

New Institutional Structures on the Internet: The Economic Design of Online Auctions

Axel Ockenfels*

Max Planck Institute for Research into Economic Systems, Jena

prepared for:

M.J. Holler, H. Kliemt, D. Schmidtchen and M. Streit (eds.),
Jahrbuch für Neue Politische Ökonomie, 20,
Tübingen: Mohr Siebeck, forthcoming.

March 1, 2002

Abstract

The internet facilitates interactive and asynchronous communication among human and artificial agents in large bilateral markets. As a consequence, new institutional structures are emerging in the new economy. I discuss a few core problems associated with the economic design of internet markets and the way they are approached by eBay.com, a popular online auction platform.

* Max Planck Institute for Research into Economic Systems, Strategic Interaction Unit, Kahlaische Straße 10, D-07745 Jena, Germany; ockenfels@mpiew-jena.mpg.de. I gratefully acknowledge the helpful comments of Daniela Dohmes-Ockenfels and Werner Güth as well as the support of the Deutsche Forschungsgemeinschaft through the Emmy Noether-Program.

I. Introduction

The new economy's electronic communication channels reduce search costs, asymmetric information and coordination problems. This makes internet markets more efficient than conventional markets with respect to transaction costs and price levels. But to look at internet markets just as low friction-versions of traditional markets is misleading. The new online technologies imply that market structures that have been proven to perform well in the old economy can be inappropriate in the internet. The availability of low-cost asynchronous and interactive trading mechanisms in large, anonymous person-to-person markets leads to new institutional solutions that can be substantially different from standard offline market structures.¹

The most successful internet market to date is eBay.com. eBay is a person-to-person internet auction market where individuals and firms can participate both as buyers and sellers. eBay only provides the auction platform including all communication channels, but does not act as seller or buyer.² Profits are made through two kinds of selling fees. The first is an insertion fee between 30 cents and \$3.30, depending on the amount of the (secret or public) reserve price. The second is a final sale price fee that generally ranges from 1.25 to 5 percent of the final price. According to auctionwatch.com, a rich source of online auction-related information, eBay's consolidated net income for the second quarter 2001 was \$24.6 million and its gross merchandise sales amounted to \$2.25 billion. Furthermore, eBay's community included 34 million registered users who submitted about 20 bids per second to the millions of auctions running at each point of time.

eBay's astonishing success in bringing traders together from all over the world is partly based on innovative details of the market design that cannot easily be found in conventional markets of the old economy. The next sections analyze three particularly interesting design features, namely automatic bidding agents that make interactive and asynchronous price discovery possible at the same time, a simple fixed auction deadline ('hard close') as a way to end asynchronous bidding, and a sophisticated computerized reputation mechanism implemented to promote trust among strangers. All features are

¹ While many studies are concerned with how the new economy challenges the agents' behavioral *strategies* developed in the old economy (see Shapiro and Varian 1998), the economic analysis of how the new economy challenges market institutions has just started.

designed to solve problems that naturally arise in large electronic markets characterized by asynchronous interactive trading. Thus, eBay’s design solutions and their performances are of interest not only to internet auction designers.

II. eBay’s automatic bidding agent

II.1 The advent of second price auctions

eBay asks its bidders to submit a ‘maximum bid’ and explains: “eBay will bid incrementally on your behalf up to your maximum bid, which is kept secret from other eBay users.” That is, each bidder is asked to decide the maximum he is willing to pay and to enter this amount. eBay’s automatic bidding agent then confidentially bids up to this amount. The automatic agent overbids earlier or subsequently incoming bids of others – not revealing the amount of the maximum bid – as long as the price does not exceed the bidder’s maximum bid. If the price exceeds the bidder’s maximum bid, eBay’s automatic agent automatically exits the auction on behalf of the bidder.

maximum bid		high bidder	price (increment = \$0.50)
bidder 1	bidder 2		
\$5	\$10	2	\$5.50
\$8	no change	2	\$8.50
\$12	no change	1	\$10.50

Table 1. Bidding example

Table 1 shows a fictitious bidding history of an eBay auction with two bidders. Bidder 1 submits a maximum bid of \$5 and bidder 2 bids \$10 early in the auction. Bidder 2’s automatic bidding agent raises the price, known to all bidders at each point of time, to \$5.50, one increment above the second highest submitted bid. Typically, eBay auctions last a week, during which bidders can adjust their maximum bid. In our example, bidder 1, not knowing bidder 2’s maximum bid but only the current price, decides to increase his

² eBay is often referred to as C2C-market, which is, however, misleading since also many businesses and professional dealers choose to offer their goods on eBay.

maximum bid to \$8 after being outbid by bidder 2's automatic bidding agent.³ This raises the price to \$8.50 but leaves bidder 2's high bidder status unchanged. Then, bidder 2 enters \$12 and becomes the high bidder.

At the end of any eBay auction, the winning bid is the highest maximum bid, but the final price paid by the winner is a small increment above the highest submitted bid by the *competing* bidders (or the seller's reservation price if the winner is the only active bidder). In the example above, bidder 2 wins and pays \$10.50 for the item. Of course, the outcome would have been unaffected if bidder 2's last bid had been even higher, say \$100, if the highest submitted bid by the loser is not changed.⁴ Hence, by using the automatic bidding agent, eBay employs what economists call a second price auction.

While many observers of eBay seem to believe that eBay invented this auction format, second price auctions were known long before the first internet auctions emerged. In his seminal work, Vickrey (1961) was the first to study second price auctions from an economic theory point of view. In 1996 he received the Nobel prize for his work, in which he showed that a sealed bid second price auction is, under certain conditions, strategically equivalent to the English ascending price auction. But while Vickrey discusses some real world applications of a couple of other auction formats, he does not indicate that (sealed bid) second price auctions, now also known as "Vickrey auctions", have ever been actually used. And in fact second price auctions were rarely used before the emergence of online trading. Moldovanu and Tietzel (1998) discovered that Johann Wolfgang von Goethe designed a second price auction in 1797 to elicit his publisher's willingness to pay for a manuscript. Lucking-Reiley (2000) discovered that some stamp auctioneers used second price rules as early as the nineteenth century. Rothkopf et al. (1990) mention a few more recent examples, but the title of their article, "Why are Vickrey auctions so rare?", is an unambiguously sign that second price auctions never really kicked off. At first glance, this is surprising because second price auctions have economically advantageous properties. In particular, as shown by Vickrey (1961) for the

³ Each submitted maximum bid must exceed both, one's earlier maximum bids, if any, and the current price.

⁴ The price can never exceed the high bidder's maximum bid. If the difference between the highest and the second highest maximum bid is smaller than the minimum increment, the price equals the highest maximum bid. If more than one bidder submitted the maximum bid, the bidder who submitted his bid first

case of independent private values environments and risk neutral bidders, sealed bid second price auctions have dominant truth-revealing equilibrium strategies, lead to economic efficiency, and produce the same expected revenue as more popular but also behaviorally more demanding auction formats, namely English, Dutch and sealed first price auctions.

Now, only twelve years after the work of Rothkopf et al. (1990) was published and seven years after eBay was founded in 1995, second price auctions are the most prevalent auction form. Each week, millions of second price auctions are conducted not only on eBay but also on the other major person-to-person online auction houses such as Yahoo and Amazon. So what is it that makes second price auctions so unattractive in the old economy and so appealing in the new economy?

II.2 Disincentives for bidders to follow truth-revealing strategies and fear of cheating in the old and the new economy

Based on theoretical and empirical findings, Rothkopf et al. (1990) highlight two explanations for the virtual nonexistence of second price auctions: disincentives for bidders to follow truth-revealing strategies and fear of cheating. The former refers to the observation that many traditional auctions are not isolated events as typically assumed by the underlying economic theory. Since a truth revealing strategy may give away valuable information to potential competitors, or to others with whom the bidder must subsequently interact or negotiate specific contract terms, bidders are reluctant to reveal their private information. As a direct consequence, the properties that make second price auctions advantageous break down. This argument loses some of its strength, however, in online auction houses such as eBay. There, the winning bid is kept secret, that is, neither the competing bidders nor the seller are informed about the highest maximum bid. Furthermore, field studies indicate that most interactions on eBay are in fact one-shot in nature (see Resnick and Zeckhauser 2001 and Section IV of this paper). Nevertheless, to the extent bidders do interact repeatedly, there are still disincentives to truthfully reveal

is the high bidder at a price equal to the maximum bid. So, in Table 1, if bidder 2 had entered \$10, bidder 1 would have been the winner at a price of \$10.

one's private value, because all bids but the highest bid are actually revealed to the public.

The second explanation put forward by Rothkopf et al. (1990), fear of cheating, refers to potential price-enhancing activities by sellers that can be a bidders' concern. As Rothkopf et al. (1990, 102) put it : "the use by auctioneers of shills and imaginary bids to force the price above the second-highest bidder's value is notorious (Cassady 1967, chap. 12).", and "[if the bidder] fears that the bid taker, after observing his price, will insert an imaginary losing bid or a real losing bid from a confederate, then he has an incentive to bid strategically." Interestingly, Vickrey (1961) was well aware of the danger of cheating activities in second price auctions and recommended the use of neutral and trustworthy intermediaries. He writes:

"To prevent the use of a 'shill' to jack the price up by putting in a late bid just under the top bid, it would probably be desirable to have all bids delivered to and certified by a trustworthy holder, who would then deliver all bids simultaneously to the seller. Under these circumstances, the seller would have no incentive to do other than sell to the top bidder, showing him as his price the second-best bid." Vickrey (1961, p. 22-23)

eBay's electronic auction format and in particular the use of automatic auctioneers seem to make collusion between auctioneers and sellers unlikely.^{5,6} However, the sellers or their confederates may inflate the second highest bid and thus the final price by placing shill bids themselves. The risk of overbidding the 'serious' bidders by shill bids and thus not selling the item is rather low, because bids can sometimes be retracted,⁷ or the item can be relisted later. So, not surprisingly, shill bidding is in fact identified a serious problem by eBay, and it is explicitly forbidden:

⁵ But note that eBay's profit positively depends on final prices so that eBay too is in principle interested in high final prices.

⁶ Also, it is difficult to collude in second price auctions with a large number of potential bidders, each only with a small probability of having a value above the reservation price. The costs of coordinating a large bidding ring may well exceed the benefits (see Mailath and Zemsky, 1991, and the references cited therein for the theory of collusion in second price auctions).

⁷ In principle every bid on eBay is binding. But there are a number of 'exceptional circumstances,' for which eBay allows to retract bids. These circumstances include if one 'accidentally' enters a wrong bid amount, if the description of an item one has bid on has changed significantly, if one cannot reach the seller, etc.

“Shill Bidding - The deliberate use of second User IDs or other eBay members to artificially raise the level of bidding and/or price of an item. Please remember that family members or housemates should not be bidding on each other's auctions. This activity is not permitted. [...] If you suspect shill bidding, please report it to eBay here: [...]” eBay.com (2001)

Neither eBay nor other online auction houses report frequencies of detected shill bidding activities. However, in 1999 eBay developed an automatic shill bidding detection tool as a way to analyze bidding patterns over multiple listings. Also, there is a market for software programs that allow eBay users to quickly review a seller's feedback history (see Section IV) and check for patterns of shill bidding.⁸ The broad use of such ‘shill catcher tools’ indicates that shill bidding is no minor problem, even though detected shill bidders are suspended from eBay and must fear criminal prosecution.⁹

II.3 Second price auctions and asynchronous interaction

The discussion in the last section indicates that both arguments put forward by Rothkopf et al. (1990) for the rare use of second price auctions in the old economy are also more or less valid in the new economy. But second price auctions have one major advantage over other auction formats that makes them particularly suitable for large markets in the new economy: second price auctions at the same time allow asynchronous and interactive price discovery. Observe that in game theoretic terms a maximum bid submitted to eBay is a whole strategy (rather than an action) that calls for dynamic execution as other bids come in. That is, the automatic bidding agents bid on their bidders’ behalf, instantaneously calculating current prices and current high bidders, and they therefore free bidders from the necessity to follow the auctions and the price discovery process themselves. As eBay (2001) puts it: “In this way, you don't have to keep an eye on your auction as it unfolds. Remember: eBay will use only as much of your

⁸ Suspicious bidding patterns may include a bidder who exclusively bids on one seller's auctions and almost never wins.

⁹ In March 2001, three bidders who created multiple fake user IDs were accused of tricking other eBay users into believing that the shill bids they placed on each other’s items were legitimate in many hundred auctions. The total value of the winning bids in all these auctions hosted and manipulated by these bidders was supposed to exceeded \$450,000 with a total value of the high shill bids in these auctions of about \$300,000. The bidders paid back \$100,000 after pleading guilty to bidding for their own products on eBay. (see www.cybercrime.gov/Fetterman_indict.htm, and www.theregister.co.uk/content/archive/18354.html).

maximum bid as is necessary to maintain your position as high bidder. Winning was never easier!”

The old economy typically provides mechanisms that allow for *either* asynchronous *or* interactive communication. Traditional auction formats such as the English auction, for instance, require bidders to participate at the same time (and, before the emergence of modern communication channels, at the same place), tying each bidder up for the entire length of the auction. This and other formats that require ‘synchronous’ activities would obviously fail to support large internet markets, such as eBay with 34 million registered users living all over the world. On the other hand, conventional market institutions that secure asynchronous communication in large markets, such as fixed (or catalog) pricing, typically rob the internet marketplace of one of its most distinctive characteristics, the possibility of *interactive* price discovery (sometimes called ‘dynamic pricing’) allowing each market participant to dynamically respond to the activities of other market participants. Such multi-directional interaction is possible, even in very large internet markets, because the electronic online communication channels do not only allow faster speed, broader reach and lower information cost, but they also permit information to flow from and to all market participants, and not only from sellers to buyers as in standard fixed pricing schemes.¹⁰ Given the technological environment of the internet, auctions are often regarded as superior trading mechanisms than, say, fixed pricing mechanisms, because the bidders’ *interaction* in a competitive environment aggregates and effectively reveals private information and therefore facilitates efficient pricing.¹¹ This holds true especially in markets for items that have no clear market value, such as antiques, collectibles and used items as also sold on eBay. (eBay is sometimes called the world’s largest garage sale.) Thus, eBay’s automatic bidding agent allows asynchronous interaction among large numbers of artificial and human buyers and sellers in fluid price

¹⁰ In an interesting case of dynamic pricing, Amazon.com, until then thought of as a fixed price retailer, has faced allegations that it has charged different consumers different prices for the same items based on information about the shoppers’ buying habits such as billing information and purchasing history. However, Amazon.com denied all allegations, and said that the price dispersion was caused by random pricing-tests. Nevertheless, the angry customers were refunded (see USA Today, September 29, 2000; <http://www.usatoday.com/life/cyber/tech/cti595.htm>).

¹¹ It goes beyond the scope of this paper to describe the merits and the superiority of auctions as price discovery mechanisms in many situations, but see McAfee and McMillan (1987) for a superb survey of auction theory, McAfee and McMillan (1999) for a short non-technical introduction, and Klemperer (2000) for an overview of the many applications of auction theory in economics.

environments so that global trade and communication, not bounded by time or space constraints, are possible and price discovery is easy and efficient. This seems to explain the overwhelming use of the second price method in C2C internet auction markets.

III. How to end asynchronous interactive communication?

III.1 The choices made by eBay and Amazon

One question in auction design is whether to employ an open auction format, where the bidding activity is (partly) made public as bids come in, or sealed bid auctions, where market participants are not informed about the competitors' bids before the auction is over. While in principle sealed bid second price auctions are compatible with asynchronous interaction through the use of programmed artificial agents, open auctions maintain two advantages in connection with *interactive* trading not available in sealed bid auctions. First, from a theoretical point of view, open auctions tend to reduce the force of the 'winner's curse' because the competitors' bidding activities may convey relevant information that the bidders can then use in revising their estimates of value. Thus, uncertainty is reduced and so is the winner's curse, and bidders can safely bid more aggressively which typically results in higher revenues in open auctions (Milgrom and Weber 1982). Second, it has been shown experimentally that the feedback delivered in open second price auctions accelerates the speed of learning (see Ariely et al. 2002). This improves the price discovery process and increases the competition among bidders so that efficiency and revenues can be enhanced, even in private value environments.¹²

Given the decision to employ an open auction format, the question arises how to end the auctions. Traditionally, auctions are driven to conclusion by the auctioneers ("Going, Going, Gone ..."). If human auctioneers are too costly or not practical, such as in large internet markets with millions of simultaneously ongoing auctions or in the complex U.S. spectrum auctions, new ending rules must be devised. For instance, Milgrom and Wilson designed the 'activity rule' that was applied to the U.S. spectrum auctions (McAfee and McMillan 1996). The activity rule requires a bidder to be 'active' (that is to be a high

¹² See also Cramton (1998) who compares the merits of open and sealed bid auctions. While he concludes that the dynamic price discovery process of an open auction most often does a better job, he also emphasizes that two aspects may favor sealed bidding, namely *ex ante* asymmetries and weak competition. Both aspects, however, appear to be of only minor importance on eBay.

bidder or to submit new bids at certain times) on a predetermined number of spectrum licenses. If a bidder falls short of the required activity level, the number of licenses it is eligible to buy shrinks. Thus, bidders are prevented from holding back so that the auction can be concluded within a reasonable time span, when no one wants to continue bidding. Unfortunately, asynchronous interaction is incompatible with such activity rules because they require bidders to be present at certain times. A new online-specific design is needed.

The simplest method to end online-auctions is to determine a fixed end (sometimes called a ‘hard close’). That is, all market participants know exactly, from the very beginning of an auction up to the second, at what time the auction will end. This is how eBay ends its auctions. Amazon.com, on the other hand, created its own auction platform apparently using eBay’s design as a model but deviated from eBay’s model with respect to the rule for ending the auction. Amazon attempts to imitate the ‘Going, Going, Gone’ feature of traditional auction houses and employs a ‘soft close’. That is, Amazon automatically extends an auction if a bid comes in late, so that all bidders always have the opportunity to respond to the opponents’ bids. In Amazon’s own words:

“We know that bidding can get hot and heavy near the end of many auctions. Our Going, Going, Gone feature ensures that you always have an opportunity to challenge last-second bids. Here's how it works: Whenever a bid is cast in the last 10 minutes of an auction, the auction is automatically extended for an additional 10 minutes from the time of the latest bid. This ensures that an auction can't close until 10 minutes have passed with no further bids.”

Amazon.com (2001)

III.2 Some theory and empirics: Details matter!

Why did Amazon adopt virtually all rules from eBay with the exception of the ending rule? The quote above suggests that the answer has something to do with the observation that bidding gets ‘hot’ at the end and that bidders should always have the opportunity to respond to each other, which is not guaranteed by a fixed deadline. Also, if bidders cannot challenge last-second bids, revenues are supposedly smaller and economic efficiency might be in danger. But why should someone want to bid late at all? Both on eBay and Amazon, a bid wins (at a price that equals the highest bid of the competing bidders plus a small increment) if and only if the bid exceeds the final price, *regardless of*

the timing of the bid. This is why eBay urges its bidders not to wait until the closing seconds of the auction before submitting the bid, and this is why the second price method is a useful method to promote asynchronous bidding as explained in Section II.

However, as shown by Ockenfels and Roth (2001), a fixed deadline as applied by eBay creates incentives to bid very late, a practice often called ‘sniping’. The reason is that late bids can avoid early bidding wars and therefore favorably affect the price paid by the winning bidder. Suppose for example that there are ‘incremental bidders’. Incremental bidders start with a low bid, well below their true willingness to pay, and then do nothing unless they are outbid in which case they incrementally increase their bid in order to regain the status of the high bidder. Bidding in the closing seconds of an auction might be a best response to this kind of incremental bidding because bidding very near the deadline of the auction would not give the incremental bidder sufficient time to respond to being outbid. By bidding at the last moment one might win the auction at the incremental bidder’s initial, low bid, even though the incremental bidder's willingness to pay exceeds your willingness to pay.

eBay Bid History for

SONY VAIO PCG-XG700K Notebook Computer NEW NR (Item # [1242933367](#))

Currently	\$2,050.00	First bid	\$9.99
Quantity	1	# of bids	37
Time left	Auction has ended.		
Started	Jun-01-01 17:37:30 PDT		
Ends	Jun-06-01 17:37:30 PDT		
Seller (Rating)	bestprices4all (871) ★ me		

[View page with email addresses](#) (Accessible by Seller only) [Learn more.](#)

Bidding History (Highest bids first)

User ID	Bid Amount	Date of Bid
lambrosa (0) 🍷	\$2,050.00	Jun-06-01 17:34:35 PDT
boydragon (6)	\$2,025.00	Jun-06-01 17:36:25 PDT
boydragon (6)	\$1,975.00	Jun-06-01 17:35:03 PDT
boydragon (6)	\$1,925.00	Jun-06-01 17:24:49 PDT

Figure 1. Late bidding to avoid bidding wars

The screen shot in Figure 1 shows the bid history of an actual eBay auction of a notebook that ended on June 6, 2001 at 17:37:30 PDT. The history reveals that until 17:34:35 of the last day of the auction boydragon was the high bidder. Then, lambrosa outbid boydragon, leaving boydragon 175 seconds to respond. Boydragon responded immediately and managed to submit two more bids both of which were outbid immediately by lambrosa's automatic bidding agent. Boydragon's apparent attempt to regain his high bidder status failed. By the same token, it seems likely that lambrosa could have paid less for the notebook by bidding even closer to the end not leaving boydragon enough time to drive up the price by \$100.¹³

Incremental bidding may be caused by bidders who are not aware or who do not want to make use of the automatic bidding agent, and rather bid 'naively' as if they were in an ascending first price auction. In fact, field (Ockenfels and Roth 2001a) and lab (Ariely et al. 2001) evidence suggests that incremental bidding decreases with experience indicating that late bidding is likely to arise in part as a response to naïve incremental bidding. Incremental bids may, however, also occur as a rational reply in common value auction contexts, in which the bidding activity of others can convey some valuable information that may cause bidders to adjust their bids over time. As a consequence, 'informed' bidders may try to protect their information by bidding so late that it is too late for 'uninformed' bidders to incorporate this information into their bids (see Ockenfels and Roth 2001a for a game theoretic model). But multiple bids and late bidding may even occur in private value auction equilibria.¹⁴ As Ockenfels and Roth (2001a) show, multiple and late bids are consistent with implicit collusion in private value equilibria, because very late bids run the risk of not being successfully submitted.¹⁵ On Amazon, on the other hand, the incentives to bid late are eliminated or severely attenuated (for proofs within a game-theoretic model see Ockenfels and Roth 2001a). For instance, on Amazon an attentive incremental bidder in a private value auction or the

¹³ The price immediately after lambrosa submitted his bid was \$1,950, one increment (\$25 in this case) above the second highest submitted maximum bid at this time (\$1,925); see Ockenfels and Roth (2001b) for more examples.

¹⁴ Peters and Severinov (2001) have recently shown, that incremental bidding may be a symmetric equilibrium strategy in a multi-unit auction environment if buyers hope to acquire a single unit only. They do not explain the tendency to bid late, however.

uninformed bidder in a common value auction can be provoked to respond whenever a bid is placed.

Field and lab data support the theoretical considerations. In a field study of the timing of bids on eBay and Amazon auctions, Roth and Ockenfels (2001) found that there is significantly more late bidding on eBay than on Amazon. For instance, 40 percent of eBay-Computers auctions and 59 percent of eBay-Antiques auctions in their sample have last bids in the last 5 minutes, as compared to about 3 percent of both Amazon-Computers and Amazon-Antiques auctions that have last bids in the last five minutes before the initially scheduled deadline or later. Also, there is significantly more late bidding in antiques auctions than in computer auctions on eBay (but not on Amazon), suggesting that the expertise effect in auctions with interdependent values is actually reflected in the timing of bids. Further analyses by Ockenfels and Roth (2001a) reveal that the impact of the bidders' feedback numbers on late bidding is significantly positive on eBay and negative on Amazon. This suggests that more experienced bidders on eBay bid later than less experienced bidders, while experience on Amazon has the opposite effect – another indication that the differences in the *strategic* environments determined by the ending rules drive the differences in the timing of bids. Experiments of Ariely et al. (2001) reproduce the differences of the timing of bids and of the experience effect in a laboratory setting in which the only difference between auctions is the ending rule. In particular, subjects are randomly assigned to different auction conditions, and bid for identical, artificial commodities for which they are paid in cash by the experimenter according to values that they know when they bid. The experiment thus controls for differences other than the closing rule that might affect behavior on Amazon and eBay, such as the number of auctions being conducted at any time, market size or self selection effects that might have biased the field results.

Further studies such as a survey conducted by Ockenfels and Roth (2001a) and a computational study based on genetic algorithms by Ünver (2001) yield additional support to the theoretical considerations, so that virtually all complementary empirical and theoretical methods point in the same direction: a fixed deadline induces strategic

¹⁵ Because the time it takes to place a bid may vary considerably due to, for instance, erratic internet traffic or connection times, last-minute bids have a positive probability of being lost. See Ockenfels and Roth (2001a) for survey and anecdotic evidence.

incentives to bid late, while an automatic extension rule removes these incentives. This, however, implies that bidders want to be in front of their computer screens towards the end of eBay auctions, so that the activities become synchronized. As a consequence, the advantages of the second price method discussed in Section II are partly removed by the hard close! Fortunately, however, this is not the end of the story. New artificial bidding agents are emerging as a response to the dilemma posed by eBay's fixed deadline rule. Esnipe.com, for instance, offers its users to place bids on their behalf during the last few seconds of the eBay auction, so that bidders do not need to be in front of their monitors when the auction ends. Esnipe (2001) nicely summarizes the merits, risks and the prevalence of late bids as follows:

“eSnipe reduces bidding wars by masking interest in auction items until the last possible moment. Because auctions on eBay take three to 10 days to close, emotional overbidding can start soon after the auction opens and last until the auction closes, instead of mere minutes as is the case in traditional offline or “outcry” auctions. eSnipe protects bidders’ privacy. For example, eSnipe is popular among buyers of collectibles, who treasure the anonymity it brings them. Even on the Web, news travels fast in communities of collectors, who prefer to avoid triggering interest in an item by bidding on it too early. eSnipe allows users to place bids when it’s convenient for them. Overseas users, for example, do not have to stay up to the middle of the night just to place a bid. Other users like eSnipe because it allows bids to be “canceled” up until five minutes before the auction close, whereas eBay’s terms of agreement treat any bid placement as a contract to buy. Manual sniping has been widely practiced for years, often by auction enthusiasts who keep multiple browser windows running with one hand and stopwatch in the other. A class of utilities that ran on the users’ machines started the automated sniping trend, but they weren’t reliable because of inaccurate system clocks, unpredictable computer performance, and outright crashes at inopportune times. [...] eSnipe is the clear market leader with 50,000 registered users, placing more than 10,000 bids a day that average about \$16 million per month.” <http://www.esnipe.com/about.asp>

III.3 A related design issues in online negotiations

The problem how to drive asynchronous interaction to conclusion is not unique to electronic auction sites. Similar problems arise in the design of online negotiation sites that promise dispute resolution (e-commerce disputes and traditional litigation) via electronic and standardized communication between the plaintiff and the defendant. These sites have been a popular alternative to lawsuits to resolve disputes, especially

those involving small amounts of money and conflicts between state and international laws.

One of the most successful online negotiation sites, [clicknsettle.com](http://www.clicknsettle.com),¹⁶ reports that it experimented with different ways of pushing the involved parties to an agreement. First, in 1999, one of their negotiation models was based on round-by-round demands and offers. In particular, the claimant and defendant would each enter three separate and distinct settlement offers and demands. These figures were then compared round by round, that is defendant offer number one versus claimant demand number one etc. But this format did not prove to be effective, because a deadline effect similar to what has been observed on eBay and to what has been observed in experimental bargaining games (Roth et al. 1988) hindered efficient negotiations: “After reviewing the early results with our clients, we discovered that in most negotiations, the first two rounds were being ‘wasted’ and the disputing parties really only had one opportunity to settle the case, the final round.” (<http://www.clicknsettle.com/onlinenegotiation.cfm>). Clicknsettle.com then decided to change the negotiation tool in several respects to avoid the end effect. First, they dropped the round-by-round negotiation and implemented the possibility to negotiate the settlement in continuous time, over a maximum of 60 days allowing an unlimited number of settlement bids. Each time a new bid is submitted by either party, it is compared to the last amount input by their adversary and a settlement calculation is performed. Second, [clicknsettle.com](http://www.clicknsettle.com) requires the parties to move 5 percent towards their adversary with each new settlement bid. To further push the speed of ‘price discovery’ both sides are paying a nominal fee each time they bid. Research is under way to examine the merits of these changes, which signify that deadline effects can pose serious problems in many interactive online exchange outside auctions.

IV. How to promote trust among strangers?

IV.1 eBay’s feedback forum

eBay cheaply connects millions of traders all over the world. But because most transaction partners are strangers, living in different cities or even countries, and because paying is typically required before the buyer can inspect or receive the item, there is a

¹⁶ Other sites include squaretrade.com, cybersettle.com or splitthedifference.com.

substantial risk involved in almost any transaction taking place on eBay. Sellers have incentives not to send the item to the buyers or to send poorer quality once they have received the money, and thus buyers must fear that their trust in sellers will not be reciprocated.

Internet auction markets face the challenge to encourage sellers to be trustworthy (to avoid moral hazard), and to discourage defective sellers (to avoid adverse selection) in an enormously large, anonymous, and spatially dispersed person-to-person market. But because eBay has no analogue in the old economy, 'offline institutions' give little guidance to solve the problem. Unlike in traditional auction houses such as Christie's and Sotheby's, none of the eBay transactions are secured, that is eBay neither attests about the sellers nor does it guarantee the items' qualities. Also, the distance in space between buyers and sellers does typically not allow for inspection an item's quality before the transaction takes place. Furthermore, reputation based on repeated interaction with the same seller cannot support the emergence of trust, because on eBay a buyer will rarely interact again with the same seller (see Resnick and Zeckhauser 2001). Finally, selling on eBay is typically cheaper than an ad in the local newspaper, so that signaling trustworthiness through 'fancy stores' is also not possible.

Very early on, eBay recognized the trust problem as its 'core challenge' and created the ingenious *feedback forum* that is copied by many other C2C online auction houses.¹⁷ It is based on the idea that the market participants continuously assess and reassess their transaction partners. Each user's assessment is made public and so forms his individual reputation that others can condition their behavior on. In this sense, eBay imitates the informal and unstructured disseminating of reputations via rumors and word of mouth in the old economy by a computer network that gathers information about past behavior swiftly and systematically, and distributes it to a broad community. In short, as eBay's founder Pierre Omidyar writes in a letter to all users:




¹⁷ The feedback forum is the most important and visible piece of a whole program to make online trading on eBay safer. Other mechanisms include internal investigations in case a user believes someone is misusing eBay, escrow services in case a users wants to involve a third party in a transaction to handle the exchange of money, verifying programs in case a user voluntarily provides identity information that will be verified by a third party for validation, fraud protection that covers most items free of charge to all eBay users 'in good standing' up to \$200 (certain restrictions apply), grading, authentication, verification, and appraisal services, in case a user demands an expert second opinion on an item, and dispute resolution by an independent, neutral third party to help resolve disputes that may arise.

“By creating an open market that encourages honest dealings, I hope to make it easier to conduct business with strangers over the net. Most people are honest. And they mean well. Some people go out of their way to make things right. I've heard great stories about the honesty of people here. But some people are dishonest. Or deceptive. [...] It's a fact of life. But here, those people can't hide. We'll drive them away. Protect others from them. This grand hope depends on your active participation. Become a registered user. Use our Feedback Forum. Give praise where it is due; make complaints where appropriate.”

<http://pages.ebay.com/community/news/founders-letter.html>

Figure 2 demonstrates with the help of a randomly chosen eBay auction how the feedback forum works.

The screenshot shows the eBay website interface for an auction. At the top, there is the eBay logo and navigation links: home, my eBay, site map, sign in, Browse, Sell, Services, Search, Help, Community, and item view. The main heading is 'Antique Hand Line Holder For Ice Fishing' with item number #1123520339. The category is 'Collectibles/Fishing/General'. The current bid is \$5.50, with 1 quantity and 1 minute and 42 seconds left. The first bid was \$1.00. The auction started on Mar-15-01 and ends on Mar-22-01. The seller is 'treebeard-in-maine' with a rating of 47. The high bidder is 'stromberg@webtv.net' with a rating of 51. Payment options include Money Order/Cashiers Checks. Shipping is fixed. The page also features 'eBay Online Payment' with VISA and MasterCard logos and an 'Instant' purchase option.

Antique Hand Line Holder For Ice Fishing		Item #1123520339		
Category: Collectibles/Fishing/General				
	Currently	\$5.50	First bid	\$1.00
	Quantity	1	# of bids	7 bid history
	Time left	1 min, 42 secs	Location	Wells, Maine
	Started	Mar-15-01 07:44:42 PST	Country	USA
	Ends	Mar-22-01 07:44:42 PST	mail this auction to a friend	
	Seller (Rating)	treebeard-in-maine (47) ★		
		view comments in seller's Feedback Profile view seller's other auctions ask seller a question		
	High bid	stromberg@webtv.net (51) ★		
	Payment	Money Order/Cashiers Checks		
	Shipping	Buyer pays fixed shipping charges, Will ship to United States only, See item description for shipping charges		
		 Instant This purchase cc		

Panel 1

home | my eBay | site map | sign in

ebay™

Browse Sell **Services** Search Help Community

overview registration buying & selling my eBay about me feedback forum safe harbor

Check out [LIVE](#) auctions on eBay. Search [tips](#)

Search titles and descriptions

Overall profile makeup

53 **positives**. 47 are from unique users and count toward the final rating.

0 **neutrals**. 0 are from users [no longer registered](#).

0 **negatives**. 0 are from unique users and count toward the final rating.

ebay ID card [treebeard-in-maine](#) (47)

Member since Friday, Mar 31, 2000 ★

Summary of Most Recent Comments

	Past 7 days	Past month	Past 6 mo.
Positive	4	13	53
Neutral	0	0	0
Negative	0	0	0
Total	4	13	53
Bid Retractions	0	0	0

[Auctions](#) by treebeard-in-maine

You can [leave feedback](#) for this user. Visit the [Feedback Forum](#) for more info on feedback profiles.

Panel 2

Items 1-25 of 53 total

= 1 = [\[2\]](#) [\[3\]](#) [\(next page\)](#)

User: oldfolksgoingbroke (90) ★ Date: Mar-21-01 17:58:39 PST
Praise: Received today - nice plane - little clean up and it will shine!!! Thank you!
User: joyeux1 (6) Date: Mar-21-01 06:32:41 PST
Praise: very nice people to work with! I highly recommend them
User: sfurney (1) Date: Mar-20-01 19:19:20 PST
Praise: Excellent description, accurate. Fast shipment- highly recommend!
User: zion@ak.net (135) ★ Date: Mar-17-01 09:15:35 PST
Praise: Thank you for payment! A pleasure doing business with!
User: prof1941 (94) ★ Date: Mar-14-01 15:57:27 PST
Praise: This item is a gem. Thanks for your service and help.
User: keila@cass.net (110) ★ Date: Mar-05-01 08:29:43 PST
Praise: Good transaction
User: marmartinez1875 (98) ★ Date: Mar-02-01 12:47:51 PST
Praise: Splendid MA Doll & 2nd transaction with Seller. Highly Reccoment!!!!

Panel 3

Figure 2. eBay’s feedback forum

Suppose a bidder wants to buy an “antique hand line holder for ice fishing” as offered by the seller “treeberad-in-maine” in Panel 1 of Figure 2. The seller’s reputation is measured by his “feedback rating” that is shown in parentheses next to each market

participant's user ID. The rating is the sum of feedback points. At the end of each transaction, the winning bidder and the seller in an eBay auction can give each other feedback points, + 1 point for a positive comment, 0 points for a neutral comment, and – 1 point for a negative comment.¹⁸ Thus, *treeberad-in-maine*'s feedback rating of 47 indicates that the seller has received a balance of 47 feedback comments from other eBay members.

A click on the feedback rating opens up a new window with a more detailed feedback profile (Panel 2 in Figure 2). This profile disaggregates the feedback rating with respect to both (algebraic) signs and over time. For instance, *treeberad-in-maine* received 4 positive feedbacks within the last seven days, and no negative or neutral feedbacks at all up to now. On an even more disaggregated level, each user has also access to the complete feedback history including the user IDs of the feedback givers, the exact time of the feedback, and a short comment by the feedback giver (such as “Excellent description, accurate. Fast shipment – highly recommended!”; see Panel 3 in Figure 2).

IV.2 Some theory and empirics

eBay interprets feedback as individual reputation: “A user's feedback is a key factor people use to determine whether or not they want to trade with that user. What feedback you give or receive is an important part of your trading reputation at eBay. [...] If you're a buyer, checking a seller's Feedback Profile before you make a bid is one of the smartest and safest moves you can make. This Feedback Profile answers many questions about how a seller does business. Is she highly recommended by other buyers? Does he sell quality merchandise?” (<http://pages.ebay.com/services/forum/feedback.html>). But can the mechanism really support trust among strangers?

A couple of theoretical papers, most notably Kandori's (1992) work (see also Milgrom et al. 1990, Ockenfels 2002, and Okuno-Fujiwara and Postlewaite 1995), apply Folk theorem arguments to matching games and show that cooperative norms can emerge among strangers if information about the opponent's reputation is honestly provided and locally transmitted. Applying a complementary indirect evolutionary approach, Güth and

¹⁸ One can leave multiple comments in someone's feedback history, but they count only once, so that no one person can ‘tip the scales’ in either feedback direction, positive or negative.

Ockenfels (2002a/b) discuss how trust in large anonymous and non-anonymous communities may evolve and relate these findings to eBay's attempts to promote trust among strangers. The key insight from these lines of research is that it is not the repeated reciprocal interaction that drives the emergence of trust and trustworthiness, but the amount and quality of information available about the respective transaction partners (see Bolton et al. 2001 for supportive experimental evidence).¹⁹ So, theoretically, a reputation mechanism that collects and distributes information about individuals' trustworthiness can enhance the efficiency of an internet marketplace. But a closer look reveals that there are potentially devastating troubles in translating this basic idea into the practical design of a feedback forum.

The main problem is that the theory does not tell us where the information necessary to promote trust comes from. It is rather typically assumed that the information is exogenously and honestly provided by some more or less unspecified information intermediaries (Kandori 1992; but see Milgrom et al. 1990). On eBay, however, feedback is provided endogenously, by the market participants themselves, which creates a number of incentive problems that may prevent honest feedback provision.²⁰ First, feedback information is a public good. While a user profits from the knowledge about the feedback rating of his transaction partner, he has no incentive to share his own experience with others. As a result, one should expect that feedback is rarely given. Second, the information value of feedback ratings may be low. For instance, a user may want to use secondary user IDs or confederate eBay members to artificially raise the level of his own feedback, or to denunciate a competitor (commonly called 'feedback bombing'). He also may not be willing to give negative feedbacks in order to avoid 'reciprocal' negative feedbacks (as sometimes threatened by the transaction partner). Or, he may offer to sell feedback, trade feedbacks undeservedly, or buy feedbacks. Finally, users can easily change their user IDs and thus disguise their trading histories if they want to. These activities undermine the effectiveness of the feedback forum. Therefore, eBay asks all market participants to report feedback offenses and promises to undertake disciplinary

¹⁹ In Güth and Ockenfels (2002), we also demonstrate how legal institutions such as courts and legal insurance may substitute type detection capabilities not available in completely anonymous communities.

action such as indefinite suspension of a user's account, temporary suspension, or a formal warning.²¹

An increasing number of field studies indicate that despite the theoretical difficulties of decentralized and endogenous feedback provision mechanisms, eBay's feedback system has a significant impact on market participants' behavior. Lucking-Reiley et al. (1999) examine a sample of 461 mint-condition Indian-head pennies for which they were able to obtain estimates of book values. They find that a seller's feedback rating on eBay has a measurable positive effect on the auction price, and that negative feedback ratings have a much greater impact than positive feedback ratings do. Dewan and Hsy (2001), Houser and Wooders (2001), Melnik and Alm (2001), Ockenfels (2002), and especially Resnick and Zeckhauser (2001) come to similar conclusions. Hasker et al. (2001) and Resnick and Zeckhauser (2001) additionally report that the marginal impact of the feedback rating is strongly decreasing.²²

Resnick and Zeckhauser (2001) report that in their sample buyers commented on sellers for 52.1 percent of the time, and sellers commented on buyers 60.6 percent of the time. Thus, a majority of eBay users is willing to contribute to the public information pool. However, there are also indications of reciprocal feedback giving. First, negative feedbacks are given extremely rarely, in less than one percent of all cases, probably indicating a concern of reciprocation. Second, as Resnick and Zeckhauser (2001) report, sellers give a negative feedback with a probability of about 20 percent after having received a negative feedback but only with probability of one percent after having received a positive feedback. The rareness of negative feedbacks may also explain that

²⁰ In the evolutionary framework it is often assumed that people are equipped with individual detection capabilities that work through face to face communication channels (Frank 1987 and 1988), that are, however, not feasible in online trading.

²¹ eBay distinguishes 4 forms of fraudulent feedbacks: defensive and offensive shill feedback (using secondary eBay User IDs or other eBay members to artificially raise the level of your own feedback or to leave negative comments for another user), feedback extortion (demanding any action of a fellow user that he or she is not required to do, at the threat of leaving negative feedback), and feedback solicitation (offering to sell feedback, trade feedback undeservedly, or buy feedback); see <http://pages.ebay.com/help/community/investigates.html>.

²² This might reflect the fact that online identities can easily and cheaply be changed, so that very small feedback ratings may mask an identity change rather than indicate a new user. Theoretical considerations therefore suggest that users with very small feedback ratings must signal their trustworthiness by accepting much lower prices than users with higher positive feedback ratings (see Ockenfels 2002).

prices respond much more sensitive to negative feedbacks than to positive feedbacks – the information value of negative feedbacks appears to be much higher.

So, in principle eBay's reputation mechanism seems to work. But is the design of the feedback forum optimal? Next to eBay's feedback forum there are dozens of alternative designs on the web differing with respect to, e.g., who is allowed to give feedback (only buyers, only sellers, both, only experts, ...), how the statistics are summarized (mean rating, cumulative rating, fragmented rating, ...), how the market participants are identified (user ID, email address, name, ...), contents of feedbacks (numerical, verbal, ...), and the access to user ratings (all, only specific individuals, ...). Furthermore, eBay continually modifies and adjusts the details of its feedback forum. The question how a good reputation mechanism for internet market platforms should look like is apparently not answered yet. But game theory is used to understand alternative models of reputation systems, lab experiments and field data are used to examine empirically the performance of systems at work (much of the research is coordinated at <http://databases.si.umich.edu/reputations/index.html>). The research aims to improve the design of online reputation systems, thereby reducing the risks inherent in interactions with strangers.²³

V. Conclusions

By an innovative design, eBay takes advantage of the technological opportunities made possible by the new online communication channels. In particular, eBay's automatic bidding agents bridge the gap between easy and asynchronous fixed pricing and the flexibility and efficiency of interactive pricing. Tendencies to convert asynchronous into synchronous interaction concentrated at the end of the auction, as reinforced by eBay fixed deadline rule, are diluted by the emergence of suppliers for new artificial bidders such as the one provided by esnipe.com.

²³ The empirical studies seem to suggest that, for instance, eBay should display positive and negative feedbacks separately, because the information value of negative feedbacks is much higher, which is not accurately captured by cumulative ratings. Also, eBay should probably only allow buyers to give feedbacks, thereby diminishing the problem of reciprocal feedbacks between sellers and buyers, and taking into account that only the sellers' ratings (but not the buyers' ratings) seem to have an economically measurable impact on behavior. Theory and the lab study by Bolton et al. (2001) suggests that 'recursive' information about the opponent's reputation performs better than straightforward information about past behavior.

It is likely that the technological development will soon accelerate the use of artificial agents well beyond bidding support and auction search engines as provided by eBay. IBM researchers believe that over the course of the next decade, the internet will be overflowed by billions of software agents (Kephart et al. 2000). Once the artificial agents are fed with preferences, they compare fixed and dynamic price solutions, autonomously interact and negotiate with other robots or human agents, aggregate the exchange-relevant information, and eventually execute the necessary transactions with other trustworthy artificial or human agents identified by online reputation systems. eBay is just the first step to online marketplaces in which artificial agents facilitate perfectly interactive and asynchronous price discovery and information aggregation.

References

- Ariely, Dan, Axel Ockenfels and Alvin E. Roth (2002). "An Experimental Analysis of Late Bidding in Internet Auctions." Working paper, Harvard University.
- Bajari, Patrick and Ali Hortaçsu (2000). "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions." Working paper, Stanford University.
- Bolton, Gary, Elena Katok and Axel Ockenfels (2001). "What's in a Reputation? Indirect Reciprocity in an Image Scoring Game." Working Paper, Harvard University.
- Cramton, Peter (1998): "Ascending Auctions." *European Economic Review*, 42(3-5), 745-756.
- Dewan, Sanjeev, and Vernon Hsu (2001). "Trust in Electronic Markets: Price Discovery in Generalist versus Specialty Online Auctions." Working paper, University of Washington.
- Frank, Robert H. (1987). "If Homo Economicus Could Choose His Own Utility Function, Would He Want One with a Conscience?" *American Economic Review*, 77, 593-604.
- Frank, Robert H. (1988). "Passions Within Reason: The Strategic Role of Emotions." New York: W.W. Norton.
- Güth, Werner, and Hartmut Kliemt (2000). "Evolutionarily Stable Co-operative Commitments." *Theory and Decision*, 49, 197-221.
- Güth, Werner, and Axel Ockenfels (2002a). "The Coevolution of Trust in Anonymous and Non-anonymous Communities." *Jahrbuch der Neuen Politischen Ökonomie*, this issue.
- Güth, Werner, and Axel Ockenfels (2002b). "The Coevolution of Morality and Legal Institutions: An indirect evolutionary approach." Working Paper, Max Planck Institute for Research into Economic Systems, Jena.
- Hasker, Kevin, Raul Gonzales, and Robin Sickles. (2001). "An Analysis of Strategic Behavior and Consumer Surplus on eBay." Working paper, Rice University.
- Houser, Daniel and John Wooders (2001). "Reputation in Auctions: Theory and Evidence from eBay." Working paper, University of Arizona.
- Kandori, M. (1992). "Social Norms and Community Enforcement." *Review of Economic Studies*, 59, 63-80.

- Kephart, Jeffrey O., James E. Hanson, and Amy R. Greenwald (2000). "Dynamic Pricing by Software Agents." *Computer Networks*, forthcoming.
- Klemperer, Paul (2000). "Why Every Economist Should Learn Some Auction Theory." Invited Paper for the World Congress of the Econometric Society 2000.
- Kollock, Peter (1999). "The Production of Trust in Online Markets." In: E.J. Lawler, M. Macy, S. Thyne, and H.A. Walker: *Advances in Group Processes*, Vol. 16, CT: JAI Press.
- Lucking-Reiley, David (2000). "Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce." *Journal of Economic Perspectives*, 14(3), 183-192.
- Mailath, George J., and Peter Zemsky (1991). "Collusion in Second Price Auctions with Heterogeneous Bidders." *Games and Economic Behavior*, 3, 467-486.
- McAfee, R. Preston, and John McMillan (1987). "Auctions and Bidding." *Journal of Economic Literature*, 25, 699-738.
- McAfee, R. Preston, and John McMillan (1999). "Game Theory and Competition." Working paper, mimeo.
- McAfee, R. Preston, and John McMillan (1996). "Analyzing the Airwaves Auction." *Journal of Economic Perspectives*, 10(1), 159-175.
- Melnik, Mikhail I., and James Alm (2001). "Does a Seller's e-commerce Reputation Matter?" Working paper, Georgia State University.
- Milgrom, Paul R., and Robert R. Weber (1982). "A Theory of Auctions and Competitive Bidding." *Econometrica*, 50(5), 1089-1122.
- Milgrom, P., D. North, and B. Weingast (1990). "The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs." *Economics and Politics*, 2, 1-23.
- Moldovanu, Benny, and Manfred Tietzel (1998). "Goethe's Second Price Auction." *Journal of Political Economy*. 106(4), 854-59.
- Ockenfels (2002). "Reputationsmechanismen auf Internet-Marktplattformen." Max Planck Institute for Research into Economic Systems, Jena.
- Ockenfels, Axel and Alvin E. Roth (2001a). "Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction." Working Paper, Harvard University.

- Ockenfels, Axel and Alvin E. Roth (2001b). "The Timing of Bids in Internet Auctions: Market Design, Bidder Behavior, and Artificial Agents." *Artificial Intelligence Magazine*, forthcoming.
- Okuno-Fujiwara, M and A. Postlewaite (1995). "Social Norms and Random Matching Games." *Games and Economic Behavior*, 9, 79-109.
- Peters, Michael, and Sergei Sereniov (2001). "Internet Auctions with Many Traders." Working paper, University of Toronto.
- Resnick, P. and Zeckhauser, R. (2001). "Trust Among Strangers in internet Transactions: Empirical Analysis of eBay's Reputation System." Draft prepared for NBER workshop.
- Resnick, P., R. Zeckhauser, E. Friedman, and K. Kuwabara (2001). "Reputation Systems: Facilitating Trust in internet Interactions. Working Paper, mimeo.
- Roth, Alvin E. (2001). "The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics," Fischer Schultz lecture, *Econometrica*, forthcoming.
- Roth, Alvin E., Murnighan, J.K., and Schoumaker, F. (1988). "The Deadline Effect in Bargaining: Some Experimental Evidence," *American Economic Review*, 78(4), 806-823.
- Roth, Alvin E., and Axel Ockenfels (2001). "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, forthcoming.
- Rothkopf, Michael H., Thomas J. Teisberg, and Edward P. Kahn (1990). "Why are Vickrey Auctions Rare?" *Journal of Political Economy*. 98(1), 94-109
- Shapiro, Carl, and Hal R. Varian (1998). "Information Rules: A Strategic Guide to the Network Economy." Boston: Harvard Business School Press.
- Smith, Michael D., Joseph Bailey, and Erik Brynjolfsson (1999). "Understanding Digital Markets. Review and Assessment." In Erik Brynjolfsson and Brian Kahin, eds. *Understanding The Digital Economy*, MIT Press, 2000, 99-136.
- Ünver, Utku (2001). "Internet Auctions with Artificial Adaptive Agents." Working paper, Koc University, Turkey.
- Vickrey, William (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*. 16, 8-37.