

Does Laboratory Trading Mirror Behavior in Real World Markets?

Fair Bargaining and Competitive Bidding on eBay

Gary E Bolton^{*}
Penn State University, USA

Axel Ockenfels
University of Cologne, Germany

24 February 2011

We conducted a framed field experiment on eBay, and examined to what extent both social and competitive laboratory behavior are robust to institutionally complex real world markets with experienced traders, who selected themselves into these markets. For buyers, the data strongly confirm the dichotomy between equitable bargaining and competitive bidding predicted by social preference equilibrium and suggested by lab evidence. Importantly, reputation building on eBay cannot explain the social behavior. We also observe that the behavioral patterns in the field experiment mirror fully naturally occurring trading patterns in the market. In particular, some sellers fail to use their commitment power as predicted by theories of both selfish and social behavior, with the pattern of deviation reflecting traders' market experience outside the experiment. These patterns further amplify the dichotomy between bilateral and competitive bidding.

* Gary Bolton, Smeal College of Business, Penn State University, University Park, PA 16802 USA, gbolton@psu.edu (<http://www.personal.psu.edu/faculty/g/e/geb3/>). Axel Ockenfels, University of Cologne, Department of Economics, Albertus Magnus Platz, D-50923 Cologne, Germany, ockenfels@uni-koeln.de (<http://ockenfels.uni-koeln.de/ao>). We thank Karin RuëtZ and Dietmar Ilsen for skillful research support. We thank colleagues and seminar participants at Barcelona, Cologne, Copenhagen, Harvard, Innsbruck, Magdeburg, Michigan, Munich, Nottingham and Penn State as well as anonymous referees for helpful comments. Bolton gratefully acknowledges the support of the National Science Foundation. Ockenfels gratefully acknowledges the support of the Deutsche Forschungsgemeinschaft. We used theory, laboratory and field experiments to help eBay (re-)designing certain aspects of their auction rules and feedback forum, partly in order to mitigate various forms of social conflict; this study is partly motivated by but independent from our collaboration with eBay.

I. Introduction

People profess preferences for fairness. Less certain is whether they act on these preferences in a consistent way, or whether these preferences bend to circumstance. The issue has implications for the measure of social welfare and so institutional design. Economists have long observed that competitive market behavior accords well with pure self-interest, and so markets arguably attenuate preferences for fairness. More recently, social preference theory argues that preferences for fairness are compatible with competitive behavior. The primary evidence for this view comes from laboratory market and bilateral bargaining games. These exhibit a pattern of competitive and fair behavior that social preferences can explain but self-interest alone cannot. The catch is that the laboratory is itself a special institution, and arguably different in many ways from real world markets.

This paper reports an experiment performed on eBay to examine whether the laboratory findings concerning fair bilateral bargaining and competitive bidding are robust to a natural trading system. The field games we examine are substantially more elaborate than their laboratory cousins (ex., Roth et al. 1991). Yet as with the laboratory games, the field games have similar equilibrium paths, assuming selfish motives and rational behavior. In the laboratory versions, observed behavior in the bilateral bargaining game is strongly influenced by equity considerations, whereas behavior in the market game follows the selfish equilibrium closely.

There has recently been a good deal of controversy over the extent to which results of laboratory studies in economics extend to more natural settings. Speaking of the differences in lab and field institutions, Levitt and List (2007) posit that, along with monetary considerations, human decisions are influenced by (1) the presence of moral and ethical considerations; (2) the nature and extent of scrutiny of one's actions by others; (3) the context in which the decision is embedded; (4) self-selection of the individuals making the decisions; and (5) the stakes of the game. The authors started a fruitful debate arguing that, "Because the lab systematically differs from most naturally occurring environments on these dimensions, experiments may not always yield results that are readily generalizable" (p. 170).

Ours is a *framed* field experiment (Harrison and List 2004, List 2006a), primarily because traders knew that they are taking part in an experiment (so we cannot derive conclusions for situations in which traders do not know that they are taking part in an experiment as in List 2006b), and because we induced public knowledge about the degree of buyer competition and the maximal economic surplus (so we cannot derive conclusions for situations in which traders do

not know). One practical reason for inducing this knowledge is that this is the standard in related laboratory environments, and so this allows us to directly compare results. More importantly, however, models of inequality aversion, which are widely applied to related laboratory environments, predict *selfish* behavior if traders expect competition or do not know the comparison standard. Laboratory evidence confirms this (e.g., Cooper and Kagel, forthcoming). This implies that in order to investigate – based on clear theoretical predictions – whether social preferences as observed in the laboratory can also be identified in our naturally occurring context, we need control what people know about competition and surplus; otherwise we could not be sure if observed selfish behavior is just due to a lack of information or due to different preferences (see section II for more details, and section IV for a discussion of related literature). Another advantage our design affords is that it captures the natural circumstances of eBay’s marketplace along the five above dimensions mentioned above: The transactions take place on eBay’s platform according to all of eBay’s naturally evolved rules. Subjects are experienced traders who self selected themselves into the eBay environment for their own, independent purposes (including into their market roles as buyers or sellers). The stakes reflect those of the average transaction on eBay. As is common in many markets, trader behavior on eBay is highly scrutinized, both by other traders and by the market’s management. EBay traders post “feedback” on their experience with other traders, information made available to all traders when deciding on future trading partners. Moral and ethical considerations are part of these postings. The actions in our experiment were official eBay transactions and became part of these feedback histories; this allows us to link some of the experimental data to naturally occurring field data. Behind the scenes, eBay management monitors the site for malfeasance, and this too plays a role in our experiment.

Our test addresses another factor that is a critical differentiator of lab and field ecologies, this being complexity. The laboratory is designed for highly controlled empirical research. Trading systems are typically stripped down to match the parsimony of the theory being tested, with the goal of minimizing complicating factors and associated (boundedly) rational responses. In contrast, real world trading systems endure because they facilitate trade for profit. They are agglomerations of formal and informal rules, norms and heuristics that, together with trader experiences, histories and networks, evolve over time. As such, the natural trading system is far more complex than what has been implemented in a laboratory, and thus also creates scope to observe more complex behavior including non-rational trading patterns.

To give an example of system complexity in the present context, eBay's trading rules have the same first mover-second mover format as do ultimatum bargaining games (ex., Güth et al. 1982) and simple market games (ex., Roth et al. 1991). In both eBay and the lab games, first movers have a significant advantage in that they have considerable commitment power, able to specify key terms of the deal. Yet the eBay platform offers the sellers who set up the trade more options than are typically studied in a laboratory (ex., auction, posted price), with subsidiary rules serving practical purposes that must be abided by. eBay is also a far noisier trading environment than the lab. For instance, computer processing time and connection speed may prevent very late bids from being accepted, traders might employ shill-bidding or other fraudulent strategies, or they may experience computer, server, health or other problems which can affect their trading. Laboratory studies typically avoid such complexities.

The present experiment exploits the natural game forms implicit in eBay's trading system in a way that – despite of the institutional complexities – enables us to classify the behavior we observe as either in-selfish-equilibrium or in-social-equilibrium or as out-of-equilibrium. We only need establish the necessary information structure and opportunity costs by manipulating the object traded (which we do without bending the natural rules of trade). This also allows us to relate behavior in our field experiment to robust laboratory findings, and when they agree we can be rather confident that behavior in our field environment is part of a more general trading inclination. Finally, eBay permits us to check for relationships between behavior inside the experiment and fully naturally occurring behavior outside the experiment. As in many markets, after the price has been determined, and independent of the degree of competition in the price discovery phase, all eBay transactions boil down to bilateral ones. To facilitate this phase, eBay makes available transaction histories of traders, including reputation, via eBay's feedback forum. Indeed, transactions in our experiment became part of our traders' official eBay trading histories. We exploit these links between bilateral exchange in and outside the experiment, to test several hypotheses, partly suggested by others concerning fully natural (eBay) trading.

Our findings complement observations from field experiments conducted in very different domains and with different degrees of control (discussed in section IV). For example, parts of our conclusions confirm the finding in List (2004), who observes competitive behavior in markets with face-to-face interaction both in laboratory and framed field experiments, and the finding of List (2006) who reports data from artefactual and framed experiments that are in line

with social preferences.¹ One of the advantages of an experiment on eBay is relative ease of replication, admitting the kind of ‘losing control in a controlled way’ research program that has proven productive in the lab.² Easily accessible field test beds control for the differences in complexity across field environments, and permit the method of replicating existing results and testing new hypotheses in a series of experiments. Online markets, being widely accessible, are one such set of test beds (also see Chen et al. 2010).³ We discuss this issue along with related research in section IV, after discussing our experiment (section II) and results (section III). Section V concludes.

II. Design of the experiment and hypotheses

The entire experiment was conducted on eBay’s German market platform, which is mostly equivalent to the U.S. eBay market (the relevant exceptions will be discussed). The experiment offered no explanation of the platform’s trading rules: Subjects were simply referred to eBay’s own explanation and presentation of the trading formats. Our traders neither showed up personally nor signed any documents or receipts before, during or after the experiment. Nor did they have to report the data to us, this being available on eBay for anyone to download.⁴ All communication between traders and experimenter was done via the Internet. Communication was automated, with the exception of answers to private questions to the experimenter (see Appendix for all standardized communication). Prior to participating, traders had to confirm on a webpage that they were available at the time of the experiment, and they had to correctly answer some questions to demonstrate basic understanding of the set-up (ex., the number of potential

¹ An artefactual experiment is a laboratory experiment with a “non-standard” subject pool. Where social preferences have less of an effect in List (2006) is when he moves from framed to natural field experiments, wherein sellers do not know that they are taking part in an experiment.

² For instance, it took many experiments, performed by many different investigators, to vet the robustness of social preferences in bilateral bargaining under fully laboratory conditions; for reviews see, ex., Roth (1995), Bolton (1999) and Camerer (2003). For a sample of, and references to, additional work see Andreoni (2003).

³ eBay’s markets are popular with a large and diverse set of traders, a total 222 million registered users in 33 countries, all in cyberspace. For a history of eBay, which touches the early evolution of trading rules and norms, see Cohen (2002). For more information on eBay, including statistics on usage, size of company, international breadth etc., see <http://investor.ebay.com/index.cfm>. For a survey on economic eBay research see Ockenfels et al. (2006). See Harrison and List’s (2004) and List’s (2006a) taxonomy of field experiments, classifying them depending on their distance between the extreme of laboratory and natural environments. Also see John List’s bibliography of field experiments at <http://www.fieldexperiments.com>. Kagel (1995) provides an introduction into laboratory auction studies.

⁴ However, after the experiment was finished, we sent out a question to selected eBay buyers in order to better understand the motives behind their behavior; see section III.3.

buyers; see Appendix). All game payoffs were wired in order to keep the money flows under our control. In addition, independent of trading success, sellers were paid all eBay transaction fees.

In total, we recruited 400 experienced eBay traders from student populations at four German universities: Bonn, Cologne, Jena and Magdeburg. The average feedback score in our subject population was 89, and the median was 44.⁵ This implies that the average (median) subject successfully completed at least 89 (44) eBay transactions as a winning bidder or as a seller, and thus has gained considerable experience with eBay's market platform. To participate, a trader had to have a valid eBay account and a feedback score of at least one. In addition, we made sure that the sellers in our experiment had a feedback score of at least 10, and received at least one of the feedback points in the role of a seller, guaranteeing some minimum role-specific trading experience.⁶

II.1 Gaining control on eBay

While eBay's trading institutions and options (the rules of the game) are substantially more detailed and complex than what is typically implemented in the laboratory, they lend themselves to relatively crisp analysis – provided that we gain control over some critical theoretical parameters. To the extent possible, we crafted the controls to mirror the pattern of natural activity. For example, the geographic dispersion of our subjects permitted us to select cohorts of traders that are anonymous to one another and unlikely to have a common history or future. In this way, our design minimizes repeated game and social interaction effects, something that is important when interpreting motivations behind behavior. At the same time, this geographic dispersion – and associated one shot trading – mirrors the norm on eBay (e.g., Resnick and Zeckhauser 2002).

Another important control is induced valuations (Smith 1982), permitting a clear interpretation of agents' pecuniary incentives. We did this by creating tradable "certificates" that had specific values to specific traders. These control the opportunity costs associated with a trade. Induced values were chosen such that the equilibrium prices in the experiment roughly match

⁵ The feedback score is the sum of all positive feedbacks received minus the negatives received. Thus, a feedback score of n implies that the trader must have successfully completed at least n eBay transactions. Because transactions may not end successfully (a seller may find no buyer and a bidder may not win), or successful transactions may end without or with negative feedback, the feedback score underestimates the traders' experience. It is, however, often used as a proxy for experience; see, e.g., Roth and Ockenfels (2002).

⁶ By the rules of eBay Germany, a seller can only choose the buy-it-now format if he has at least 10 positive feedback points, so that this restriction also made sure that all sellers had the same offer formats available (<http://pages.ebay.de/help/sell/fixprice-faq.html>).

average prices on eBay. The equilibrium price in both treatments will be about 20 €(see section II.4), or about \$24 at the time of the experiment. According to eBay, in 2005 (when we conducted our experiment), about 72 million active eBay users generated a gross merchandise volume of \$44.30b in 1.88b listings, yielding an average price of \$23.56.⁷

We made buyer and seller valuations and the degree of competition (that is, the number of buyers with positive valuations for a given certificate) public knowledge. This choice was not driven by a goal to maximize the realism of our setting, but rather to gain sufficient control for a clean derivation of hypotheses. It has been shown, both in theory and in the laboratory, that to the extent traders expect competition or do not know the economic surplus to be distributed, social preferences are consistent with selfish behavior (see, e.g., Fehr and Schmidt 1999, Bolton and Ockenfels 2000, Card et al. 2011, and Ockenfels et al. 2011 and the references therein). So, in order to rule out the possibility that selfish behavior is the result of limited information rather than selfish preferences, we provided the relevant information. Also, these features align our setting with a large class of well-known laboratory bargaining and price competition games with publicly known valuations (section II.4), and thus allows us to link our results back to lab phenomena.

On eBay, the distributions of valuations and number of bidders vary widely, given the wide array of goods and trader sub-populations on the platform. Regarding the valuations, there often appears to be considerable uncertainty, although we do not know of any study that measures how the degree of uncertainty varies across bidders and products (we get back to this in our concluding section). Regarding competition, the number of bidders in eBay auctions follows a power law distribution, which implies that if there are bidders at all, one-bidder interaction is the most likely scenario (e.g., Yang et al. 2003). For eBay Germany, where our experiment took place, Namazi and Schadschneider (2006) found that about 8 percent of all eBay auctions have exactly one bidder. So, while usually an eBay-seller cannot be sure ex-ante about the number of active bidders, our experimental conditions capture two simple scenarios, one or multiple bidders, which are both typical for eBay transactions and so must both be taken into consideration by (rational) sellers when devising selling strategies.

⁷ See http://investor.ebay.com/downloads/fund_Metrics.pdf.

II.2 Bilateral trading environment

The experiment involves a bilateral trading environment and a competitive trading environment. In each encounter of the bilateral trading environment (BT), one seller and one buyer met to transact a single certificate, with a value of zero to the seller and 20 € to the buyer. Each participant was invited to engage in exactly one trading encounter. All this information was part of the general instructions sent out to all traders participating in the BT treatment.

eBay rules require that the seller first chooses a trading format along with a price offer. The trading format can be an auction, buy-it-now (BIN), which is essentially a posted price, or a combined auction plus BIN.⁸ The *price offer* is a start price if the seller chooses the auction format, a fixed price if the BIN format, or both if the hybrid format. If the seller chooses the auction format, a buyer can start bidding in an open, dynamic second-price auction.⁹ In case of a buy-it-now offer, a buyer can accept the fixed-price specified by the seller. Sellers are also free to choose the duration of the offer, from 3 to maximally 10 days. By eBay's rules, however, all offers must end at a publicly known and predetermined time.

Altogether, there were 50 buyer-seller pairs involving 100 participants. The start time was restricted in order to insure that offers were made independently (Monday, 10/10/2005, 3pm). No offer could be submitted after this time; in particular, sellers were not allowed to submit a second offer in case the first one did not end in a transaction. Also, sellers were told to post their offers under a specific eBay category, and to use a specified item description, in order to suppress both tacit and explicit communication between traders (known to have an effect in lab studies).¹⁰ Finally, sellers were not allowed to use any marketing options, such as bold font, highlight etc., or to ask for shipping or handling costs.

⁸ In the latter case, a buyer who chooses the BIN-option ends the auction immediately. If, however, the first buyer submits an auction bid, the BIN-option disappears and the auction proceeds normally. About one third of all eBay transactions were completed via BIN-offers in 2005, the rest via auctions. At the time of our experiment, eBay's "best-offer" format already available in the U.S. (which allows 'hidden' negotiations) was not available in Germany.

⁹ More specifically, eBay asks the bidders to submit maximum bids (called "proxy bids") and explains that "eBay will bid incrementally on your behalf up to your maximum bid, which is kept secret from other eBay users." That is, once a bidder submits his (proxy) bid, eBay displays the currently winning bid as the minimum increment above the previous high proxy bid or, in case there is only one bidder, equal to the auction start price. At the end of the auction, the bidder who submitted the highest bid wins the auctioned item and pays a price equal to the second-highest bid plus the minimum increment. See eBay.com or eBay.de for more details and Ockenfels and Roth (2006) for a description of eBay's auction rules from an auction theoretic perspective.

¹⁰ The eBay category was "Sammeln & Seltenes > Technik & Geräte > Wissenschaft & Medizin > Sonstige." ("Collectibles > Technology & Tools > Science & Medicine > Miscellaneous"). Translated from German, the item description was: "This offer is placed as a part of a research project. The certificate is only valuable to you when you are registered as a participant of the project and are informed about the article number of this offer. Please refrain from asking the seller questions."

Once the seller informed the experimenter about the offer identity number, we checked if everything was in line with the rules and only then did we inform a randomly selected buyer from a different university population about the offer identity number (so the buyer could not communicate with the seller prior to the offer posting). The buyer could then accept the offer by bidding (at least) the start price in the auction or by accepting the BIN-offer, respectively. If the buyer chooses to let the offer time expire without accepting the price offer, both buyer and seller end up with zero payoffs.

While BT roughly resembles laboratory bilateral bargaining games such as the ultimatum game, it differs in important respects. First of all, the eBay seller has more options than typically allotted a first mover in a laboratory bargaining game. Also, unlike in laboratory bargaining games, we did not explicitly point buyers to the option of ‘rejecting’ an offer – traders were simply told they were free to take any action as long as it is in line with eBay’s trading rules. Not being given an explicit option of rejecting an offer may affect the likelihood of choosing it.¹¹ Other differences include that buyers were free to choose the timing of their bids; other traders (without induced valuations for the ‘certificate’, including shill bidders) can submit bids; traders could in principle contact each other (though, we do not know of any such instance); they could see the opponents’ reputation, and, if the offer is accepted, leave feedback on each other. All of these features can have and some do have implications for the theoretical analysis and the actual trading behavior, as demonstrated in sections II.4 and III.

II.3 Competitive trading environment

In the competitive trading environment (CT) one seller met 9 buyers to transact one trading certificate. As in BT, the value of the certificate to the seller was zero, while the value to each buyer was 20 €. Altogether, there were 30 markets (300 subjects). Each participant was invited to engage in exactly one trading encounter. No subject participated in the bilateral trading encounters. In each market, bidders were recruited from all four universities. All offers started on Monday, 11/14/2005 at 3pm. Everything else was as in the bilateral trading environment.

The game resembles various laboratory price competition games, but it too differs from the laboratory settings in important ways. First of all, typical laboratory Bertrand games (e.g., Dufwenberg and Gneezy 2000, Roth et al. 1991) and laboratory second-price auctions (see Kagel

¹¹ Also, one loses some control over *why* an offer is not accepted; we will come back to this when we analyze the data.

1995 for a survey) involve sealed bidding, while eBay involves open bidding, allowing the possibility of coordinated bidding (section II.4). Also, in laboratory price competition among buyers, sellers typically do not assume any active role. However, on eBay, sellers choose formats and price-offers, which turn out to be theoretically and behaviorally relevant (section II.4). Another difference is that the “market game” by Roth et al. (1991) allows the seller to reject the best offer, while the highest bid on eBay automatically establishes a binding contract.

II.4 Theory and hypotheses

Our two central hypotheses differ regarding the extent of advantage sellers gain from their commitment power across BT and CT environments. Our null hypothesis is derived under the assumption that traders in the market act to maximize their pecuniary gains. The alternative hypothesis is derived under the assumption that traders maximize social preferences, as often supported in laboratory environments. Proofs of the hypotheses are based on perfect Bayesian equilibrium, which reduces to perfect equilibrium when we assume selfish preferences since then games are complete information.¹²

There are many social preference models. Here, we restrict our attention to the models by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000), who are outcome-oriented and thus tend to be simpler than models of reciprocity that are based on psychological game theory, such as Rabin (1993) and Dufwenberg and Kirchsteiger (2004).¹³ There are also many ways to model eBay. Again, we restrict ourselves to two prominent approaches. One is to abstract away from all dynamics and model eBay as a second-price sealed bid auction. In fact, Bajari and Hortaçsu (2003) show that under certain assumptions, equilibrium in a dynamic eBay auction model is formally equivalent to the equilibrium in a second-price sealed-bid auction model. Ockenfels and Roth (2006), on the other hand, develop a game theoretic model of eBay’s dynamics, in which early bids are modeled as taking place at times t on the half open interval $[0,1)$, while late bids happen at time $t = 1$. Thus there is always time to follow an early bid with another bid, but late

¹² We caution that because the theoretical predictions are quite demanding and because we do not give our traders learning opportunities within our framed experiment (which includes atypical aspects such as fictitious commodities and public knowledge of values and degree of competition), one should not expect that equilibrium predictions are fully supported by the data. In fact, as we will argue, the evidence suggests that out-of-experiment experience and mistakes matter.

¹³ This is not an unusual approach in the experimental economics literature, and seems especially justified in our context, because most social preference models predict the same or similar behavioral patterns in standard ultimatum bargaining environments – although the underlying motives and so behavioral predictions may differ significantly in non-standard games; see Cooper and Kagel (forthcoming) for the usage, differences and limitations of social preference models in experimental economics.

bids happen simultaneously, when it is too late to submit a subsequent bid in response. Early bids are transmitted with probability 1, while late bids are successfully transmitted with some probability p that may be smaller than one. For our private value environment, this model implies a continuum of equilibria.

Our hypotheses describe the most straightforward equilibrium patterns. In addition, we outline the implications of other equilibria, resulting from the dynamics of eBay bidding and other institutional complexities. While our two hypotheses serve as a first benchmark, we discuss our data with respect to the full set of equilibria.

Null hypothesis (based on selfish preferences): Independent of the treatment, sellers are indifferent between the BIN-format, the auction format and the hybrid format. Regardless of format, seller choice of start- or BIN-price (both in hybrid offer-format) is equal to the buyer valuation of 20 €.¹⁴ At least one buyer accepts the proposed price.

Proof: eBay requires sellers to stipulate an initial price offer either in the form of a start price for the auction or a BIN price, or both. The price offer is known to the potential buyers, because we chose an eBay category that does not allow a hidden auction reserve price. Moreover, the price-offer is a credible commitment, because when trade fails to be executed, sellers were not allowed to start a second offering. This was public information available to all traders. That the null hypothesis is a perfect equilibrium regardless of the degree of competition then follows from straightforward backward induction.

There is also a ‘selfish equilibrium’ in CT with a low auction start price and at least two bidders successfully bidding their full values (and so again, a final price of 20 €). But if we consider eBay’s dynamics, a low start price would allow bidders to coordinate on low-price equilibria through early or late bidding. Ockenfels and Roth (2006) demonstrate, in an eBay private-value second-price auction model that matches the bidder situation in our CT condition in every relevant detail, that there are ‘late bidding equilibria’, in which all buyers bid values in the

¹⁴ Strictly speaking, of course, an offer of 20 € makes the buyer indifferent and thus may be rejected, in which case the equilibrium price in our experiment is 19.99 €. Also, if buyers require more than one cent to break the indifference, an alternative null hypothesis would be that the seller payoff is smaller than 19.99 € – but the same in bilateral and competitive trading. (As we will see, both hypotheses will be rejected.) For simplicity of exposition, however, we drop equilibria with $20 - \epsilon$ payoffs for the rest of this theory section. We will come back to such equilibria in the data section.

closing seconds of the auction.¹⁵ Because late bids run the risk of not being successfully transmitted to eBay, it can happen with positive probability that no bids come through, which implies zero revenues, or that only one bid comes through, which because of eBay's second-price rule implies that the auction revenue is equal to the auction start price.¹⁶ However, all of these late bidding equilibria imply an outcome that gives positive chance to the seller receiving less than the entire pie, while in our CT treatment, the seller can *unilaterally force* the null hypothesis equilibrium, thereby capturing the entire pie with probability 1 by setting the start- or BIN-price to 20 €. So, before collecting the data, we decided to stick to the null hypothesis as our central benchmark. (Yet, as we will see later, the data speak more in favor of the prediction that sellers rely on buyer competition and that this reliance is justified.)

We next turn to examine behavior under the assumption that traders value relative as well as pecuniary payoffs (they have social preferences) as modelled by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000). Both models posit heterogeneity with respect to how players trade-off pecuniary and relative payoffs, making these games of incomplete information. The specifics of the alternative hypothesis are derived from the Bolton and Ockenfels model, although similar results obtain from the Fehr and Schmidt model (discrepancies are pointed out in the footnotes). In the Bolton and Ockenfels model, individual utility is defined as $u(x, \sigma)$, u increasing in own payoff of x and decreasing as relative payoff σ moves away from $1/n$, where $n = 2$ in BT and 10 in CT; if total player payoffs, c , are non-zero, $\sigma = x/c$, otherwise $\sigma = 1/n$. As before, to economize on space, we refer the reader to proofs in the literature where possible.

Alternative hypothesis (based on social preferences): (1) In BT, the sellers' average price offer is smaller than predicted by the null hypothesis but is always at least 10 €. All buyers accept price offers of 10 € but some do not accept price offers of more than 10 €. (2) In CT, it is convenient to

¹⁵ See their Theorem on p. 303, in which they assume that all values are drawn from a degenerate distribution with all its mass at H (= 20 € in our experiment). Their proof also implies that there is no late bidding equilibrium if the start price is equal to H , in which case there is a continuum of equilibria in all of which bidders push the price up to H at some time 'early' in the auction.

¹⁶ Any 'early' bid immediately triggers a price war in which all bidders bid their values early. Because bidding value early is also an equilibrium strategy, this threat is credible. Similar late bidding equilibria could in principle also exist in the bilateral trading environment. At the time of the experiment, eBay allowed sellers in certain categories to lower one's start price as long as no bid has been submitted. This creates incentives for buyers, when no BIN-price is offered, to delay the bid until very late, when there is a positive probability of conflict, in the hope that the seller will revise his minimum demand (e.g., Ma and Manove 1993; see also Roth et al. 1988). However, in the categories chosen for the experiment, this option was not available.

distinguish selfish and fair-minded sellers: (i) A selfish seller uses the auction format with a start price of 18 €. The final price of the auction is bid to 20 €. (ii) Fair-minded sellers use the BIN-format; a price-offer at or below 18 € is accepted.

Proof of (1): Given the Bolton and Ockenfels utility function, an equal split of the bargaining pie is always preferred to a split where the other bargainer receives more than half of the (same) pie. For this reason, the seller should never choose a start price of less than half the pie, 10 €. The rest of formal proof then straightforwardly follows the proof of Statement 2 in Bolton and Ockenfels (2000). Intuitively, inequitable offers in the BT treatment will not be accepted by buyers who prefer the equitable (zero) allocation that results if the offer is not accepted, and for this reason, the optimal price offer is less than 20 €

Proof of (2i): A price of 18 € would yield a payoff of 2 € to the buyer, which is equitable and thus acceptable to all bidders.¹⁷ The rest of the formal proof then follows the proof of Statement 9b in Bolton and Ockenfels (2000). Intuitively, at a price offer of 18 € all buyers are better off bidding and possibly winning, than not bidding and either having the no-trade outcome or allowing another buyer to make the trade. Competition aligns absolute and relative payoffs so that buyers behave as if they are selfish.¹⁸

Proof of (2ii): In this case, the exact seller price offer depends on the trade-off between equity and selfish concerns of the seller, and may differ across models.¹⁹ If the seller's ideal selling price is 18 € or less (not below 2 € in the Bolton and Ockenfels model), then the optimal strategy is to set a BIN-price at the ideal level. This offer is acceptable to all buyers (and using

¹⁷ An offer of 2 € is equitable and thus acceptable in the Bolton and Ockenfels model because it is the share the buyer would get if the efficiency gain from trade is equally distributed among all market participants (see Bolton and Ockenfels 1998 for an illustration in a laboratory 3-person bargaining game). If one applies other fairness norms or uses a different definition of the reference group, other payoffs may be interpreted as fair (see, e.g., Bolton and Ockenfels 2005 and 2006 for discussions of different reference points for equity), which however would not affect the qualitative predictions.

¹⁸ Suppose a trader's utility is increasing in one's pecuniary payoff but also depends on a relative payoff component. The relative payoff is defined as one's share of the total payoff of all involved n traders, and is equal to $1/n$ in case all traders get zero payoffs. Now suppose that utility is increasing in relative payoffs as long as the relative payoff is smaller than $1/n$. Then, losing in the competitive trading environment means that others will share the surplus, implying the worst possible outcome: zero pecuniary and zero relative payoffs.

¹⁹ Bolton (1991), for instance, postulates that traders do not care about equity when they are in a favoured situation. As a consequence, CT-sellers in Bolton's model, who are in a favoured position, behave just like sellers in the standard model and maximize their revenues as assumed in the discussion above. The models by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) are consistent with price-offers well below 20 € but differ with respect to the formulation of equity reference point; in Fehr and Schmidt traders evaluate their payoffs in comparison to each other trader, whereas in Bolton and Ockenfels traders only care about their relative status with respect to the whole reference group. Social utility models based on psychological game theory also differ with respect to equity reference points.

the auction format risks a higher final price). If the seller's ideal selling price is strictly between 18 € and 20 €, then the seller must weigh the risk of having such an offer not accepted (see discussion below) against having the competitive outcome or an outcome of 18 € or less (the latter two can be had with probability 1), and the decision reached can differ depending on the precise specification of the utility function. However, it is always the case that the seller uses the BIN-option to make the price offer, since using the auction format risks a higher price settlement than the fair-minded seller would like.

As stated in the alternative hypothesis, social preference models imply that sellers may want to choose an auction with a start price below valuations even though valuations are common knowledge. To understand this more clearly, observe that the null hypothesis equilibrium offer of 20 € is not a reasonable perfect Bayesian equilibrium offer with Bolton and Ockenfels social preferences because accepting is weakly dominated by not bidding at all: If nobody accepts, the outcome is equitable (all market participants receive zero payoffs), and if one or more buyers accept, all buyers still receive zero pecuniary but now inequitable payoffs with respect to the seller's payoff. (Price offers strictly between 20 and 18 € also run a risk of attracting no bids, depending on the exact preferences of the bidders.)

As under selfish preferences, there are coordinated bidding equilibria in the open eBay format.²⁰ Importantly, demanding the whole cake is (unlike in the scenario with selfish preferences) no resort for the seller in models with social preferences. Fair-minded buyers have good reasons to reject very greedy offers, but even fair-minded bidders dissipate all buyer rents in an auction once they get locked into competition.

We finally note that the analyses above abstract away from some influences that might potentially also affect behavior. Most notably, one might think that eBay's feedback system opens the door for other equilibria that may confound social preference predictions, and that the transaction costs in our field settings are larger than in corresponding laboratory settings. We discuss these issues when we present and interpret the evidence.

²⁰ For a coordinated early-bidding equilibrium, observe that an early bid of 20 € by another bidder implies zero pecuniary payoffs and zero relative payoffs regardless of one's own bid (as long as it does not lead to negative payoffs), so that the incentive to compete against a 20 € bid is weak with social as well as selfish preferences.

III. Trading behavior

We first investigate seller behavior in both treatments, and then separately analyze buyer behavior in competitive and bilateral trading environments. Section III.4 discusses the potential effect of reputation building via eBay’s feedback forum on social concerns, and vice versa.

As with lab experiments, we had some ‘no-shows’: In the bilateral trading environment, 4 (out of 50) sellers did not submit an offer to eBay. In the competitive trading environment, 1 (out of 30) sellers did not submit an offer. Analyses are based on the remaining observations. Analyses of payoffs and revenues are net of eBay’s fees, which were paid by the experimenter.

III.1 Seller behavior: format choice and price offer

Table 1. Tabulation of offer channel selected by sellers

	auction	BIN (buy-it-now)	auction with BIN-option	Total
BT	12	16	18	46
CT	17	1	11	29

Table 1 tabulates the choices of the offer channels and shows that the sellers’ choices depend on the treatment (chi-square, $p = 0.002$). More specifically, when we neglect the hybrid format, there is a slight but not significant tendency to use a fixed price format in the bilateral trading environment BT (binomial, $p = 0.572$), while there is a strong and highly significant tendency to make use of competitive bidding in the competitive environment CT ($p < 0.001$). This pattern of format choice is suggested by the alternative hypothesis in the sense that buyer competition in CT-auctions can avoid problems associated with buyer resistance to accept high BIN price offers (see section II.4). It is, however, also consistent with the null hypothesis, because every choice is in line with the selfish equilibrium outcome.

The pattern of format choices in our experiment may not only be affected by isolated equilibrium considerations. It has been reported that outside experience influences a subject’s behavior within an experiment (Burns 1985, Dyer et al. 1989, Cooper et al. 1999, Bolton et al. 2011). In line with this literature, we observe that a seller’s trading experience and preferences that evolved in the field, outside of the experiment, also influences format choice. After the experiment, on 12/12/2005, we downloaded transaction data of our subjects available on eBay. This data on trading histories include all transaction data if the transaction took place not earlier

than 9/9/2005 and when the transaction partner left feedback (on eBay Germany, other data were not available). The history data show that those BT-sellers who chose the auction (BIN) format have a significantly higher probability of choosing an auction (BIN) in non-experimental eBay encounters (Mann-Whitney U , two-tailed, $p = 0.028$ for the auction and 0.000 for the BIN format; the corresponding effect for auction plus BIN-option is with $p = 0.286$ not significant). Analogous effects hold for CT-sellers (p -values are 0.003, 0.014, and 0.009). This demonstrates that the trading strategies in our experimental setting mirror the sellers' individual trading patterns on eBay in naturally occurring field encounters.

Figure 1. Tabulation of sellers' initial price offers

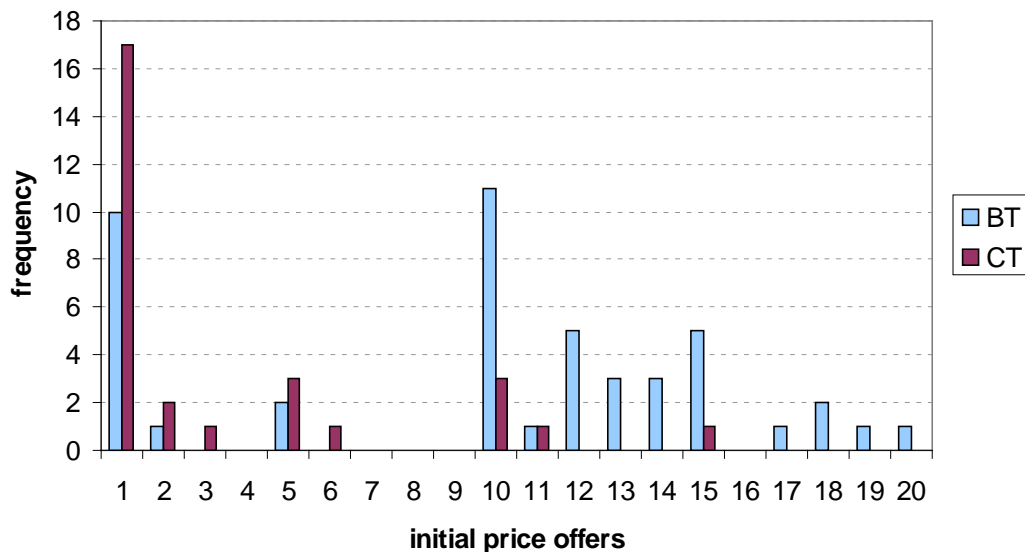


Figure 1 shows the initial price offers posted by sellers in bilateral and competitive trading.²¹ For the instances in which the seller chooses the hybrid format involving an auction start price and a fixed price, we define the price offer as the minimum of the two, which by eBay rules, must be the auction start price. Overall, only a single price offer is consistent with our null hypothesis (19.99 or 20 €). The modal offer is 10 € the equal split, in BT and 1 € the default start price on eBay, in CT.²² If we include selfish equilibria where sellers rely on buyer competition driving the price up to 20 €, all offers in CT are in equilibrium, yet all but one offer in BT are not. The alternative hypothesis does only somewhat better. While ‘social equilibrium’ is in principle consistent with offers strictly below 20 € in both CT and BT, price offers below the

²¹ As will become clear later, these price offers are not necessarily identical with the price offers eventually accepted or not accepted (e.g., Figure 5), so we call them “initial.”

²² Price offers in BT are higher than in CT (Mann-Whitney U , two tailed, $p < 0.001$).

equal split in BT (28 percent) are out-of-equilibrium regardless of whether preferences are social or selfish.²³

Offers of below 50 percent are sometimes observed in laboratory games, where the specific pattern of behavior suggests they are mostly mistakes rather than attempts to be generous. For example, in the first two rounds of Roth et al.'s (1991) large ultimatum game experiment, 9.3 percent of the observed offers were below 50% of the pie, but by final rounds 9 and 10 none of the offers were below 50%. This suggests that the initial offers below 50% in our experiment were due to inexperience rather than preference. Prior to trading, we quizzed sellers to ensure they understood there would be only one potential buyer (see questionnaire in Appendix). So the rate of these kinds of mistakes is likely due to sellers who did not know from previous experience on eBay how to handle the lack of competition in BT (in section III.2 we will see that a low start price in CT is a quite sensible strategy).

This view is weakly supported by a closer look at sellers' eBay histories from outside the experiment. Those BT-sellers who placed a price offer smaller than the equal split were generally less experienced with the use of auctions and start prices. For instance, BT-sellers who chose auction start prices below the equal split generally placed smaller start prices on eBay than the other sellers (an average start price of 1.92 € versus 7.06 €, Mann-Whitney U , two-tailed, $p = .021$), suggesting that those sellers are less familiar with utilizing start prices. Also, BT-sellers with offer prices below the equal split are somewhat less inclined to choose the BIN format, a more 'natural' choice in bilateral bargaining (an average proportion of BIN-offers relative to all eBay offers of 4.5 percent versus 25.8 percent; Mann-Whitney U , two-tailed, $p = .111$). Finally, the average feedback score of those BT-sellers with price offers below the equal split is 95 versus 224 for the other sellers, and the average number of auctions conducted is 4.5 versus 10.2, though these differences are not statistically significant.

One interesting individual example along these lines is a seller who dealt with the lack of buyer competition by shill bidding. Shill bidding is an illegal attempt by sellers to raise the price by bidding oneself or through a confederate just below the buyer's maximum bid. The seller chose a start price of 1 € and later submitted shill bids to increase the final price to 5.50 €. Why do we suspect shill bidding? For one, the first part of the name of the bidder, who submitted the shill bids and who was not registered at our experiment, is identical with the last name of the

²³ We treat a 9.99 € offer here and elsewhere in the paper as equitable. Such low offers also cannot be rationalized by maximin or efficiency preferences as, e.g., proposed by Charness and Rabin (2002).

seller. Second, we found four non-experimental auctions on eBay run by our seller, where the same bidder submitted non-winning bids. So it seems likely that the seller adopted shill bidding from his eBay experience outside the experiment.²⁴


Summing up, we find that with only one exception, sellers do not try to skim all rents directly through their price offer, as predicted by our null hypothesis. As we will see in the next subsection, some of this can be rationalized by sellers' reliance on buyer competition in CT or by sellers' anticipation of buyers' resistance to accept high price offers in BT. However, even when taking this into account, some BT-sellers still fail to make use of their commitment power: they choose price offers too low to be consistent with any selfish or social equilibrium. The data provides some evidence that this behavior is related to a lack of previous experience on eBay, gained outside our experiment, how to handle the lack of competition in BT.

III.2 Buyer behavior in competitive trading

Figure 3 shows the final prices (seller profits) realized in the 29 competitive markets. Observe that whenever bidders competed in auctions (the 20 dark bars in Figure 3), the final price was above 19 €. In 14 of these auctions the price was as predicted by equilibrium theory (19.99 € or 20 €). In 5 auctions, prices were in a range between 19.50 € and 19.98 €. This price range describes stable, competitive outcomes of the bidding process, because the minimum bidding increment is 0.50 € implying that no bidder can profitably alter the outcome once the price reached this range. Only in one CT-auction was the final price below 19.50 €, namely 19.09 €. The close match between competitive market data and theory is also observed in the laboratory (ex., the competitive markets reported in Roth et al. 1991).

²⁴ We did not include shill bidding in our theory, because it is neither allowed nor advantageous for sellers to employ shill bidding strategies in our environment. For instance, as explained in section II.4, CT-sellers can always guarantee a price of 20 € without any shill bidding and regardless of whether the bidders are selfish or fair-minded. Moreover, CT-buyers can protect themselves against shill bidding by matching the seller's initial price offer in the sense that any shill bid would then have to outbid the buyer. Because all rules of eBay applied in our experiment, we excluded neither the seller nor the bids of the shill bidder when computing the final payoffs. If eBay had closed the auction because of shill bidding (as it happened once in our competitive trading environment), the game payoffs to all participants would have been zero, just like on eBay.

Figure 3. Realized prices (seller profit) in the 29 competitive trading encounters



The success of auctions with respect to revenues and the large agreement with prices predicted by equilibrium theory – for both selfish and social preferences – is remarkable, given that sellers gave more control away over the price discovering process than theory says is necessary. In particular, observe that in 15 out of the competitive 20 auctions, the start price was 1 €. From a selfish equilibrium perspective, low start prices run the risk of low revenues via coordinated early or late bidding equilibria. The lower the start price the higher the potential gains from collusion.²⁵ While the social preference models suggest that a seller may choose a low auction price, acceptable to buyers, in order to get competitive bidding going, an offer of 1 € is smaller than what is needed to attract bidding (because it would give the buyer more than his fair share). So, in this sense, it may seem surprising that sellers do not make more use of their commitment power – regardless of whether preferences are selfish or social.²⁶ As it turns out, however, buyers were not able to take advantage of sellers’ low price-offers: Virtually all buyer payoffs were competed away. This rationalizes the use of even very low offers in CT.

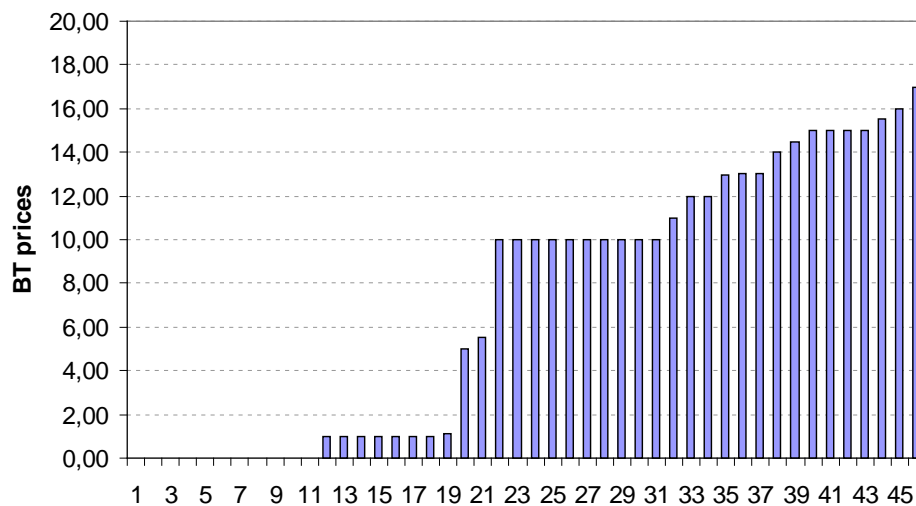
The occurrence of low revenues (the eight light bars in Figure 3) is not the result of failed competition. It is rather due to the corresponding sellers’ choices not to make use of buyer competition by offering a BIN price. All trades that ended with the acceptance of a BIN-offer ended with seller profits below or equal to 18 €. There are at least two possible reasons for

²⁵ This follows directly from the analysis by Ockenfels and Roth (2006); see also Klemperer (2000), who forcefully argues in the context of spectrum auctions that high reserve prices may prevent collusion.

²⁶ Three other sellers choose the equal split as auction start price, which may be interpreted as a ‘fair’ offer to attract bidders. (The other two start prices were 2 € and 3 €)

choosing BIN-prices below the competitive level. First, as explained before, fair-minded sellers may be willing to leave some money for the buyers. Second, sellers may underestimate the power of competition and believe that a BIN-offer below competitive level may yield higher revenues. The data do not allow separating between these reasons,²⁷ but there is evidence that some traders underestimated the implication of competition for prices. For one, there were three buyers who rejected the BIN-offer when the sellers chose the hybrid offer format, and instead submitted a bid which started the auction. In all three cases, the winning buyers later faced a final price that was well-above the initially offered BIN-price. This behavior is inconsistent with both standard and social preference theory. It appears that these buyers wrongly thought that they could get better prices in the auction. Also, the one market that ended with zero payoffs to the seller (Market 1 in Figure 3) was removed by eBay before the market was supposed to end, because of shill bidding on the part of the seller. As a consequence, all game payoffs were zero. This seller, too, seemed to have underestimated the robust power of competition, apparently seeing a need to push competition.

Figure 4. Realized prices (seller profit) in the 46 bilateral trading encounters



III.3 Buyer behavior in bilateral trading

Figure 4 sorts the final prices (seller profits) in the 46 BT encounters. Zeros are due to not accepted offers. The average final price (seller payoff) is 7.17 € compared to 17.76 € in CT. The hypothesis of equal revenues is strongly rejected (Mann Whitney U , two-tailed, $p < 0.001$).

²⁷ For instance, there is no evidence that those CT sellers who used BIN are less experienced, as measured by the feedback score.

There is also a much wider variance in BT final prices. Eleven encounters ended in zero profits, the maximum seller profit was 17 €. None ended close to the 20 € seller profit expected under selfish preferences.

Figure 5. Buyer responses in bilateral trading, by price offers

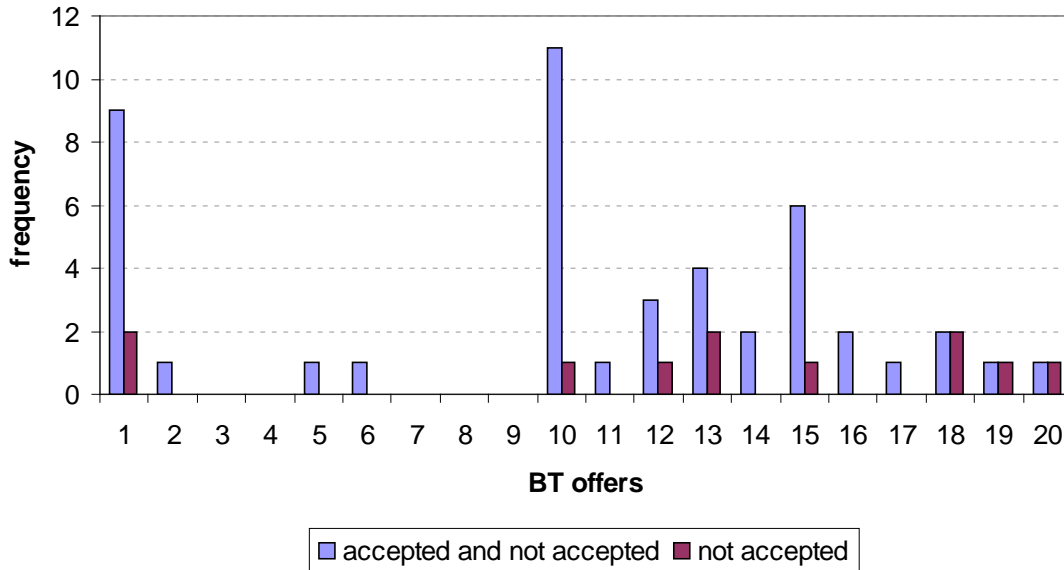


Figure 5 helps to understand the large variance in final price. It displays the pattern of offers accepted and not accepted in bilateral trading. An offer is defined as accepted when the buyer accepted the BIN price or submitted a winning bid in the auction, respectively. An offer is defined as not accepted when neither a BIN-offer is taken nor any valid auction bid is submitted. Of course, we do not know for sure why any particular offer was not accepted. One might think, for instance, that transaction costs could have deterred some bidders to accept an offer. Thus, after the experiment was completed, we asked all BT-buyers to briefly explain their strategy. We discuss the results at the end of this subsection.

Figures 4 and 5 mix a pattern of behavior common to laboratory observation with a pattern rarely seen in laboratories. Starting with the latter, there are 10 sellers who earn profits which lie strictly in between 0 € and the equal split. This phenomenon is the result of accepted price offers in this range, which seem to be partly explained by inexperienced sellers, inept at handling the lack of competition (as discussed in section III.1). Two of these price offers (both at 1 € as well as one equal split offer) are not accepted, even though they would have yielded high

absolute respectively equitable payoffs to the buyers.²⁸ While rejections of price offers below the equal split are sometimes observed also in laboratory research and are not necessarily inconsistent with social utility models,²⁹ the rejection of the fair offer is not captured by the social utility models. In any case, all three buyers told us after the experiment that they actually intended to submit a bid but then unfortunately missed eBay’s deadline for bidding. So, following the buyers’ explanation, not accepting those offers cannot be interpreted as “rejections”, but must rather be attributed to the loss of control in our field experiment. In *all* other cases of non-acceptance, the buyers who responded to our question told us that, given the low offer, they did not *want* to submit a bid (two buyers did not respond to our question). The rest of the buyer behavior looks pretty much like what has been observed in laboratory bargaining experiments. Overall, 24 percent (11/46) of the offers were not accepted, leading to zero seller and zero buyer payoffs. If one takes into account only offers equal or larger than the equal split, 26 percent (9/34) of the offers were not accepted. Furthermore, all four ‘greedy’ offers starting at 18 € were not acceptable to buyers. There is a correlation between the price offer and whether it was accepted (Spearman rho = .28, $p = 0.060$, two-tailed); when we drop those three offers that were unintentionally rejected or if we restrict ourselves to offers above the equal split, the correlation becomes stronger (rho = .50, $p = 0.001$, and rho = .379, $p = 0.027$, respectively).³⁰

The reader may have noticed that the distribution of offers accepted and not accepted in Figure 5 is in parts somewhat shifted to the right compared to the distribution of the initial price offers in Figure 1. There are two reasons. The first reason is shill bidding: one of the sellers with an auction start price of 1 € submitted shill bids to increase the price to 5.50 € (see section III.1). So, the initial price offer differed from the final offer eventually accepted by the buyer. The second reason is that in five out of the 18 cases in which BT-sellers selected the hybrid selling format (auction with BIN-option), buyers accepted the BIN price even though it was higher than

²⁸ Because eBay employs second-price auctions, the final price in an auction with only one bidder is equal to the start price, regardless of the number of bids and the bid amounts, as long as the highest bid is larger than the start price. So, in these cases, even fair-minded buyers could not increase the price in order to get more equitable outcomes.

²⁹ Several other experimental studies including Bornstein and Yaniv (1998), Güth et al. (2007), Gehrig et al. (2007), Hennig-Schmidt et al. (2001), Mitzkewitz and Nagel (1993), Bolton et al. (2006) and Roth et al. (1991) report rejections of offers that are advantageous to the buyer. Bolton and Ockenfels’ (2000) social utility model is consistent with this phenomenon.

³⁰ Our rejection rate is close to what is observed in the Roth et al. (1991) laboratory ultimatum bargaining game, where 26.4 percent of the offers were rejected (there was no significant second mover experience effect). However, because there is heterogeneity across standard laboratory ultimatum games (e.g., Slonim and Roth 1998), we do not make quantitative claims here; calibrations of social utility functions are presented in Fehr and Schmidt (1999) and DeBruyn and Bolton (2008), among others.

the auction start price.³¹ Thus, because price offers in Figure 1 are defined as the minimum price that can be accepted and Figure 5 only shows offers that are eventually accepted, the distributions differ.

We suspect that the phenomenon of BT-buyers who take the higher BIN price could be partly caused by risks of ‘irregular’ bidding in the field, which are controlled away in the laboratory. The risks come from fraudulent shill bidding or from active bidders who are not registered for the experiment and just bid for fun or out of curiosity.³² In fact, we did not only have two rather unambiguous cases of shill bidding, but four more eBay users, not registered at our experiment, who submitted bids to our experimental offers (none of the bids won). The BIN offer might command a premium to protect the BT-buyer from these risks (although the risk is small: there is only one instance of bidding by unregistered traders in BT). This explanation implies that those who took BIN offers would prefer to close the transaction quickly. In fact, BIN-offers were significantly more likely to be accepted earlier than start-price offers as measured in seconds before the offer terminates (Mann-Whitney U , two-tailed, $p = 0.023$).³³

Finally, we comment on the hypothesis that uncontrolled transaction costs might play in the interpretation of our field experiment. In fact, if buyers differ with respect to transaction costs associated with bidding, the probability that a BT-offer is accepted would decrease with the price-offer, because a buyer’s payoff must be larger than the corresponding transaction costs for a trade to occur. At the same time, prices are high in CT if there is a sizeable fraction of bidders in the population with negligible transaction costs. So, some of the patterns we observe in the experiment might be due to factors not related to preferences.

However, it is very unlikely that transaction costs explain the behavior of our BT-buyers. As we mentioned before, buyers who did not accept a small share of the gains from trade explicitly told us that they were rejecting an unfair offer or unfair behavior; the three exceptions concern the price offers at or below the equal split. Social concerns also explain why the modal offer is the equal split, which would not easily be explained by transaction costs arguments. Even more importantly, transaction costs of bidding were negligible by the experimental design. For one, the overall transaction costs of participating in the experiment are lower than

³¹ The average difference between BIN and auction start price in these 5 cases is 3.30 €

³² We followed a chat room discussion of two eBay bidders, not registered at the experiment, discussing the possibility of submitting bids on our auctions just out of curiosity.

³³ One reason for eBay buyers to prefer BIN-offers is impatience. In our experiment, all subjects were paid at the same time after all auctions ended, so in this sense impatience cannot be part of the explanation.

participating in typical laboratory experiments (while the stakes are significantly higher). Traders did not have to be physically present at a certain time or certain place, and all communication was done asynchronously and electronically; in particular, bidding could typically take place at any time during a whole week.³⁴ An experienced eBay buyer could have completed the whole experiment, including registration, reading instructions and bidding, in probably 5-15 minutes. The actual bidding, once bidders are registered and sit in front of the computer, may take probably between 10 seconds up to a minute, mostly depending on whether the bidder is already logged in on eBay, or not.³⁵ (Unlike for eBay transactions outside our experiment, there is no need for searching the platform, researching the value of the offered product, or assessing the ‘trustworthiness of the offer’.) Recall also that we only invited experienced eBay traders, who already had an eBay account and successfully completed transactions on eBay, and who were available at the time of the experiment (which was explicitly required in the invitation email). In order to further minimize transaction costs, we conducted the experiment during a University teaching period in Germany so that traders were unlikely to be on vacation and had easy and frequent access to the Internet (either at the University or at home), we wired all payoffs with no confirmation or signature required, we paid all pecuniary transactions costs (eBay and bank fees), and there was no shipping and handling involved in our eBay transactions. It seems hard to imagine a field experiment that establishes lower or less controlled opportunity costs of trading. Finally, complementary laboratory data reinforce our point: transaction costs are not necessary to produce social bargaining patterns – although transaction costs could of course reinforce the behavior.

III.4 Reputation building behavior

eBay employs an electronic reputation mechanism that enables buyers and sellers to leave feedback (“positive”, “negative”, or “neutral”) on each other after a successful transaction.³⁶

³⁴ One might argue, however, that if bidders have an incentive to bid late, the transaction costs associated with bidding at a certain time can be quite high. However, as we have mentioned before, there is no incentive to bid late in our BT environment, and in fact we do not see last-minute bidding in BT. Moreover, if bidders still wanted to bid late but could not do so because of high transaction costs, their best response is not to reject a low offer, but rather to bid early, when transaction costs are small. (Ockenfels and Roth 2005 show that bidding early is an equilibrium.) Alternatively, they could use artificial sniping agents that automatically place bids in the closing seconds of eBay auctions.

³⁵ Traders knew the stakes involved and the ‘games being played’ *before* registering. So if transaction costs prevented some traders to take part in the experiment, one could expect that these traders chose not to go through the registration procedure, which was much more time consuming than the actual bidding.

³⁶ See Resnick and Zeckhauser (2002) for an early description and analysis of eBay’s mechanism.

Leaving feedback on the trading partner can on the one hand confound social behavior, and on the other hand be another effective way of expressing social behavior. Furthermore, the “reputation” of a trader, that is, the percentage of positives gained in non-experimental ‘natural’ transactions, may have predictive value for behavior in our experiment. Here, we investigate both hypotheses.

A number of field and lab studies show that eBay’s feedback score tends to affect both future revenues and probabilities of sale (Dellarocas 2006, Resnick et al. 2006, Bolton et al. 2004). Furthermore, some eBay observers inferred from field data that not giving feedback may be an indicator of an unsatisfactory transaction. Dellarocas and Woody (2008), for instance, state that their “results confirm the wide-spread belief that eBay traders are substantially more likely to post feedback when satisfied than when dissatisfied” (see also Klein et al. 2007). Part of the reason is fear of retaliatory negative feedback: giving negative feedback increases the probability of receiving negative feedback (Bolton et al. 2010). If our experiment captures natural trading patterns, these field observations suggest that we should observe buyers in our bilateral bargaining environment to leave positive feedback when their bargaining payoff was high and to leave no or negative feedback when the payoff was small.³⁷

Overall, in the bilateral trading environment, we observe that the propensity to give feedback in our experiment is roughly in line with Resnick and Zeckhauser's (2002) field observations. In their huge data set obtained from eBay, buyers commented on sellers for 52 percent of the items (52 percent in our experiment), sellers on buyers 61 (73) percent of the time, and 99.1 (98.4) percent of the feedback provided by buyers was positive.

In our BT condition, 11 offers were rejected, so we have 35 encounters where feedback could be given. From these, 16 buyers left no feedback on the respective seller, 18 left positive feedback and one left a neutral feedback;³⁸ 25 BT-sellers left feedback. We find that the average buyer payoff when no feedback was given is 9.18 € while the average payoff when positive feedback was given is 11.76 €. The difference is statistically significant (Mann Whitney *U* test, one-tailed, $p = 0.042$). This observation supports the hypothesis in the literature – that “silence”

³⁷ Of course, in more complex transaction, satisfaction with a transaction may depend on many dimensions, such as quality of the item, quality of communication, shipping time and costs etc. In our experiment, however, the only dimension which can plausibly be the basis of feedback giving is the distribution of pecuniary payoffs.

³⁸ The neutral one occurred in the bargaining with shill bidding; the buyer left a comment “I suspect shill bidding. Otherwise ok.” (translation from German).

correlates with dissatisfaction – in a controlled field environment. It also shows that BT-buyers do not only punish by rejecting but also use the feedback mechanism as a reciprocation device.³⁹

One might at first glance suspect that a concern for a good reputation is sufficient to explain the social behavior we observe in BT. The prospect of being rewarded with a good feedback may be a good incentive for traders to behave ‘fairly’. But such reputational concerns *cannot* explain the key phenomenon on the buyer side: rejections in the BT-treatment. The reason is that on eBay, traders can only give feedback after the seller’s offer is actually accepted. This implies that by rejecting unfair offers buyers not only forgo monetary profits, but they also make it impossible to improve their own reputation, and to use eBay’s feedback forum to articulate one’s dissatisfaction about the seller. That is, rejections occur even though they *prevent* reputation building.⁴⁰

The way reputation building interacts with trading behavior raises the question whether a trader’s reputation information carries information about individual trading inclinations, which in turn have predictive value for the trading strategies employed in our experiment. In particular, one might hypothesize that ‘greedy’ BT-sellers tend to generally have more problems with eBay buyers, outside experimental control, and thus are more likely to have lesser eBay reputations. However, the effect, if any, can be expected to be small, because unfair offers may not be accepted in the first place and traders with relatively low reputations are avoided and probably even crowded out. In fact, the data point to the right direction, but the effect is not quite significant at conventional levels. The average reputation (that is, percentage positive) of those who offer the equal split is 99.8 percent, while the average reputation of those sellers who demand more is 99.6 percent, which is however not significant when applying a Mann Whitney

³⁹ In CT, we have 28 encounters in which feedback could have been given. 14 buyers and 21 sellers gave feedback, and all feedback was positive. There is no significant correlation between buyer payoffs and feedback giving.

⁴⁰ Also, observe that the trader who naturally wants to build up a reputation in our bargaining game is the buyer: he could profit from being known as somebody who rejects low offers. However, on eBay, the seller cannot condition his offer on the buyer’s reputation. Rather, the seller must make the offer without knowing the buyer and the buyer’s reputation, so that this kind of reputation building is not valuable to buyers. Our experimental design reflects this. Sellers did not know anything about buyers before actually posting their binding offer on eBay. (Moreover, we insisted that all offers are made public simultaneously, up to the second, so to keep seller behavior independent.) Regarding anonymity issues, we minimized the probability that traders know each other, or have a common history or future on eBay by matching traders from different cities. Finally, all traders were asked not to communicate with each other: (a) sellers were told to post their offers under a specific eBay category, and to use a specified item description, in order to suppress any communication between traders, and (b) all traders were asked to send any questions to us. Accordingly, we have no indication that a trader tried to directly contact other traders that could have diminished anonymity among traders or that matching partners did know more about opponents than what is revealed on eBay, which is eBay-identities and feedback.

U test. The Spearman rank correlation between seller reputation and corresponding price offers yields -0.205 at a one-tailed significance level of 0.086 .

An alternative interpretation of the correlation could be that those who demand more for themselves are more experienced, and that more experience implies a higher probability of having received a negative feedback, just because more experience implies more completed transactions. In fact, the correlation between reputation and feedback score (as a measure of experience) is significantly negative (Spearman $\rho = -0.270$, one-tailed $p = 0.035$). Moreover, the average feedback score of those who offer the equal split is 366 , while the average score of the others is $131:153$ for those who demand more than the equal split, and 95 for those who demand less than the equal split (Mann Whitney *U* test, one-tailed, $p = 0.099$). More experience may make you more familiar with the strategic incentives, but it may also make you more sensitive to the concerns of the trading partner.

Finally, we ask whether and how buyers' acceptance behavior is correlated with reputation scores. Are buyers who do not accept an offer in our bilateral bargaining environment more or less experienced than accepting buyers? It can be argued either way. On the one hand, rejecting may be a kind of mistake by inexperienced buyers. On the other hand, there is evidence indicating that those who are experienced with markets may be more fair-minded (Henrich et al. 2001). Measuring experience by the feedback score, there is, however, no significant difference in the data. Do buyers who do not accept have a better or worse reputation? It may be that they are harder to make happy as a buyer (the resulting conflict leading to more negative feedback). Or maybe they behave in a fairer manner as a seller and expect the same treatment themselves. Measuring reputation as the percentage of positive feedback (as done by eBay), however, again yields no significant differences in the buyer data.

IV. Related literature and discussion

One of the first online-auction experiments was conducted by Lucking-Reiley (1999), who tested Vickrey's revenue-equivalence theory by auctioning off "Magic game cards" to 'real' traders on self-engineered Internet auction platforms. His methodological approach differs from ours in several respects. Most importantly, Lucking-Reiley was willing to give up significant experimental control. As he puts it (p. 1078): "My field experiments conduct unconditional tests of a theory's predictions, because I do not observe whether or not the underlying assumptions of the theory are true. ... Field tests assess the practical predictive power of a theory, since most

theoretical assumptions in economic models are intrinsically unobservable in practice.” This reasoning is compelling when, as in Lucking-Reiley’s study, holding the underlying model assumptions (for auctions concerning beliefs over the distribution of values and number of potential buyers, risk preferences, etc.) constant across treatments is sufficient to obtain a clear test. However, such an assumption is not sufficient for testing the hypotheses we deal with here; the reason being that the nature of the underlying values is precisely what is in question. For example, if we had not induced (pecuniary) values for the trading object, then the interpretation of a non-acceptance in the bilateral treatment as evidence of social preferences would be confounded with the possibility that it simply reflects a low pecuniary value for the traded good of the respective buyer. Which interpretation is valid has both theoretical and practical consequences for trader strategies and market efficiency: whether or not, in market circumstances with little competition, trader transgression of accepted norms of fairness leads to market inefficiency.

How then can we know both that an experimental result is reflective of results in the fully natural environment and that the result stems from the underlying causes posited by the model? One way is a series of complementary field experiments, with controls ranging from tight to very loose. If we observe essentially the same behavior across these experiments, we can be confident that the result reflects both natural behavior and posited causes. While a complete picture has yet to emerge, our study is complementary, in this manner, to other field (as well as the lab experiments cited earlier) already in the literature.

Our study is a “framed field experiment”, because of its “non-standard” subject pool, its field context, and because of subjects knowing that they are participating in an experiment (List 2006a). The study is closely related and mostly complementary to a series of framed field experiments by List. First, building upon the classroom experiments by Chamberlain (1948) and Smith (1962; see also Plott, 1982, and Roth, 1995b), List (2004) observes that neoclassical competitive market theory explains the equilibrating forces in a double-auction both in laboratory markets and in framed field experiment conducted in sports card and collector pin marketplaces. This finding is further reinforced by our finding in the CT environment that competitive behavior is robust to the anonymous and complex electronic trading institutions as evolved on eBay, and to noisy, partly out-of-equilibrium seller behavior. One difference to List’s work is that he induces competition on *both* market sides, and then has buyers engage in unstructured, face-to-face bilateral bargaining with sellers until they make a contract or the respective trading period

terminates. One implication of the design is that his and the corresponding earlier experiments lack the control necessary to make the kind of conclusions we do. For example, in his specific market environment, it is difficult – if not impossible – both to derive clear-cut (non-cooperative) game theoretic equilibria, even if we rule out further complicating social preferences, and to isolate the respective roles of the bilateral bargaining and multilateral competition institutions on behavior and outcomes – which is a major, complementary goal of our study.

In a second related paper, List (2006b) bridges lab and field behavior in carefully controlled lab and field gift-exchange environments, using sports card traders in a real sports card market.⁴¹ His data suggest that the laboratory patterns are rather stable in framed field experiments, that is, against changes in subject pools, communication channels, framing and information scenarios. Again, our study reinforces this finding in a different institutional and behavioral context. For instance, List (2006b) studies positive reciprocity while our work is mainly concerned with negative reciprocity (e.g., a buyer can reject if the seller's price-offer is perceived as unfair).⁴² Moreover, our framed field experiment also supports List's observation in his natural field experiment as well as the result of numerous laboratory experiments, that reputational concerns may play an important role in market transactions in that they can trigger behavior that is consistent with social preferences.⁴³ While our BT-traders use their influence on the trading partner's reputation to reward fair and punish relatively unfair behavior, which may create additional incentives to make fair offers, reputation effects cannot explain rejections of unfair offers, the kind of negative reciprocity in bilateral bargaining that we observe in our study (see section III.4).

Finally, List (2006b) also conducted natural field experiments, where sellers did not know that they are taking part in an experiment, and where he identifies a much lesser role for social preferences. On the other hand, Falk (2007), Chen et al. (2010), Card et al. (2010) and Ockenfels et al. (2011) found, in very different contexts, a larger role for social preferences in natural field experiments. We did not conduct a natural field experiment and so cannot directly contribute to

⁴¹ See also Gneezy and List (2006), who investigated social preferences in a field experiment of a labor market, which resembles a laboratory gift-exchange game environment.

⁴² Laboratory results suggest an asymmetry regarding positive and negative reciprocity in that cooperation in positive reciprocity games, such as trust games or public goods games, tend to be less stable (Gneezy and List 2006 for gift exchange games, Ledyard 1995 for public good games, Bolton et al. 2004 for trust games), while negative reciprocity tends to be more robust (Roth 1995 for ultimatum games, Fehr and Gächter 2000 for public good games with punishment).

⁴³ For laboratory experiments see, e.g., Camerer and Weigelt (1988), McKelvey and Palfrey (1992), Jung et al. (1994), Andreoni and Miller (1999), Brandts and Figueras (2003), and Bolton et al. (2004).

this question. However, on the one hand even natural field experiments like List's (2006b) insert controls which might influence behavior and outcomes,⁴⁴ and on the other hand it is almost inevitable to lose some control in natural compared to framed experiments.⁴⁵ Therefore, we propose to also follow a complementary research venue: by exploiting the information about past trading patterns available on eBay, we investigate whether there are behavioral patterns in the controlled experiment that correspond to fully naturally occurring trading strategies devised by the same individual. For instance, we have shown that the choice of the trading format and other strategies in the experiment can be partly linked to trading patterns in the uncontrolled field. Moreover, the fact that there is a correlation between BT price offers and receiving positive feedback in the framed experiment on the one hand, and feedback received earlier in the fully natural environment on the other hand (namely the seller reputation, although this latter correlation is only weakly significant) seems to suggest that some of the social patterns observed under experimental control mirrors fully natural behavior. However, more research employing these different methodological venues is needed to robustly establish or reject the various potential links.

There is also a related and fast-growing literature dealing with the economics of the eBay market (e.g., Lee and Malmendier, forthcoming; see Bajari and Hortacsu 2004, and Ockenfels et al. 2006 for surveys). Most papers in this literature test auction theory, both in the uncontrolled field (e.g., Ockenfels and Roth 2006, Bajari and Hortacsu 2003) and in the fully controlled laboratory (e.g., Ariely et al. 2005). Garrat, Walker and Wooders (2004) also investigated bidding behavior of actual eBay buyers and eBay sellers. They self-engineered an experimental second-price sealed-bid auction platform on the Internet and found that whether or not traders have experience with online auctions affect bidding behavior. A number of papers examine the use of the BIN-option on eBay. The literature mainly focuses on impatience and time preferences (e.g., Mathews 2004) and risk with respect to the uncertainty of buyer valuations (e.g., Reynolds and Wooders 2003). On the other hand, our study suggests that risk with respect to 'irregular' bidding on the buyer side, underestimating the power of competition and social

⁴⁴ For instance, in List's (2006b) natural field experiment, buyers were instructed to approach 10 different, pre-selected dealers, to make a pre-determined offer for a specified card, and to avoid haggling.

⁴⁵ For instance, in List's (2006b) natural field experiment, sellers' beliefs about the degree of seller competition was not controlled for (sellers may have wrongly thought that if they do not close the current deal at a good profit, some other seller would do it), whereas in the corresponding laboratory experiments it was known that each buyer was allowed to only approach one seller, more closely controlling the competitive environment and social embeddedness – as it is also typically the case in laboratory gift-exchange games.

preferences on the seller side may also contribute to the success of eBay's BIN-format. A few papers have used the actual eBay platform for controlled experiments. Jin and Kato (2008; see also Jin et al. 2010) conducted a field experiment on eBay, examining the link between price, quality, seller claims and seller reputation in markets for baseball cards. They find, among other things, that some buyers are misled by non-credible claims of quality, paying higher prices but do not receive better quality. They also partly attribute their data patterns to loopholes in the eBay reputation mechanism, which also is an area of active research (Dellarocas 2006 provides a survey). Here too, one stream of papers uses uncontrolled field data (e.g., Resnick and Zeckhauser 2002), and others brought eBay's reputation system into the fully controlled laboratory (e.g., Bolton et al. 2004). One paper which is close to ours in methodology, is Resnick et al.'s (2006) field experiment on eBay, investigating the impact of (manipulated) differential feedback reputation scores on revenue. The design of the feedback system is investigated with the help of laboratory and natural field experiments by Bolton et al. (2010). None of these online market studies investigates the impact of social preferences on bidding and trading behavior, and none is conducted to identify trading patterns that connect controlled laboratory and fully uncontrolled market behavior.

V. Conclusions

The data from our framed field experiment on eBay demonstrate that the social preference behavior observed for buyers under lab conditions extend to experienced traders operating in their natural trading system. These social preferences may be obscured in an uncontrolled environment by competition or incomplete information about the value of goods being traded, factors that cloud the visibility of social preferences in theory and in the lab. Our data confirm that buyers in the bilateral trading environment are prepared to reject unfair price offers. The social behavior cannot be explained by reputational concerns. Buyer competition, on the other hand, robustly yielded highly competitive outcomes, dissipating all buyer rents and giving sellers 2.5 times higher revenues than without buyer competition.

However, neither selfish nor social equilibria satisfactorily organize *all* behaviors. Maybe the most notable out-of-equilibrium phenomenon in our field study is that some sellers fail to fully use their commitment power and their information about bidder numbers and valuations. In particular, many auction start prices in both trading environments are too low – regardless of whether measured by standard or social preference theory. The out-of-equilibrium seller

behavior largely reflects traders' market (in)experience outside the experiment. While this hurts in bargaining, it turns out not to hurt in the competitive trading environment. In fact, it tends to strengthen the dichotomy observed in the laboratory. On the one hand, bargaining revenues are even more dispersed than in the laboratory. On the other hand, competitive trading not only masks social preferences but is also robust against the increased strategic and institutional complexities on eBay. There are miscalculations, noise and malfeasance behind otherwise competitive looking behavior.

One of the principal lessons the present work has for the vetting of laboratory results in the field is the need for experiments that accommodate clear (game) theoretical predictions, and thus losing control in a very controlled way. The greater complexity of the field environment produces a far richer pattern of behavior than observed in the lab. Some of this behavior is an optimal strategic response to the more elaborate rules of the natural market. Yet greater complexity also opens the door to increased out-of-equilibrium mistakes and noise, so that the underlying reasons for behavior can appear different than they actually are – unless the experiment is sufficiently controlled to compare and parse explanations.

Control implies abstracting away from real-world complexities. In particular, the eBay traders in our controlled field experiment know more about the potential economic surplus from successful transactions than a typical eBay trader outside our experiment. Social preference models and laboratory data suggest that therefore social behavior is less an issue in the uncontrolled environment than it is in our controlled environment. However, this design choice was motivated by a desire to separate two competing explanations for selfishness in less controlled field studies: lack of social preferences or lack of information about comparison standards.

Moreover, we find that the behavior in our controlled field experiment is intertwined with uncontrolled field behavior. A priori experience with eBay predicts our sellers' format choice, the price offer and other trading strategies. Past trades on eBay also affect the sellers' eBay reputation, which is (weakly) related to the experimental bargaining behavior. The influence also goes in the other direction; behavior in the experiment affects the sellers' eBay reputation, and thus the outcome of future transactions. Thus, the trading patterns we observe in our experiment predict and are predicted by trading patterns employed outside any experimental control; they mirror 'natural' trading patterns. At the same time, our experiment demonstrates that, while eBay's reputation building stage cannot explain the social behavior we observe, it is used as an

additional opportunity to bring social concerns into play. This way, preferences and market institutions interact to promote social behavior. While our experimental set-up intervenes in the field in various ways, it is exactly this experimental control that allows us to make the connection between the controlled laboratory world and the uncontrolled world of economic transactions.

A full understanding of the role of social preferences in markets requires a series of complementary studies. Our study is open to relatively easy replication, as well as extension in various directions. Some extensions involve a changing or further loosening of control to check for the marginal effects of such factors as incomplete information about values or selling ‘real’ items, allowing more social interaction and communication between traders, endogenizing the number of traders, not telling subjects that they are participating in an experiment etc. For instance, social behavior has more dimensions than outcome-based social comparison, including reciprocity and procedural fairness. It would be interesting to investigate whether and how the role of these other social concerns changes as information about others' payoffs is gradually taken away, in a controlled way. With the emergence of online trading platforms, a highly controlled and flexible analysis of field behavior is feasible.

References

- Andreoni, James, Marco Castillo, Ragan Petrie (2003), “What Do Bargainers' Preferences Look Like? Experiments with a Convex Ultimatum Game.” *American Economic Review*, 93(3), 672-685.
- Andreoni, J., and J.H. Miller (1993). “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma: Experimental Evidence.” *Economic Journal*, 103, 570-85.
- Ariely, D., A. Ockenfels and A.E. Roth (2005). "An Experimental Analysis of Ending Rules in Internet Auctions." *RAND Journal of Economics*, 36(4), 790-809.
- Bajari, P. and A. Hortaçsu (2004). “Economic Insights from Internet Auctions.” *Journal of Economic Literature*, XLII, 457-486.
- Bajari, P. and A. Hortaçsu (2003). "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions." *RAND Journal of Economics*, 34(2), 329-55.
- Bolton, G.E. (1991). "A Comparative Model of Bargaining: Theory and Evidence." *American Economic Review*, 81(5), 1096-1136.
- Bolton, G.E. (1999). “Weak and strong equity effects.” In D. Budescu, I. Erev and R. Zwick, eds., *Games and Human Behavior*, Erlbaum: Mahwah, NJ, 203-236.
- Bolton, G.E., J. Brandts and A. Ockenfels (2005). “Fair Procedures: Evidence from Games Involving Lotteries.” *Economic Journal*, 115 1054–1076.

- Bolton, G.E., B. Greiner and A. Ockenfels (2010). "Engineering Trust: Reciprocity in the Production of Reputation Information." Working paper, University of Cologne.
- Bolton, G.E., E. Katok, and A. Ockenfels (2004). How Effective are Online Reputation Mechanisms? An Experimental Study. *Management Science*, 50(11), 1587-1602.
- Bolton, G.E., and A. Ockenfels (1998): "Strategy and Equity: An ERC-Analysis of the Güth-van Damme Game." *Journal of Mathematical Psychology*, 42, 215-26.
- Bolton, G.E., and A. Ockenfels (2000). "ERC - A Theory of Equity, Reciprocity and Competition." *American Economic Review*, 90(1), 166-193.
- Bolton, G.E. and A. Ockenfels (2005). "A Stress Test of Fairness Measures in Theories of Social Utility." *Economic Theory*, 25(4), 957-82.
- Bolton, G.E. and A. Ockenfels (2006). "Measuring Efficiency and Equity Motives - A Comment on Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments." *American Economic Review*, 96, 1906-1911.
- Bolton, G.E., A. Ockenfels, and U. Thonemann (2011). Managers and Students as Newsvendors. Working paper, University of Cologne.
- Bornstein, G., and I. Yaniv (1998). "Individual and Group Behavior in the Ultimatum Game: Are Groups More Rational Players?" *Experimental Economics*, 1, 101-8.
- Brandts, J., and N. Figueras (2003). "An Exploration of Reputation Formation in Experimental Games." *Journal of Economic Behavior and Organization*, 50, 89-115.
- Burns, P. (1985), "Experience and Decision Making: A Comparison of Students and Businessmen in a Simulated Progressive Auction." In V. Smith, ed., *Research in Experimental Economics Vol. 3*, Greenwich: JAI Press, 139-157.
- Camerer, C.F. (2003). "*Behavioral Game Theory: Experiments in Strategic Interaction*." Princeton University Press: Princeton.
- Camerer, C., and K. Weigelt (1988). "Experimental Test of a Sequential Equilibrium Reputation Model." *Econometrica*, 56, 1-36.
- Card, D., A. Mas, E. Moretti, and E. Saez (2010). "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction." NBER Working Paper 16396.
- Chen, Y., M. Harper, J. Konstan and S.X. Li (2010). "Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens." *American Economic Review*, 100(4), 1358-98.
- Chamberlain, E. (1948). "An Experimental Imperfect Market." *Journal of Political Economy*, 56, 95-108.
- Charness, G., and M. Rabin (2002). "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics*, 117, 817-869.

- Cooper, D., and J. Kagel (forthcoming). "Other-regarding preferences." In: J. Kagel and A. Roth (eds.), *The Handbook of Experimental Economics*, Volume 2, in preparation.
- Cooper, D., Kagel, J., Lo, W., and Gu, Q., (1999). "Gaming against Managers in Incentive Systems: Experimental Results with Chinese Students and Chinese Managers." *American Economic Review*, 89(4), 781-804.
- Cohen, A. (2002). *The Perfect Store: Inside Ebay*. Little, Brown and Company.
- DeBruyn, A., and G.E. Bolton (2008). "Estimating the Influence of Fairness on Bargaining Behavior." *Management Science*, 54, 1774-1791.
- Dellarocas, C. (2006). "Reputation Mechanisms." *Handbooks in Information Systems*, Vol. I, Handbook on Economics and Information Systems, edited by Terrence J. Hendershott, Elsevier, 629-660.
- Dellarocas, C. and C. A. Wood (2008). "The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias." *Management Science* 54(3), 460-476.
- Dufwenberg, M., and G. Kirchsteiger (2004). "A Theory of Sequential Reciprocity." *Games and Economic Behavior*, 47, 268-98.
- Dufwenberg, M. and U. Gneezy (2000). "Price Competition & Market Concentration: An Experimental Study." *International Journal of Industrial Organization*, 18, 7-22.
- Dyer, D., J. Kagel and D. Levin (1989), "A Comparison of Naive and Experienced Bidders in Common Value Offer Auctions: A Laboratory Analysis." *Economic Journal*, 99(394), 108-115.
- Engelbrecht-Wiggans, R. (1987). "On Optimal Reservation Prices in Auctions." *Management Science* 33 (6), 763-770.
- Falk, A. (2007). "Gift-exchange in the Field." *Econometrica*, 75 (5), 1501-1511.
- Fehr, E. and S. Gächter (2000) "Cooperation and Punishment in Public Goods Experiments." *American Economic Review*, 90(4), 980-994.
- Fehr, E., and K.M. Schmidt (1999). "A Theory Of Fairness, Competition, And Cooperation." *The Quarterly Journal of Economics*, 114(3), 817-868.
- Forsythe, R., J.L. Horowitz, N.E. Savin and M. Sefton (1994) "Fairness in Simple Bargaining Experiments." *Games and Economic Behavior*, 6, 347-69.
- Garratt, R., M. Walker, and J. Wooders (2004). "Experienced bidders in online second-price auctions." Working paper, University of California.
- Gehrig, T., W. Güth, V. Levati, R. Levinsky, A. Ockenfels, T. Uske, T. Weiland (2007). "Buying a pig in a poke: An experimental study of unconditional veto power." *Journal of Economic Psychology*, 28(6), 692-703.
- Gneezy, U., and J.A. List (2006). "Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets using Field Experiments." *Econometrica*, 74(5), 1365-1384.

- Güth, W., C. Schmidt, and M. Sutter (2007). "Bargaining outside the lab - a newspaper experiment of a three-person ultimatum game." *Economic Journal*, 117, 449-469.
- Güth, W., R. Schmittberger and B. Schwarze (1982). "An Experimental Analysis of Ultimatum Bargaining." *Journal of Economic Behavior and Organization*, 3(4), 367-388.
- Harrison, G., and J.A. List (2004). "Field Experiments." *Journal of Economic Literature*, XLII, 1009-1055.
- Hennig-Schmidt, H., Z.-Y. Li, and C. Yang (2001). "Non-monotone strategies in ultimatum bargaining: first results from a video experiment in the people's Republic of China." *Proceedings of the International Congress of Mathematicians Game Theory and Applications Satellite Conference*, Qingdao Publishing House, 225-31.
- Henrich, J., R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, and R. McElreath, (2001). "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies." *American Economic Review*, Papers and Proceedings, 91, 73-78.
- Jin, G.Z., and A. Kato (2008). "Price, Quality and Reputation: Evidence from an Online Field Experiment." *RAND Journal of Economics*, 37 (4), 983 - 1005.
- Jin, G.Z., A. Kato, and J. List (2010). "That's News to Me! Information Revelation in Professional Certification Markets." *Economic Inquiry*, 48(1), 104-122.
- Jung, Y.J., J.H. Kagel, D. Levin (1994). "On the Existence of Predatory Pricing: An Experimental Study of Reputation and Entry Deterrence in the Chain Store Game." *The RAND Journal of Economics*, 25, 72-93.
- Kagel, J. H. (1995) Auctions: A Survey of Experimental Research. In: J. H. Kagel & A. E. Roth (Eds.). *The Handbook of Experimental Economics*. 501-585. Princeton University Press. Princeton, NJ.
- Kagel, J., C. Kim and D. Moser (1996). "Fairness in Ultimatum Games with Asymmetric Information and Asymmetric Payoffs." *Games and Economic Behavior*, 13(1), 100-110.
- Klein, T. J., C. Lambertz, G. Spagnolo, and K. O. Stahl (2007). Reputation Building in Anonymous Markets: Evidence from eBay. Working paper, University of Mannheim.
- Klemperer, P. (2004). "Auctions: Theory and Practice." Princeton University Press.
- Ledyard, J. (1995). "Public goods: a survey of experimental research." In: Kagel, J. and Roth, A. E. (Eds.), *Handbook of Experimental Economics*, Princeton, 111-194.
- Lee, H., and U. Malmendier (forthcoming). "The Bidder's Curse." *American Economic Review*.
- Levitt, S.D. and J.A. List (2007). "What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World?" *Journal of Economic Perspectives*, 21(2), 153-174.

- List, J.A. (2006a). "Field Experiments: A Bridge between Lab and Naturally Occurring Data." *Advances in Economic Analysis & Policy*, 6(2), available at: <http://www.bepress.com/bejeap/advances/vol6/iss2/art8>.
- List, J.A. (2006b). "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions." *Journal of Political Economy*, 114(1), 1-37.
- Lucking-Reiley, D. (1999). "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet." *American Economic Review*, 89(5), 1063-1080.
- Ma, C.A., and M. Manove (1993). "Bargaining with Deadlines and Imperfect Player Control." *Econometrica*, 61(6), 1313-39.
- Mathews, T. (2004). "The Impact of Discounting on an Auction with a Buyout Option: a Theoretical Analysis Motivated by eBay's Buy-It-Now Feature." *Journal of Economics*, 81(1), 25-52.
- McKelvey, R.D. and Palfrey, T.R. (1992). "An Experimental Study of the Centipede Game." *Econometrica*, 60, 803-836.
- Namazi, A. and A. Schadschneider (2006). "Statistical properties of online auctions." *International Journal of Physics*, C17, 1485.
- Ochs, J., und A.E. Roth (1989). "An Experimental Study of Sequential Bargaining." *American Economic Review*, 79, 355-84.
- Ockenfels, A. and Roth, A. E. (2006). "Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction." *Games and Economic Behavior*, 55(2), pp. 297–320.
- Ockenfels, A., D. Reiley and K. Sadrieh (2006). "Online Auctions." *Handbooks in Information Systems*, Vol. I, Handbook on Economics and Information Systems, edited by Terrence J. Hendershott, Elsevier, 571-628.
- Ockenfels, A., D. Sliwka and P. Werner (2011). "Bonus Payments and Reference Point Violations." Working paper, University of Cologne.
- Plott, C. (1982). "Industrial Organisation Theory and Experimental Economics." *Journal of Economic Literature*, 20, 1485-1527.
- Rabin, M. (1993) "Incorporating Fairness into Game Theory and Economics." *American Economic Review*, 83, 1281-302.
- Reiley, D. H. (2006). "Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet. *RAND Journal of Economics*, 37(1), 195-211.
- Resnick, P., and R. Zeckhauser (2002). "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System." In: Michael R. Baye, ed. *The Economics of the Internet and E-Commerce (Advances in Applied Microeconomics, Vol. 11)*. JAI Press.

- Resnick, P., R. Zeckhauser, J. Swanson, and K. Lockwood (2006). "The Value of Reputation on eBay: A Controlled Experiment." *Experimental Economics*, 9 (2), 79-101.
- Reynolds, S. S., and J. Wooders (2003). "Auctions with a Buy Price." Working Paper. University of Arizona.
- Roth, A.E. (1995a). "Bargaining Experiments." In John H. Kagel and Alvin E. Roth, eds., *Handbook of Experimental Economics*. Princeton: Princeton University Press, 253-348.
- Roth, A.E. (1995b). "Introduction to Experimental Economics." In John H. Kagel and Alvin E. Roth, eds., *Handbook of Experimental Economics*. Princeton: Princeton University Press, 3-109.
- Roth, A. E. and A. Ockenfels (2002). "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet." *American Economic Review* 92(4). 1093-1103.
- Roth, A.E.; V. Prasnikar, M. Okuno-Fujiwara, and S. Zamir (1991). "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo." *American Economic Review*, 81(5), 1068–95.
- Roth, A.E., J.K. Murnighan, and F. Schoumaker (1988). "The Deadline Effect in Bargaining: Some Experimental Evidence." *American Economic Review*, 78(4), 806-23.
- Slonim, R., and A.E. Roth (1998). "Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic." *Econometrica* (66)3, 569-596.
- Smith V.L. (1962). "An experimental study of competitive market behavior." *Journal of Political Economy* 70, 111-137.
- Smith, V. (1982). "Microeconomic Systems as an Experimental Science." *American Economic Review*, 72 (5), 923-55.
- Yang, I, H. Jeong, B. Kahng, and A.-L. Barabasi (2003). "Emerging Behavior in Electronic Bidding." *Physical Review E*, 6

Appendix

Instructions sent by email to subjects in BT [Translation from German]

(Instructions in CT are completely analogous.)

Subject: Invitation eBay-Experiment
To: (...)
From: eBay-Experiment <ebay@lab.uni-koeln.de>

Dear (...),

Thank you very much for your registration and your interest in our eBay Project. Today we offer you to take part in an eBay experiment. Please read this email carefully to the end. Then decide whether you want to participate or not.

In the experiment you will either be a seller or a potential buyer of a fictitious good (called "certificate") on eBay. Information on your role will be sent to you in another email at the beginning of the experiment.

Rules for sellers:

- The certificate has no intrinsic value. If the certificate is not sold, the seller will not receive any payoff. If the certificate is sold, we will transfer the final price to the seller's bank account.
- The seller can choose the selling format on eBay (auction, auction with Buy It Now Option, or fixed price (Buy it Now)). For further information on the selling options, go to: <http://pages.ebay.de/help/sell/ia/angebotsformate.html>
- All eBay fees including the Listing Fee, the Final Value Fee, the Buy It Now fee, and the Scheduled Listings Fee, will be paid by us, regardless of whether the certificate is sold. We do not refund any expenses for optional features that promote the good, however, such as subtitle, eBay Picture Service, or Listing Designer.
- The offer must start on Monday October 10th 2005 at 3pm (Scheduled Listings Option) and must not run longer than 10 days. The category and the item description will be determined by us.
- The seller can post her listing only once. There will be no second chance to sell the certificate if the first attempt fails.
- Otherwise, all eBay rules apply.
- After having posted the listing on eBay, the seller will inform us about the eBay item number. We will check the offer and verify its compliance with our requirements. Then we will send the web address of the listing to a randomly chosen buyer.
- Important: There will be only 1 (one) potential buyer for the certificate!

Rules for buyers:

- The value of the certificate to the buyer is €20. If the buyer chooses to buy we will transfer €20 minus the price the certificate sold for to the buyer's account. If the buyer does not buy, there will be no money transfer.
- All eBay rules apply.
- Important: If the payable price exceeds €20, the buyer pays more than the certificate is worth to him or her. In that case the buyer must pay the difference to us! (All eBay rules apply; in particular all bids are binding!)

Rules for all participants:

Our payment obligation does only apply to the sellers' listings that we verified and transferred to a randomly chosen buyer. Payments will be made only if you use the eBay user ID you submitted upon registration. If you sell or buy a different certificate or use a different eBay user ID, you will not receive any money from us!

To facilitate the procedure, we will transfer the final price plus all fees to the registered seller, and €20 minus the final price to the registered buyer, as soon as the experiment is completed. You do not have to pay any eBay fees or any fees for the money transfer.

If you are interested in taking part in the experiment and are available during the required time (Oct 10th to Oct. 20th 2005), please go to the following link and sign up; note that with your agreement to participate, you claim to have understood the above mentioned rules and conditions and you agree to them.

http://www.lab.uni-koeln.de/ebay/teilnahmefragebogen.php?id=ux_

Before Friday October 7th 2005 we will inform you whether you are a seller or a buyer in our experiment.

Please do not hesitate to contact us at any time if you have any question: ebay@lab.uni-koeln.de

Thank you!

This email was automatically sent to you by the experiment participant system

If you want to update your information or sign off, please send us an email to the following address: ebay@lab.uni-koeln.de

Cologne Laboratory for Economic Research

<http://www.lab.uni-koeln.de>

Subject: eBay-Beginning of Experiment
To: (...)
From: eBay-Experiment <ebay@lab.uni-koeln.de>

Dear (...),

Thank you very much for your willingness to participate in our research project. You are a BUYER. We will randomly select a seller for you and will inform you when the offer is posted. This will not be before Monday October 10th 2005 at 3pm.

Best regards

Your Experiment Team

Please do not hesitate to contact us at any time if you have any question: ebay@lab.uni-koeln.de

This email was automatically sent to you by ...

Subject: eBay-Beginning of Experiment
To: (...)
From: eBay-Experiment <ebay@lab.uni-koeln.de>

Dear (...),

Thank you very much for your willingness to participate in our research project. You are a SELLER. Please create an offer taking into account the rules we outlined to you in our previous email. The listing must be posted in the category:
Collectibles > Technology & Tools > Science & Medicine > Miscellaneous

The offer must start on October 10th 2005 at 3pm (Scheduled Listing). The offer must end at the latest on October 20th 2005 at 3pm. Apart from that, the duration of the listing can be chosen freely.
The text in the offer must be:

Item title: "Certificate ..."

Item description: "Certificate ..."

This offer is part of a research project. The certificate is only of value to you, if you are a chosen participant and if you were informed about the item's ID number by the project manager. Please do not address any questions to the seller."

If there are any questions, you can forward them to us. We will answer them.

Delivery: "No Shipping – item needs to be picked up"

Shipping: You must not ask for shipping costs.

Please send us the eBay item ID number before October 10th 2005 at 12pm.

Best Regards,
Your Experiment Team

Please do not hesitate to contact us at any time if you have any question: ebay@lab.uni-koeln.de

This email was automatically sent to you by ...

Subject: Your eBay-offer ... has been forwarded
To: (...)
From: eBay-Experiment <ebay@lab.uni-koeln.de>

Dear (...)

Thank you very much for your email.

Your offer on eBay follows our specifications. The item ID number has just been forwarded to a randomly chosen buyer

Best Regards,
Your Experiment Team

Please do not hesitate to contact us at any time if you have any question: ebay@lab.uni-koeln.de

This email was automatically sent to you by ...

Subject: Your eBay offer ...
To: (...)
From: eBay-Experiment <ebay@lab.uni-koeln.de>

Please confirm that you received this e-mail by sending a response e-mail to this mailing address. Thank you.

Hello (...),

Thank you very much for your willingness to participate in our research project.

Your auction has been posted here: ...

The corresponding eBay auction ID is: ...

The eBay offer started Oct 10th 2005 at 3pm and ends in 10 days

You are bidding for the certificate ...

We would like to remind you of the following:

If you win the certificate at a final price higher than your value of €20, you have to pay the difference to us. If you do not place a successful bid, you will not receive any payoff. (All eBay rules apply without exceptions; all bids are binding!)

Thank you!
Your Experiment Team

Please do not hesitate to contact us at any time if you have any question: ebay@lab.uni-koeln.de

This email was automatically sent to you by ...

Subject: End of eBay experiment: ...
To: (...)
From: eBay-Experiment <ebay@lab.uni-koeln.de>

Dear (...)

Thank you very much for your participation in our eBay experiment

If you sold the certificate, you will receive the final price plus all eBay fees. If there was no deal, you will only receive the fees.

To simplify the payment process WE will do all money transfers, and YOU do not transfer any money. Please be patient. For accounting reasons, we can only transfer money, when we have received bank account information from all participants.

There are two more things we would like you to do:

1. Please go the following link and answer some questions regarding the experiment

http://www.lab.uni-koeln.de/ebay/verkaeuferfragebogen_mx.php?id=ux_

2. Please forward the following emails to us that you received from eBay: "Sold eBay item" email or "eBay item not sold" email. If you have already done so, you do not need to do it again. We need the information to complete the experiment's documentation.

We hope you enjoyed our experiment.

Best regards

Your Experiment Team

PS: eBay rules apply to the evaluation process also: participants can evaluate each other (only in the case of successful trades)

Please do not hesitate to contact us at any time if you have any question: ebay@lab.uni-koeln.de

This email was automatically sent to you by

Subject: End of eBay experiment: ...
To: (...)
From: eBay-Experiment <ebay@lab.uni-koeln.de>

Dear (...),

Thank you very much for your participation in our eBay experiment

If you bought the certificate, you will receive the difference between the final price and your value of €20. If there was no deal, you will not receive any money.

To simplify the payment process WE will do all money transfers, and YOU do not transfer any money. Please be patient. For accounting reasons, we can only transfer money, when we have received bank account information from all participants.

There are two more things we would like you to do:

1. Please go the following link and answer some questions regarding the experiment

http://www.lab.uni-koeln.de/ebay/kaeuferfragebogen_mx.php?id=ux_

2. If you bought a certificate please forward the following email to us that you received from eBay: "Bought eBay item" email. If you have already done so, you do not need to do it again. We need the information to complete the experiment's documentation.

We hope you enjoyed our experiment.

Best regards

Your Experiment Team

PS: eBay rules apply to the evaluation process also: participants can evaluate each other (only in the case of successful trades)

Please do not hesitate to contact us at any time if you have any question: ebay@lab.uni-koeln.de

This email was automatically sent to you by ...

Questionnaire before the start of the experiment

Before registering at our experiment, please read carefully your email and answer the following questions.

1. How many potential buyers are there?
2. What is the starting time of the offer?
3. What is the value of the certificate to the buyer?

Example of screen shot:



[Main](#)

[Anmelden](#)

[Informationen](#)

[Datenschutz](#)

[Impressum](#)

Teilnahmebestätigung

Bevor Sie sich verbindlich zu unserem Experiment anmelden, bitten wir Sie noch ein paar Fragen zu beantworten.
Bitte lesen vor Beantwortung der Fragen Ihre eMail noch einmal sorgfältig durch.

1. Wie viele potenzielle Käufer gibt es?

2. Wann soll das Angebot beginnen?

Am . . 2005 um : Uhr

3. Welchen Wert hat das Zertifikat für den Käufer?

Euro

Für Fragen wenden Sie sich bitte an ebay@lab.uni-koeln.de.