit bids has a mass point at $\bar{b}$, the other bidders have an incentive to bid above $\bar{b}$ on their first unit with positive probability. If there is no such bidder, and wlog $\bar{b}$ is the supremum of the support of the distribution first-unit bids by bidder $l_1 \neq i$, then there is a bidder $l_2 \neq i, l_1$ who with positive probability has an incentive to bid in the interval $[\bar{b}, v^h)$, where we use the requirement that his first-unit bids are bounded away from v^h. If such bids are in the interior of this interval with positive probability, we have a contradiction because this violates the definition of $\bar{b}$. Otherwise, we must have a mass point at $\bar{b}$, which would take us back to the earlier case.

Since bidder i's bid on his mth unit is less than v^h with positive probability, there is a bidder other than bidder i who bids with positive probability in the interval $(0, v^h)$. Thus, from the foregoing argument, with the role of bidders exchanged, it follows that *with positive probability all bidders bid below v^h on their first unit. Hence, with positive probability all bidders bid above zero on their last unit.*

Step 4. Consider three distinct bidders l_1, l_2 and l_3. Suppose bidder l_1 bids with positive probability at or above v^h on his first unit. Then the support of the distribution of bids of bidder l_2 on his last unit must be bounded away from v^h because otherwise bidder l_1 with positive probability would win a unit at a price above his valuation for that unit. Denote the supremum of this support by $\bar{b}_m$. From *Step 3*, we know that $\bar{b}_m \in (0, v^h)$. Then bidder l_3 has an incentive to bid with positive probability at or above $\bar{b}_m$ on his first unit. If such bids with positive probability are below v^h, we have a contradiction because then bidder l_2 would have an incentive to bid with positive probability above $\bar{b}_m$ on his last unit in violation of the definition of $\bar{b}_m$. Continue then with the case where both bidders l_1 and l_3 with positive probability bid at or above v^h on their first unit. Then the support of the distribution of bids of bidder l_2 on his $m-1$th unit must be bounded away from v^h. Denote the supremum of this support by $\bar{b}_{m-1}$ and note that $\bar{b}_{m-1} \in [\bar{b}_m, v^h)$. Then bidder l_3 has an incentive to bid with positive probability at or above $\bar{b}_{m-1}$ on his first two units. If in this case bidder l_3's bids on his second unit are below v^h with positive probability, we have a contradiction because then bidder l_2 would have an incentive to bid above $\bar{b}_{m-1}$ on his $m-1$th unit. Continue then with the case where bidder l_1 bids with positive probability at or above v^h on his first unit and bidder l_3 bids with positive probability at or above v^h on his first two units. Iterating this argument, we find that the support of the distribution of bidder l_2's bids on his first unit must be bounded away from v^h. Denote the supremum of this support by $\bar{b}_1$ and note that $0 < \bar{b}_m \leq \bar{b}_{m-1} \leq \ldots \leq \bar{b}_1 < v^h$. Then bidder l_3 has an incentive to bid with positive probability at or above $\bar{b}_1$ on all of his units. If in this case bidder l_3's bids on his mth unit are below v^h with positive probability, we have a contradiction because then bidder l_2 would have an incentive to bid above $\bar{b}_1$ on his first unit in violation of the definition of $\bar{b}_1$. But at the same time, it is impossible for bidder l_3 to bid with positive probability at or above v_h on all of his units because then with positive probability there would be $m + 1$ bids at or above v^h and therefore at least one bidder with positive probability would win an item at a price above his valuation for that item. Since the choice of the three bidders was arbitrary, it follows that *the probability that any bidder bids at or above v^h on any of his units is zero.* Combined with the earlier observation that all bidders bid with positive probability in $(0, v^h)$ on their last unit, this implies that *all bidders bid with positive probability in* $(0, v^h)$ *on* all *of their units.*

Step 5. Pick an arbitrary pair of distinct bidders l_1 and l_2. Suppose that the support of the distribution of bidder l_1's bids on his last unit is bounded away from v^h. Denote the supremum of this support by $\overline{h}$. From *Step 4*, $0 < \overline{h} < v^h$ and from *Step 1*, bids at or above $\overline{h}$ win at least one unit with positive probability. Hence bidder l_2 must bid with positive probability in $[\overline{h}, v^h)$ on his first unit, as he cannot be bidding at or above v^h by *Step 4*. This in turn implies that bidder l_1 has an incentive to bid with positive probability in $(\overline{h}, v^h)$ on his last unit, contrary to the definition of $\overline{h}$. Since the choice of bidders was arbitrary, it follows that *for all* $\epsilon > 0$, *there is positive probability that* all *bids of* all *bidders are in the interval* $(v^h - \epsilon, v^h)$.

Step 6. Suppose there is an interval $(e, e') \subseteq (c, v^h)$ in which no bidder bids on his first unit with positive probability. Then there is a maximal such interval $(\underline{e}, \overline{e}) \subseteq (c, v^h)$ by exactly the same argument as in Lemma A3 of Blume and Heidhues (2004). From *Step 5*, we know that $\overline{e} < v^h$ and that with positive probability every bidder bids in $(\overline{e}, v^h)$ on all of his units.

Denote the infimum of the support of the distribution of bids on any bidder's last unit in the interval $[\overline{e}, v^h]$ by $\underline{b}_m$. Suppose that $\overline{e} < \underline{b}_m$. Without loss of generality, suppose that $\overline{e}$ is the infimum of the support of the distribution of bidder l_1's first unit bids in the interval $[\overline{e}, v^h]$ and that $\underline{b}_m$ is the infimum of the support of the distribution of bidder l_2's mth-unit bids in the same interval. Then, if bidder l_3 has all of his valuations in the interval $(\overline{e}, \underline{b}_m)$, he will have all of his bids in the interval $(\overline{e}, \underline{b}_m]$; this follows from *Step 2*. If in this case he bids with positive probability in $(\overline{e}, \underline{b}_m)$ on his mth unit, we have the desired contradiction. Otherwise, bidder l_3 with positive probability bids at $\underline{b}_m$ on his last unit. But then, if bidder l_2 has all of his valuations in the interval $(\overline{e}, \underline{b}_m)$ he will have all of his bids in the same interval, and we have again reached a contradiction. *We conclude that we must have* $\overline{e} = \underline{b}_m$.

Then take a bidder all of whose valuations are in $(\underline{e}, \overline{e})$ *to derive a contradiction to the assumption that no one bids with positive probability in* $(\underline{e}, \overline{e})$ *on his first unit. Distinguish the cases where* $\underline{e} = c$ *and* $\underline{e} > c$:

Step 6A. Consider the case where $\underline{e} = c$. Without loss of generality, suppose that $\overline{e}$ is the infimum of the support of the distribution of bidder i_1's mth-unit bids in the interval $[\overline{e}, v^h]$. Then by *Step 2* a bidder i_2 other than i and i_1 who has all of his valuations in the interval $(c, \overline{e})$ has an incentive to bid on his first unit in the interval $(c, \overline{e}]$. If such bids are in the interval $(c, \overline{e})$, we have the desired contradiction. Therefore, suppose the distribution of bidder i_2's bids on his first unit when all of his valuations are in the interval $(c, \overline{e})$ has a mass point at $\overline{e}$. *Step 1* implies that the distribution of bidder i_1's bids on his mth unit cannot have a mass point at $\overline{e}$; otherwise bidder i_2 with positive probability would win a unit at a price above his valuation for that unit.

Continuing with the condition that the distribution of bidder i_2's bids on his first unit has a mass point at $\overline{e}$, there are two subcases of the case $\underline{e} = c$ to consider:

Step 6Ai. First, if $i_1 = i$, then, as we just saw, for all bidders other than bidder i their distributions of first-unit bids must have mass points at $\overline{e}$.

Step 6Aii. Second, if $i_1 \neq i$, then since i_1's distribution of bids on his mth unit does not have a mass point at $\overline{e}$, for any $\epsilon > 0$, there is positive probability that bidder i bids in the interval $(\overline{e}, \overline{e} + \epsilon)$ on all of his units. But this implies that $\overline{e}$ is the infimum of the support of the distribution of bidder i's bids on his last unit in the interval $[\overline{e}, v^h]$.

Thus, in either case all bidders other than bidder i all of whose valuations belong to the interval $(c, \overline{e})$ must have a mass point at $\overline{e}$ in their distributions of first-unit bids. But then, consider the situation where bidder i has $m-1$ valuations above $\overline{e}$ and his mth valuation below $\overline{e}$. Since with positive probability all of his rivals bid at $\overline{e}$ on their first unit and there are at least two such rivals, i has to bid for $m-1$ units above $\overline{e}$ to ensure that he receives those units whenever their prices are below his values for those units. At the same time, *Step 2* implies that he will bid below $\overline{e}$ for his last unit. Let the remaining bidders all have valuations in $(c, \overline{e})$ such that they bid at $\overline{e}$ on their first unit. Then there is a bidder with valuation on his first unit below $\overline{e}$, who wins one unit and pays $\overline{e}$. Since this case has positive probability, this gives the desired contradiction.

Step 6B. Consider the case where $\underline{e} > c$. Without loss of generality, suppose that $\underline{e}$ is the supremum of the support of the distribution of bidder i_1's 1st-unit bids in the interval $[c, \underline{e}]$. Similarly, without loss of generality suppose that $\overline{e}$ is the infimum of the support of the distribution of bidder i_2's mth-unit bids in the interval $[\overline{e}, v^h]$. Then a bidder i_3 other than i_1 and i_2 who has all of his valuations in the interval $(\underline{e}, \overline{e})$ has an incentive to bid on his first unit in the interval $[\underline{e}, \overline{e}]$. If such bids are in the interval $(\underline{e}, \overline{e})$, we have the desired contradiction. This leaves us with the possibility there is a mass point for such bids at either $\underline{e}$ or at $\overline{e}$.

Step 6Bi. Therefore, suppose the distribution of bidder i_3's bids on his first unit has a mass point at $\overline{e}$. Note that the distribution of bidder i_2's bids on his mth unit cannot have a mass point at $\overline{e}$; otherwise bidder i_3 with positive probability would win a unit at a price above his valuation for that unit. There are two cases to consider:

Step 6Bi(a). First, consider $i_2 = i_1$. Then for all bidders other than bidder i_1, their distributions of first-unit bids must have mass points at either $\overline{e}$ and/or $\underline{e}$. If there is no mass point at $\underline{e}$, the argument for the case $\underline{e} = c$ applies. If on the other hand, there is a bidder i_4 other than i_1 whose distribution of first-unit bids has a mass point at $\underline{e}$, then there is a bidder (other than i_4 and i_1) who with positive probability prefers to bid in $(\underline{e}, \overline{e})$ on his last unit. But this implies that bidder i_4 sometimes wants to outbid this bidder with his first unit bid when all his valuations are in the interval $(\underline{e}, \overline{e})$. Since, by assumption he is not bidding in the interior of this interval, his distribution of first-unit bids must also have a mass point at $\overline{e}$. Hence, all bidders other than i_1 must have distributions of first-unit bids with mass points at $\overline{e}$.

Step 6Bi(b). Second, if $i_2 \neq i_1$, then since i_2's distribution of bids on his mth unit does not have a mass point at $\overline{e}$, for any $\epsilon > 0$, there is positive probability that bidder i_1 bids in the interval $(\overline{e}, \overline{e}+\epsilon)$ on all of his units. But this implies that $\overline{e}$ is the infimum of the support of bidder i_1's bids on his last unit in the interval $[\overline{e}, v^h]$, which takes us back to *Step 6Bi(a).*

Thus, in either case all bidders other than bidder i_1 must have a mass point at $\overline{e}$ in their distributions of first-unit bids, and therefore the discussion for the case $\underline{e} = c$ applies, which rules out this possibility.

Step 6Bii. Now suppose there is no bidder, like i_3 above, whose distribution of first-unit bids has a mass-point at $\overline{e}$. From the foregoing argument, it is without loss of generality to focus on the case where $i_2 = i_1$. Then, for all bidders other than i_1, their distribution of first-unit bids must have a mass point at $\underline{e}$. Then there are at least two bidders, which are different from i_1, who for a positive-probability set of values prefers to bid in the interval $(\underline{e}, \overline{e}]$ on their mth unit. If such bids with positive probability are in $(\underline{e}, \overline{e})$, then we get a contradiction because at least one bidder other than i_1 would have to have a mass point at $\overline{e}$ in his distribution of first-unit bids. Hence, both bidders' distribution of mth-unit bids must have a mass point at $\overline{e}$, which is impossible because then there would be positive probability that a bidder wins a unit at a price above his valuation for this unit.

Hence, we may conclude that in every open interval above c there is some bidder who bids in this interval with positive probability on his first unit.

Step 7. Suppose there is an interval $(d, d') \subset (c, v_h)$ and a bidder l_1 who does not bid with positive probability in (d, d') on his first unit. Then by *Step 6* for every open subinterval of (d, d') there is a bidder other than bidder l_1 who bids with positive probability in this subinterval on his first unit. Hence, by *Step 2,* bidder l_1 must bid truthfully on his mth unit over this interval. Then, using *Step 2* once more, any bidder $l_2 \neq l_1$ must bid truthfully on his first unit over the same interval. Hence, any bidder $l_3 \neq l_1, l_2$ must bid truthfully on both his first and his mth unit over (d, d'). But then bidder l_1 must bid truthfully on his first unit over the interval (d, d'), which leads to a contradiction. *Therefore every bidder bids with* *positive probability in every interval* $(d, d') \subset (c, v_h)$ *on his first unit. Hence, every bidder bids truthfully on his mth unit over the interval* (c, v^h)*. This implies that every bidder bids truthfully on his first unit over the interval* (c, v^h)*.*

Step 8. Suppose that for some bidder l and unit k, we have $v_k^l \in (c, v^h)$. Consider two bidders l_1 and l_2 different from bidder l. From *Step 7,* for any $\epsilon > 0$, there is positive probability that all of bidder l_1's bids are in the interval $(v_k^l - \epsilon, v_k^l)$ and that all of bidder l_2's bids are in the interval $(v_k^l, v_k^l + \epsilon)$. Thus, if bidder l were to bid below $v_k^l - \epsilon$, he would run the risk of not winning his kth unit when it is available at a price below his valuation for that unit. Similarly, if bidder l were to bid above $v_k^l + \epsilon$, he would run the risk of winning a kth unit at a price exceeding his valuation for that unit. *Therefore, for any valuation* $v_k^l \in (c, v^h)$*, it is uniquely optimal for bidder* l *to bid truthfully on his* k*th unit.* □

We are ready to prove Proposition 1.

Proof of Proposition 1: We have established that all strategy profiles in the class S^1 are equilibria in the text. It remains to show that any equilibrium in which $b_j^i \in \left(0, v^h\right)$ with positive probability for some bidder i and unit j corresponds to a profile of bid functions in the class S^1 that describes the behavior of each bidder for almost all valuations.

Suppose that $b_j^i \in \left(0, v^h\right)$ with positive probability for some bidder i and unit j. Let

$$b^* := \inf \left\{ b \in (0, v^h) | \exists i, j \text{ s.t. } \forall \epsilon > 0, \Pr \left\{ b_j^i \in [b, b + \epsilon) \right\} > 0 \right\}.$$

For all $v_j^i \in (b^*, v^h)$, bidders bid truthfully by Lemma 1, i.e. $\beta^i(\{\mathbf{b}^i \in B | b_j^i = v_j^i\}, \mathbf{v}^i) = 1$ for all $\mathbf{v}^i$ with $v_j^i \in (b^*, v^h)$. If $b^* = 0$, we are done.

Thus consider the case where $b^* > 0$. Whenever a bidder has his valuation for a unit in $(0, b^*)$, then he bids in $[0, b^*]$ for this unit; otherwise, by Lemma 1 there would be positive probability of this bidder winning a unit at a price above his valuation for that unit. Trivially, such bids cannot be in $(0, b^*)$ with positive probability.

Suppose there are two distinct bidders l_1 and l_2 (and possibly others) who with positive probability submit a bid b^* on their first unit. Then there exists a bidder $l_3 \neq l_1, l_2$ who bids, with positive probability, for exactly $m-1$ units above b^* and for his last unit below b^*. This implies that with positive probability l_1 wins exactly one unit for a price b^* when his valuation for this unit is in $(0, b^*)$. Thus, there is at most one bidder who bids on his first unit at b^* with positive probability.

If no bidder were to bid with positive probability at b^* on his first unit, then all bidders would bid with positive probability at zero on their first unit. Hence, such bids would win with positive probability. But in that case, bidders can gain from deviating to bidding their valuation. Thus, there is exactly one bidder, say bidder i, who bids on his first unit at b^* with positive probability.

The remaining bidders must bid zero on all of the units for which their valuation is less than b^*, for otherwise they run the risk of winning those units at prices above their valuations for those units. As a consequence bidder i bids b^* on all units for which his valuation is in the set $(0, b^*)$. $\square$

Proof of Proposition 2: We have shown in the main text that all strategies in the class S^2 are equilibria. It remains to show that for any equilibrium β in which no bidder bids in $(0, v^h)$ with positive probability, there is a strategy in S^2 that agrees with β for all i and almost all $\mathbf{v}^i$.

Suppose that, with positive probability, the number of bids at or above v^h is smaller than m. Then there is a bidder who bids zero for some unit and wins this unit with positive probability less than one. This bidder can raise his payoff by switching to always bidding his value.

Suppose that, with positive probability, the number of bids at or above v^h is greater than m. Then there exist a bidder who, with positive probability, wins one unit for a price greater or equal v^h. This bidder can raise his payoff by switching to always bidding his value.

This implies that the number of bids at or above v^h is equal to m with probability one. Since bids are independent across bidders, it follows that if bidder i bids at or above v^h for his j - th unit with positive probability, he must bid at or above v^h for his j - th unit with probability one. $\square$

Proof of Remark 3: Select any truthful equilibrium allocation of all units among the grand coalition N, and let $k^i(N)$ be the number of units player i gets in this allocation. Next, for every coalition $C \subset N$, select an efficient allocation of all units among the members of the coalition C, and let $k^i(C)$ be the number of units player i gets in this allocation. Then for any coalition C including the seller (i.e. in which $0 \in C$),

$$w(C) = \sum_{i \in C \backslash \{0\}} \sum_{j=1}^{k^i(C)} v_j^i.$$

If there are multiple efficient allocations for a given coalition C, the value of the coalition $w(C)$ is the same for all chosen efficient allocations. In such a case, we select an efficient allocaction in which bidder i receives $k^i(C) \geq k^i(N)$ units.

Let $(\pi^0, \ldots, \pi^n)$ be the profile of payoffs resulting from the truthful equilibrium. Since the truthful equilibrium is efficient, $\sum_{i \in N} \pi^i = w(N)$. Let p^i denote the payment of bidder i in the selected truthful equilibrium allocation. For any C including the seller, the coalition's payoff in the truthful equilibrium is

$$\begin{aligned} \sum_{i \in C} \pi^i &= \sum_{i \in C \backslash \{0\}} \left(\left(\sum_{j=1}^{k^i(N)} v_j^i \right) - p^i \right) + \sum_{i \in N \backslash \{0\}} p^i \\ &= \sum_{i \in C \backslash \{0\}} \left(\sum_{j=1}^{k^i(N)} v_j^i \right) + \sum_{i \in N \backslash C} p^i. \end{aligned}$$

Thus,

$$\begin{aligned} \sum_{i \in C} \pi^i - w(C) &= \sum_{i \in C \backslash \{0\}} \sum_{j=1}^{k^i(N)} v_j^i + \sum_{i \in N \backslash C} p^i - \sum_{i \in C \backslash \{0\}} \sum_{j=1}^{k^i(C)} v_j^i \\ &= \sum_{i \in N \backslash C} p^i - \sum_{i \in C \backslash \{0\}} \sum_{j=k^i(N)+1}^{k^i(C)} v_j^i, \end{aligned}$$

where the second equality uses $k^i(C) \geq k^i(N)$.

In the truthful equilibrium $\sum_{i \in C \backslash \{0\}} \sum_{j=k^i(N)+1}^{k^i(C)} v_j^i$ is the sum of the $m - \sum_{i \in C} k^i(N)$ highest losing bids from coalition C, which are outbid by bidders belonging to $N \backslash C$. Because in the truthful equilibrium i obtains $k^i(N)$ units, his payment p^i is the sum of the $k^i(N)$

highest losing bids from $N\backslash\{i\}$. Suppose that $i\in N\backslash C$. Then p^i is weakly higher than the sum of the $k^i(N)$ highest losing bids from C. Therefore, $\sum_{i\in N\backslash C}p^i$ is weakly higher than the highest $\sum_{i\in N\backslash C}k^i(N)$ losing bids from C. Since $\sum_{i\in N\backslash C}k^i(N)=m-\sum_{i\in C}k^i(N)$, this implies that

$$\sum_{i\in N\backslash C}p^i\geq\sum_{i\in C\backslash\{0\}}\sum_{j=k^i(N)+1}^{k^i(C)}v^i_j.$$

Thus, $\sum_{i\in C}\pi^i\geq w(C)$. $\square$

BIBLIOGRAPHY

Ausubel, L.M., Milgrom, P., (2002). "Ascending Auctions with Package Bidding" *Frontiers of Theoretical Economics 1*, Article 1.

Avery, C., Resnick, P. and Zeckhauser, R., (1999). "The Market for Evaluations" *The American Economic Review* Vol. 89, pp. 564–584.

Banerjee, A., (1992). "A Simple Model of Herd Behavior" *The Quarterly Journal of Economics*, Vol. 107, pp. 797–817

Bergemann, D. and Välimäki, J., (2000). "Experimentation in Markets" *Review of Economic Studies*, Vol. 67, pp. 213–234

Bikhchandani, S., Hirshleifer, D. and Welch, I., (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades" *Journal of Political Economy*, Vol. 100, pp. 992–1026

Blume, A., Heidhues, P., (2001). "All Equilibria of the Vickrey Auction" Working Paper. Retrieved from: http://www.wiwi.uni-bonn.de/heidhues/Paper/characterization2.PDF

Blume, A., Heidhues, P., (2004). "All Equilibria of the Vickrey Auction" *Journal of Economic Theory* Vol. 114, pp. 170–177.

Blume, A., Heidhues, P., (2008). "Modeling Tacit Collusion in Auctions" *Journal of Institutional and Theoretical Economics* Vol. 164, pp. 163–184.

Bolten, P. and Harris, C., (1999). "Strategic Experimentation" *Econometrica* Vol. 67, pp. 349–374

Bolton, G., Katok, E. and Ockenfels, A., (2004). "How Effective are Electronic Reputation Mechanisms? An Experimental Investigation." *Management Science* Vol. 50, pp. 1587–1602.

Cai, H., Chen, Y., and Fang, H., (2008). "Observational Learning: Evidence from a Randomized Natural Field Experiment" Forthcoming in *The American Economic Review*

Chen, Y., Harper, M., Konstan J., and Li, S.X., (2008). "Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens." Forthcoming in *The American Economic Review*

Chevalier, J. and Mayzlin, D., (2006). "The Effect of Word of Mouth on Sales: Online Book Reviews" *Journal of Marketing Research*, Vol. 43, pp. 345–354

Dellarocas, C., Awad, N.F., and Zhang, X., (2004). "Exploring the Value of Online Reviews to Organizations: Implications for Revenue Forecasting and Planning", Forthcoming in the *Journal of Interactive Marketing*

Duan, W., Gu, B., and Whinston, A.B., (2008). "Do Online Reviews Matter? - An Empirical Investigation of Panel Data" *Decision Support Systems*, Vol. 45, pp. 1007–1016.

Farrell, J. and Klemperer, P., (2006). "Coordination and Lock-In: Competition with Switching Costs and Network Effects" *UC Berkeley: Competition Policy Center*, Retrieved from: http:// www.escholarship.org/uc/item/9n26k7v1

Fehr, E. and Gächter, S., (2000). "Cooperation and Punishment in Public Goods Experiments" *The American Economic Review* Vol. 90, pp. 980-994

Fehr, E. and Schmidt, K.M., (1999). "A Theory Of Fairness, Competition, and Cooperation" *Quarterly Journal of Economics* Vol. 114, pp. 817–868

Fischbacher, U., (2007). z-Tree: Zurich Toolbox for Ready-made Economic Experiments, *Experimental Economics* Vol. 10, pp. 171–178.

Holzman, R., Monderer, D., (2004). Characterization of Ex Post Equilibrium in the VCG Combinatorial Auctions. *Games and Economic Behavior* Vol. 47, pp. 87–103.

Houser, D. and Wooders, J. (2005). "Reputation in Auctions: Theory, and Evidence from eBay" *Journal of Economics and Management Strategy*, Vol. 15, pp. 353–370.

King, S. (1995). "Search With Free-Riders" *Journal of Economic Behavior and Organization*, Vol. 26, pp. 253–271

Levine, David K (1998). Modeling Altruism and Spitefulness in Experiments, *Review of Economic Dynamics*, Vol. 1, pp. 593–622.

Li, Xinxin and Lorin Hitt (2008). "Self Selection and Information Role of Online Product Reviews" *Information Systems Research*, Vol. 19, pp. 456–474

Milgrom, P., (2004). "Putting Auction Theory to Work" *Cambridge University Press*, Cambridge, UK.

Milgrom, P., (1981). "Rational Expectations, Information Acquisition, and Competitive Bidding" *Econometrica* Vol. 49, pp. 921–943.

Milgrom, P., Weber, R.J., (1985). "Distributional Strategies for Games with Incomplete Information" *Mathematical Operations Research* Vol. 10, pp. 619–632.

Morris, S. and Song Shin, H. (2002). "Social Value of Public Information" *American Economic Review*, Vol. 92, pp. 1521–1534

Mulpuru, S. (2009). "US eCommerce Forecast, 2008 to 2013"
Forrester Research Report

Plum, M., (1992). "Characterization and Computation of Nash–Equilibria for Auctions with Incomplete Information" *International Journal of Game Theory* Vol. 20, pp. 393–418.

Satterthwaite, Mark A. (1979). "Consumer Information, Equilibrium Industry Price and the number of sellers" *Bell Journal of Economics*, Vol. 10, pp. 483–502

Scherer, F.M., Ross, D., (1990). "Industrial Market Structure and Economic Performance" *Houghton Mifflin Company*, Boston, MA.

Servátka, M. "Separating Reputation, Social Influence, and Identification Effects in a Dictator Game" forthcoming in *European Economic Review.*

Tan, G., Yilankaya, O., (2006). "Equilibria in Second Price Auctions with Participation Costs" *Journal of Economic Theory* Vol. 130, pp. 205–219.

Vickrey, W., (1961). "Counterspeculation, Auctions and Competitive Sealed Tenders" *Journal of Finance* Vol. 16, pp. 8–37.

Xiao, E. and Houser, D. (2008). "Emotion Expression in Human Punishment Behavior" *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 102, pp. 7398–7401