Does Laboratory Trading Mirror Behavior in Real World Markets?

Fair Bargaining and Competitive Bidding on EBay

GARY E. BOLTON
AXEL OCKENFELS

CESIFO WORKING PAPER NO. 2241
CATEGORY 10: EMPIRICAL AND THEORETICAL METHODS
MARCH 2008

CESifo *Working Paper No. 2241*

Does Laboratory Trading Mirror Behavior in Real World Markets?
Fair Bargaining and Competitive Bidding on EBay

Abstract

We conducted a controlled field experiment on eBay and examined to what extent both social and competitive laboratory behavior is robust to institutionally complex real world markets with experienced traders, who selected themselves into these markets. EBay's natural trading system provides bridges between lab and field environment that can be exploited to explore differences in behavior in the two environments. We find that many sellers do not make use of their commitment power as predicted by standard theories of both selfish and social behavior. However, a concern for equity strongly affects outcomes and reputation building in bilateral bargaining, while buyer competition effectively masks this concern and robustly yields equilibrium outcomes. The dichotomy of behaviors mirrors observations in laboratory research. Furthermore, we find that behavioral patterns in the field experiment mirror fully naturally occurring trading patterns in the market.

JEL Code: C93, D44.

Keywords: eBay, auctions, behavioral economics, trust, market design.

Gary E. Bolton
Smeal College of Business
Penn State University
University Park
USA - PA 16802
gbolton@psu.edu

Axel Ockenfels
University of Cologne
Department of Economics
Albertus Magnus Platz
Germany - 50923 Cologne
ockenfels@uni-koeln.de

16 August 2007
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 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; this study is partly motivated by but independent from our collaboration with eBay.

I. Introduction

Laboratory experimentation in economics has produced a large and rapidly expanding body of research. One of the important unresolved questions is how to validate lab findings out in the field. What makes this question particularly challenging is that the lab's most singular findings – those where tight laboratory control is most important to clear observation – are precisely those that are most difficult to pin down under field conditions.

This paper reports an experiment performed on eBay to examine whether well known laboratory findings concerning fair bilateral bargaining and competitive bidding are robust to a natural trading system. The two field games we examine are more elaborate than their laboratory cousins. Yet as with the laboratory versions, the two field games have similar equilibrium paths and outcomes, assuming selfish motives and rational behavior. In the laboratory versions, however, observed behavior in one game is strongly influenced by equity considerations whereas behavior in the other game follows the selfish equilibrium closely. The laboratory has been instrumental in understanding why behavior differs in games like this because it affords the control necessary to separate preferences from factors such as strategic opportunity or cognition.

The main contributions of our study are two. First, the results of the experiment speak to a current controversy concerning the extent to which behavior under lab conditions mirrors behavior under field conditions. For instance, with regard to social preferences, List (2006, p. 32) conducted lab and field versions of the gift exchange game and concluded that, "The finding that agents behave differently in tightly controlled laboratory experiments than in their naturally occurring environment poses an important challenge to laboratory studies that measure individual propensities." Second, we develop and implement a framework for the systematic vetting of laboratory results in field environments. Experimental economics has a long tradition of challenging the robustness of previous findings *within* the lab setting, including tests of artefactual features.[1] The framework we propose extends the principles underlying this tradition, in somewhat modified form, to the vetting of results *outside* the lab.

Towards a dialogue on how lab and field ecologies differ, Levitt and List (2007) put forward a baseline model. It posits that, along with monetary considerations, human decisions are influenced by (1) the presence of moral and ethical considerations; 2) the nature and extent of

[1] Much of the first *Handbook of Experimental Economics* (Kagel and Røth, eds., 1995) was devoted to the documentation of these kinds of experiments. The chapter by Roth on bargaining, games, documents the robustness of the laboratory bargaining game results referred to in this paper with respect to stake size, subject pool, and experimenter observation effects, among many other factors.

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 argue 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). [2]

Our experiment captures the natural circumstances of eBay's marketplace along all the stated dimensions: 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 management. EBay permits a trader to post "feedback " on his experience with another trader, open information for 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. (We use this information to link behavior in the field experiment to behavior in the fully naturally occurring environment.) Behind the scenes, eBay management monitors the site for unethical behavior, and this too plays a role in our experiment (see below).

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. Institutions, such as trading systems, are typically stripped down to match the parsimony of the theory being tested, with the goal of minimizing confounding considerations not accounted for by the model. In contrast, real world trading systems endure because they facilitate trade for profit. They are agglomerations of formal and informal rules, norms and heuristics, which 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.

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

[2] An earlier version of this paper "What do Laboratory Experiments Tell Us About the Real World?" was debated by Colin Camerer, Ernst Fehr, John List, Charlie Plott, and Alvin Roth in an AEA session 2006 in Boston, chaired by Gary Becker.

specify key terms of the deal. But the eBay platform offers the sellers who set up the trade mechanism more options (ex., auction, posted price, etc.) than are typically studied in a laboratory, with subsidiary rules serving practical purposes that must be abided by. EBay is also a far noisier trading environment than the lab. For instance, bidders and sellers may contact each other before the offer terminates, 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 important point is that, in checking laboratory results in the field, we need design the experiment to account for the added complexity. The design of the present experiment enables us to classify the behavior we observe as either in-selfish-equilibrium or in-social-equilibrium or as out-of-equilibrium. One of the findings will be that, contrary to what is often presumed, the level of out-of-equilibrium behavior among the self selected traders in the field environment is higher than commonly observed in the parallel lab situations with naïve subjects. Much of this difference is attributable to the miscalculations, noise and even malfeasance present in the field environment. We will see that, in the face of this complexity, the robustness of equilibrium predictions differs substantially across competitive bidding and bilateral bargaining conditions.

The design of our experiment is guided by three principles. Each speaks to some facet of the complexity issue. Together, they provide the basis for a dialogue on the robustness of lab results to field conditions. The first two principles extend principles that have long guided the vetting the robustness of experiments under fully laboratory conditions:

Losing control in a controlled way. Given differences in complexity, it is not in-and-of-itself surprising if behavior in lab and field differ. The issue is *why* they differ. The null hypothesis of our test is that they differ because the equity considerations in the lab versions of the experiment fail to the field environment, as suggested by Levitt and List. A natural competing hypothesis, offered by social preference theory, implies that competition in markets can mask social preferences, much as cooperation in repeated games can mask selfish preferences. By this hypothesis preferences are robust to the field, but may *appear* to be attenuated in certain strategic circumstances (see section II.4).

EBay's trading system defines a natural game form that lends itself to separating these two hypotheses, so long as we establish the necessary information structure and opportunity costs,

which we do by manipulating the object traded (without bending the natural rules of trade). As we will demonstrate, this loss of control in a controlled way enables a three-way link between highly controlled laboratory, partly controlled field and entirely uncontrolled 'real world' trading patterns. The results also provide a base for further, systematic study, say of the marginal effect of a further loss of control. Which leads to the second principle:

A replicable test bed. No single experiment is likely to decide the robustness of social preferences in the field. Field environments that are accessible to many investigators facilitate the research dialogue. It took many experiments, performed by many different investigators, to vet the robustness of social preferences in bilateral bargaining under fully laboratory conditions (see Bolton 1999 for a review).[3] A major reason for this is that the design and data from a single experiment almost always leads to new hypotheses about the behavior observed. Easily accessible field test beds would permit the method of replicating existing results and testing new hypotheses that has proven so productive in the lab, to go forward in the field. Online markets, being widely accessible, are one such set of test beds. 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. [4]

The third principle is *complementarity*. Given the differences in complexity and the proximity to natural phenomena, laboratory and field methodologies have complementary strengths. The problem is how to jointly leverage these strengths. Our solution is an experiment designed as a bridge between polls of lab, field, and fully natural environments. The design picks up on Harrison and List's (2004) taxonomy of field experiments, classifying them depending on their distance between the extreme polls of laboratory and natural environments (see section IV for a detailed discussion of related field experiments[5]). Rather than leaping from shore-to-shore, a bridge provides a more continuous, smoother transition, and one that can be travelled in both directions.

At one end of the bridge are the laboratory results on bilateral bargaining. Because the natural game form embedded in eBay's trading system echoes these lab games, we can exploit the lab findings in a number of ways. First, the laboratory tells us what to look for in the field.

[3] For reviews of the many experiments in the laboratory bargaining literature see, ex., Roth (1995) and Camerer (2003). For a sample of, and references to, additional work see Andreoni (2003).
[4] 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 recent survey on economic eBay research see Ockenfels et al. (2006).
[5] Also see John List's bibliography of field experiments at http://www.fieldexperiments.com/. Kagel (1995) provides an introduction into laboratory auction studies.

Equity in decision making, as observed in the lab, does not look like equity in its philosophical incarnations (ex., Güth et al. 1982, Ochs and Roth 1989). Social propensities tend to be asymmetric – people tend to care more about being treated fairly than treating others fairly – and many people are willing to trade-off at least some equity for pecuniary gains (so equity tends to be treated as a good, not as an inviolable principle).[6] The lab experiments also delineate the limitation of equity's influence. Maybe the most important of these results is that it seems to vanish from competitive markets, even when the game form is quite similar to ultimatum bargaining (ex., Roth et al. 1991). These observations have given birth to social preference theory that – together with the laboratory findings – provides guidance for our basic approach: We examine, under parallel field conditions, bilateral bargaining (where theory is divided on social preferences) and competitive bidding (where neither theory nor experiment suggests we should see any). This way, we gain better control for effects of complexity than we would if we compared field bilateral bargaining solely back to a laboratory treatment. Second, laboratory studies facilitate the interpretation of field behavior because the lab results tell us that fair bilateral bargaining results are robust to a host of factors within the lab (see footnote 1). For instance, experimental studies show that fair bilateral bargaining is not an artefact emerging from transaction costs in the trading environment, because transaction costs are typically controlled away in the laboratory. This suggests that analogous bargaining patterns in the field are unlikely to be fully explained by transaction costs, even though in the field transaction costs are difficult to be controlled away. That is, to the extent we can relate behavior in our field experiment to robust laboratory findings, we can be rather confident that the field behavior is part of a more general trading inclination (and not just the result of uncontrolled, peculiar factors in the field environment), robust to both laboratory and natural trading systems (see section IV).

At the other end of the bridge, eBay permits us to check for relationships between behavior inside the experiment and fully naturally occurring behavior outside the experiment. This is possible because eBay provides detailed individual trading histories online. Furthermore, 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

[6] One of the key techniques used to separate social propensities from other considerations has been close manipulation of the options and the information facing bargainers. These studies show how equity manifests itself in predictable ways with regard to both strategy (ex., Forsythe et al. 1994) and information conditions (ex,. Mitzkewitz and Nagel 1993, Kagel et al. 1996). While our field experiment does not address the effect of information, it speaks to the effect of strategic options on trading behavior.

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 hypotheses suggested by others concerning fully natural eBay trading.

II. Design of the experiment and hypotheses

We test the robustness of fair bilateral bargaining and competitive bidding using self-selected traders, trading in their marketplace by the marketplace's rules. The entire experiment was conducted on eBay's German market platform. 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.[7] 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 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 buyers; see Appendix). All game payoffs were wired. 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. [8] 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

[7] 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.

[8] 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).

least one of the feedback points in the role of a seller, guaranteeing some minimum role-specific trading experience.[9]

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. This way, our design minimizes repeated game and social interaction effects, which is important when interpreting motivations behind behavior.[10] 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, which control the opportunity costs associated with a trade. Induced values were chosen such that the equilibrium prices in the experiment roughly match 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.[11]

We also made buyer and seller valuations and the degree of competition (that is, the number of buyers with positive valuations for a given certificate) a matter of public knowledge. This removed any ambiguity in theoretical predictions originating from uncontrolled risk preferences. At the same time, it aligns 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

[9] By the rules of eBay Germany, a seller can only choose the BIN 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/fixedprice-faq.html).

[10] This is not to imply that we controlled all possible communication between subjects. Losing some control is inevitable when moving from lab to field environment. Still, eBay allows us to minimize some of the loss.

[11] See http://investor.ebay.com/downloads/fund_Metrics.pdf.

number of bidders likely vary widely, given the wide array of goods and trader sub-populations on the platform. For instance, 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. Our experimental conditions capture simple scenarios, differing by number of bidders (one or multiple).

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.[12] 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.[13] 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

[12] 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.

[13] 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.

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, both in order to suppress tacit and explicit communication between traders (known to have an effect in lab studies).[14] Finally, sellers were not allowed to use any marketing options, such as bold font, highlight etc., or to ask for shipping or handling costs.

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 many 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 – they were simply told to be free to take any action as long as it is in line with eBay's trading rules. 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 €. The game resembles various forms of laboratory price competition games,

[14] 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." The last sentence was included because an earlier, separate study indicated that sellers were frequently asked questions by non-participants about the nature and goals of the research. It may also have contributed to minimize buyer-seller interaction after the offer went online.

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 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 best offer on eBay automatically establishes a binding contract.

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 subject populations. All offers started on Monday, 11/14/2005 at 3pm. Everything else was as in the bilateral trading environment.

II.4 Theory and hypotheses

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

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 all potential buyers, because we chose an eBay category that does not allow a hidden reserve price in auctions. 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 known to all traders. Thus, in perfect equilibrium assuming standard pecuniary preferences, the seller is able to claim the entire surplus in either game *regardless of the degree of competition*. In both treatments, BT and CT, the seller chooses either a BIN-offer of 20 €, or an auction offer with a start price of 20 €. In case a seller chooses

the hybrid offer-format, both the auction start price *and* the BIN-offer must be 20 €.[15] In perfect equilibrium, all offers must be accepted by (at least one) buyer.

One might think that bidding in CT will always end at prices equal (or close) to 20 €, independent of the start price, so that the competitive equilibrium outcome we characterize for CT is robust against the auction start price. In fact, there is also an equilibrium in CT with a low start price and at least two bidders bidding their full values. However, because of eBay's dynamics, a low start price would allow bidders to coordinate on low-price equilibria. To see why, observe first that in Ockenfels and Roth's (2006) model of eBay's private-value second-price auction environment, which exactly matches our CT condition, all bidders will 'sooner or later' bid their values in any equilibrium in undominated strategies.[16] In particular, there are 'late bidding equilibria', in which all buyers bid values in the closing seconds of the auction.[17] Because late bids run the risk of not being successfully transmitted to eBay, however, 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.[18] That is, even when all bidders bid value, if they delay bidding until the bid risks not being accepted, revenue may be small. Furthermore, there are also 'early bidding equilibria', in which the first bidder who sees the auction bids his value, 20 € (or more), and then no one else bothers to bid because later bids can only win at a price above value (related equilibria also exist in Vickrey's sealed-bid second-price auction). While not bidding is a weakly dominated strategy (Ockenfels and Roth 2006), the early bidding equilibrium appears a plausible alternative in our setting. For one, eBay's dynamics allow a 'natural' way of selecting a winner based on the timing of bids. Second, because it is publicly known that all values are identical, all traders know that competitive strategies imply zero payoffs.

[15] 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.

[16] A strategy that calls for bidding above value is dominated.

[17] 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.

[18] In the model of Ockenfels and Roth, the equilibrium is stabilized by the threat that 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. See Ockenfels and Roth (2006) for the details.

All this said, in the face of uncertainty about the bidders' willingness and ability to coordinate on low price equilibria, sellers can avoid low revenues by choosing a start price equal to the bidders' valuation.

The foregoing discussion gives rise to our *null hypothesis*: Independent of the treatment, sellers are indifferent between the BIN-format, the auction format and the hybrid format. However, even when sellers choose the auction format in CT, they do not rely on buyer competition: All start- and BIN-prices are equal to the buyer valuations (20 €). All buyers are willing to accept the proposed price.

We next turn to examine behavior under the assumption that traders are willing to trade-off some of their pecuniary payoff for increased relative payoff (they have social preferences). As a result, inequitable offers in the BT treatment will not be accepted by buyers who prefer the equitable allocation that results if the offer is not accepted. For instance, in the models by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000), an offer of 20 € is never accepted because rejection does not reduce pecuniary payoffs but strictly increases relative payoffs.[19] Both models postulate that there is a great deal of heterogeneity with respect to how players trade-off pecuniary and relative payoffs. While equal splits are always accepted, more inequitable offers are associated with a higher probability of not being accepted. As a consequence, in perfect Bayesian equilibrium, the BT-sellers' average price offer (or start price) is smaller than what is predicted by the null hypothesis, and because of incomplete information about buyer preferences, there is a positive probability that 'unfair' offers are rejected. (The formal proof is similar to the proof of Statement 2 in Bolton and Ockenfels 2000.)

For the analysis of the impact of social preferences in CT, let us start by assuming that the seller wants to maximize *pecuniary* revenues (the asymmetric social preferences in Bolton's 1991 model yields equivalent results without additional assumptions on seller motives). Due to the heterogeneity of bidder preferences, the probability that at least one out of nine bidders is willing to accept a given (unequal) price offer is higher than the probability that a single bidder accepts a given offer. Thus, competition in CT allows the seller to make more aggressive fixed price offers than in BT. Nevertheless, the equilibrium price offer of 20 € under the null hypothesis is not a reasonable perfect Bayesian equilibrium offer with social preferences. The reason is that

[19] We take these models as examples, because since they are outcome-oriented, they tend to be simpler than models of reciprocity that are based psychological game theory, such as Rabin (1993) and Dufwenberg and Kirchsteiger (2004). However, in our context, most these models seem to yield qualitatively similar predictions.

acceptance of the offer by bidding 20 € is weakly dominated by not bidding at all (rejection): If nobody accepts the 20 €, 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.[20]

Yet, by running an auction instead of offering a fixed price, a CT-seller can, in social equilibrium, exploit buyer competition to push the price up to 20 €. The trick is to choose an auction start price that, if it were the final price, would be acceptable to bidders, independent of their strength of concern for relative payoffs. In the model of Bolton and Ockenfels (2000), for instance, a price of 18 € would yield a payoff of 2 € to the buyer, which is equitable and thus acceptable to all bidders.[21] But once the offer is accepted, the bidders' behavior is identical to the behavior of selfish traders. The reason is that a buyer is better off – both in terms of relative and pecuniary payoffs – competing and winning than not competing and allowing another buyer to make the trade. That is, the pursuit of absolute and relative payoffs is perfectly aligned so that competition makes buyers to behave as if they are selfish.[22] As a consequence, in the auction case with an 18 € reservation price, all buyers choosing undominated strategies will sooner or later bid values: If nobody else bids, a bid of 20 € yields a fair payoff of 2 €. If other bidders submit bids, a bid of 20 € makes sure that the bidder can get at least some share of the cake whenever possible; that is, whenever other bidders' bids are sufficiently low. (The formal proof is similar to the proof of Statement 9b in Bolton and Ockenfels 2000 and to the models in Ockenfels and Roth 2006.)

So, social preference models imply a rationale for why sellers may want to choose an auction with a start price below valuations, even when valuations are common knowledge. While a start price below valuations risks low revenues through early or late bidding strategies,[23]

---

[20] Bidding above 20 € is also a weakly dominated strategy.

[21] An offer of 2 € is equitable 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).

[22] 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.

[23] For the early-bidding equilibrium with social preferences, for instance, 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 – as it is with standard selfish preferences.

demanding the whole cake is (unlike in the scenario with selfish preferences) no resort 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.

If, on the other hand, the *seller* is fair-minded, he may *want* to trade at a price well below 20 €, so that sellers and buyers get on average more equitable payoffs. The exact offer then depends on the trade-off between equity and selfish concerns of the seller, and on the nature of these equity concerns in favourable relative situations. The trade-off and the nature of the concerns differ across models.[24] We will come back to this issue in the data section. What is suggested by our framework, however, is that if sellers are willing to sacrifice revenue out of equity concerns, they should not do so via the auction format, because they could then lose control over their share of the efficiency gain from trade (which depends on bidders' willingness and ability to choose competitive or collusive bidding strategies).

This discussion gives rise to our *alternative hypothesis*: In BT, the sellers' average price offer is smaller than what is predicted by the null hypothesis, and greedier price offers run a higher risk of being rejected. In CT, selfish sellers can exploit buyer competition in an auction to push up the price; too greedy BIN-offers may get rejected, but CT-auctions with more than one active bidder reach a final price of 20 €. The auction start price, then, can be below 20 €, but (in the Bolton and Ockenfels model) not below 18 €. Fair-minded sellers, on the other hand, may choose price-offers below 18 €, but should then prefer the BIN-format.

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 effect of social concerns on reputation building via eBay's feedback forum, after a successful trade.

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

[24] 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.

(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 ($p < 0.001$). While this pattern of format choice is not suggested by standard theory, it is suggested by social preference theory. In social preference theory, there is no reason to prefer one or the other format in BT, but buyer competition in CT-auctions can avoid problems associated with buyer resistance to accept high price offers (see section II.4).

The format choice is also affected by a seller's trading experience that evolved in the field, outside of the experiment. 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 was 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 experiences gained on eBay in naturally occurring encounters.

Figure 1. Tabulation of sellers' initial price offers

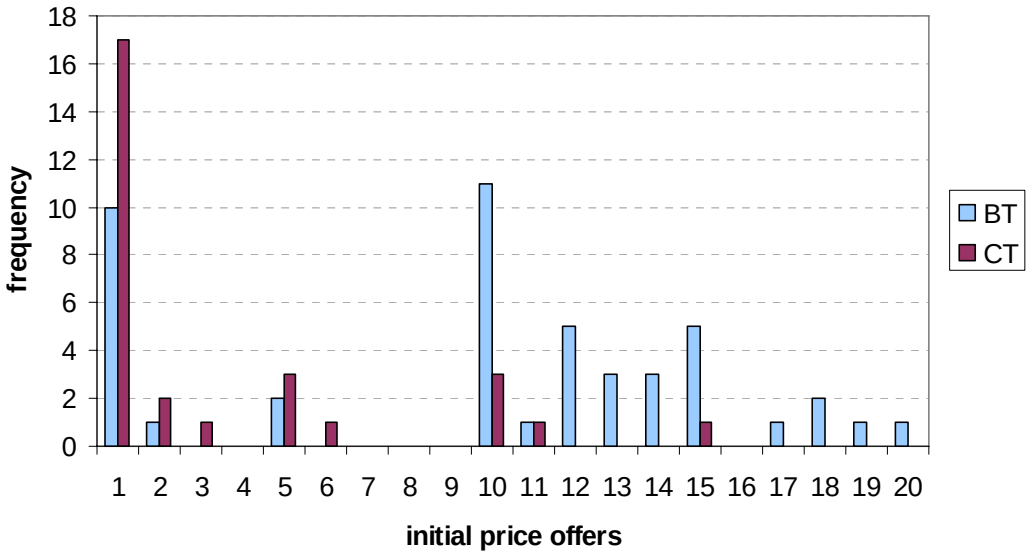

Figure 1 shows the initial price offers posted by sellers in bilateral and competitive trading.[25] For the instances in which the seller chooses the hybrid format involving an auction start price and a fixed price, we define the price offer as the minimum of the two, which by eBay rules, must be the auction start price. Overall, only a single price offer is consistent with selfish equilibrium (19.99 or 20 €). The modal offer is 10 €, the equal split, in BT and 1 €, the default start price on eBay, in CT.[26] Thus, price offers are inconsistent with our null hypothesis.

Social equilibrium is in principle consistent with the offers strictly below 20 € in both BT and CT. Nevertheless, the theory is inconsistent with a non-negligible number of price offers (13, 28 percent) below the equal split in BT.[27] These offers are out-of-equilibrium regardless of the model. We made every effort to make sure sellers understood there would be only one potential buyer (see questionnaire in Appendix). For this reason it appears that at least some of these sellers did not know from their previous experience on eBay how to handle the lack of competition in BT, perhaps not knowing how to handle the full array of tools available to them. This view is supported by a closer look at these traders' background as reflected in our history data.

[25] 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."
[26] Price offers in BT are higher than in CT (Mann-Whitney *U*, two tailed, $p < 0.001$).
[27] Strictly speaking, 14 offers are below the equal split, but we treat a 9.99 € offer here and elsewhere in the paper as equitable.

The history data suggest that 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 in our experiment generally placed smaller start prices on eBay than the other sellers who chose the auction format; the average start price in our history data is 1.92 €, while the other sellers placed a start price of 7.06 € (Mann-Whitney *U*, two-tailed, $p = .021$). Also, BT-sellers with offer prices below the equal split are somewhat less acquainted with the BIN format, which is the more 'natural' choice in bilateral bargaining, compared to other sellers; the average proportion of BIN-offers relative to all eBay offers in our history data is 4.5 percent compared to 25.8 percent of the other sellers (Mann-Whitney *U*, two-tailed, $p = .111$). Finally, we observe that the average feedback score of those BT-sellers with price offers below the equal split is 95, much smaller than the 224 of the other sellers, and the average number of auctions conducted by those BT-sellers who chose an auction start price below the equal split is 4.5, much smaller than the 10.2 of the other sellers, though these differences are not significant.

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, Roth et al. (1991) report data from a very large ultimatum game experiement done over four countries (Israel, Japan, Slovenia and the U.S.). Individual bargainers played the game 10 times, each time with a different partner. In the first two rounds of play, 9.3% of the observed offers were below 50% of the pie, but by rounds 9 and 10 none of the offers were below 50%. This first mover experience effect suggests that the initial offers below 50% were due to inexperience rather than preference. Plausibly, the rate of the these kinds of mistakes is higher in the field environment, probably because the seller's options in the field are more complex, effectively a superset of those facing a first mover in a simple lab ultimatum game.

The overall picture that emerges from combining the field experimental, the historical trading and the lab data is that the low price offers are at least partly due to a lack of experience. One interesting individual example along this line 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. 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 seller. And second, we found four non-

experimental auctions on eBay run by our seller, where the same bidder, who submitted the shill bids in our experiment, has submitted non-winning bids. This suggests that the seller employed a shilling strategy in our experiment to increase competition that was adapted from his eBay experience outside the experiment.[28]

Summing up, we find that with only one exception, sellers do not make use of their commitment power as predicted by selfish preferences for BT and CT. Furthermore, a non-negligible share of BT-sellers fails to make use of their commitment power as predicted by both selfish and social equilibrium. As expected by social preference theory, auctions are more frequently used in CT. However, for these auctions too, it appears that sellers do not make full use of their commitment power, in that the modal offer is a 1 € auction start price, chosen by 58 percent of the CT-sellers (see the discussion in the next section).

III.2 Buyer behavior in competitive trading

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

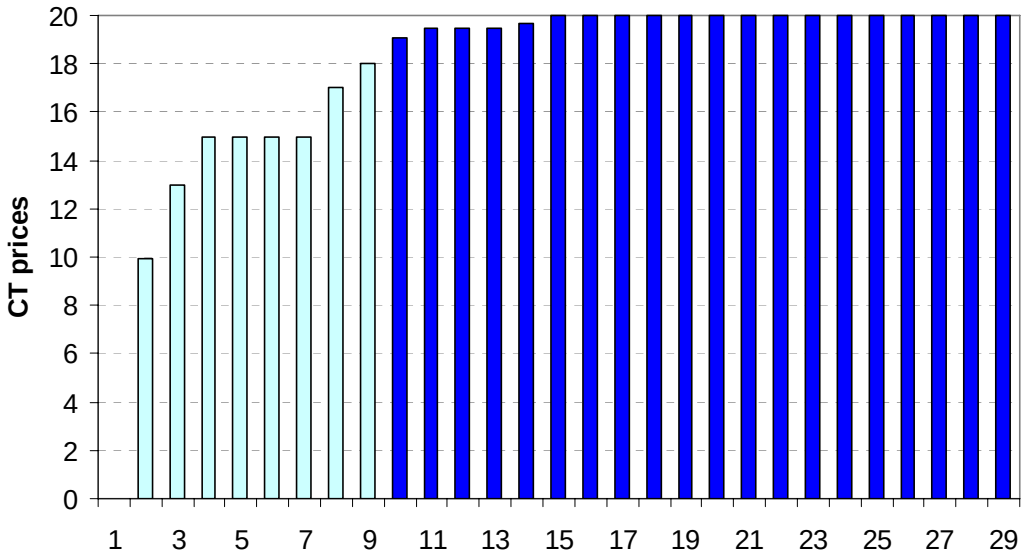

Figure 3 shows the final prices (seller profits) realized in the 29 competitive markets. Observe that whenever bidders competed in auctions (the 20 dark bars in Figure 3), the final price was above 19 €. In 14 of these auctions the price was as predicted by equilibrium theory (19.99 € or 20 €). In 5 auctions, prices were in a range between 19.50 € and 19.98 €. This price range

[28] 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.

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).

The success of auctions with respect to revenues and the large agreement with prices predicted by equilibrium theory – both in standard and in social preference models – is remarkable, because seller behavior is largely inconsistent with our hypotheses. In 15 out of all 20 auctions, the start price was 1 €, while our null hypothesis suggests 19.99 or 20 €, but it also speaks against social preference models. Social preference theory suggests that a seller may choose a low auction price that buyers are willing to accept in order to get competitive bidding going. Since an offer of 1 € is smaller than what is needed to attract bidding (because it would give the buyer more than his fair share), and since at the same time low start prices run the risk of low revenues via early or late bidding equilibria, sellers seem to give away some control over the price discovery process. In this sense, they do not make full use of their commitment power – regardless of whether preferences are selfish or social.[29] As it turns out, however, buyers were not able to take advantage of sellers' low price-offers: Virtually all buyer payoffs were competed away.

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 18 €. There are at least two possible reasons for 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,[30] 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

[29] 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 €.)

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

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 power and robustness of competition, because (s)he apparently saw a need to push competition.

III.3 Buyer behavior in bilateral trading

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

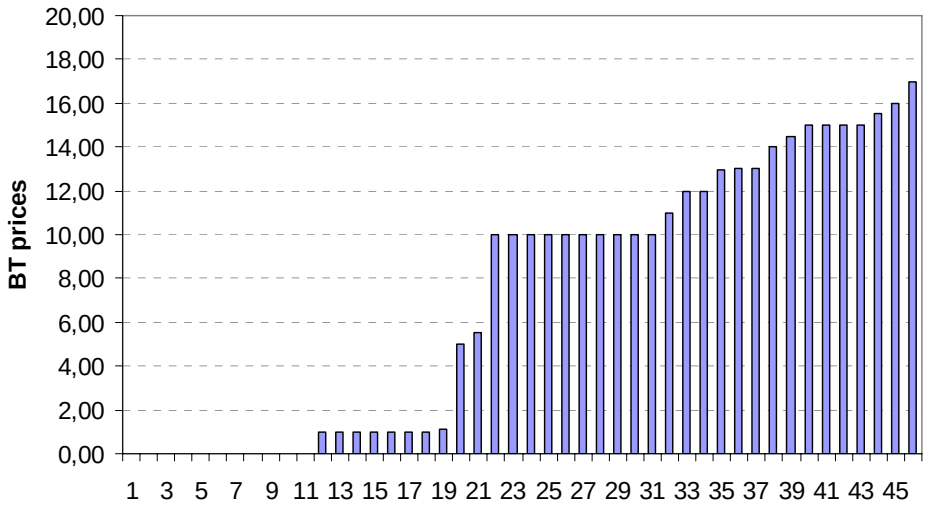

Figure 4 sorts the final prices (seller profits) in the 46 BT encounters. The average final price (seller payoff) is 7.17 €, compared to 17.76 € in CT. The hypothesis of equal revenues is strongly rejected (Mann Whitney *U*, two-tailed, $p < 0.001$). There is also a much wider variance in BT final prices. Eleven encounters ended in zero profits, the maximum seller profit was 17 €. None ended close to the 20 € seller profit expected by selfish equilibrium.

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

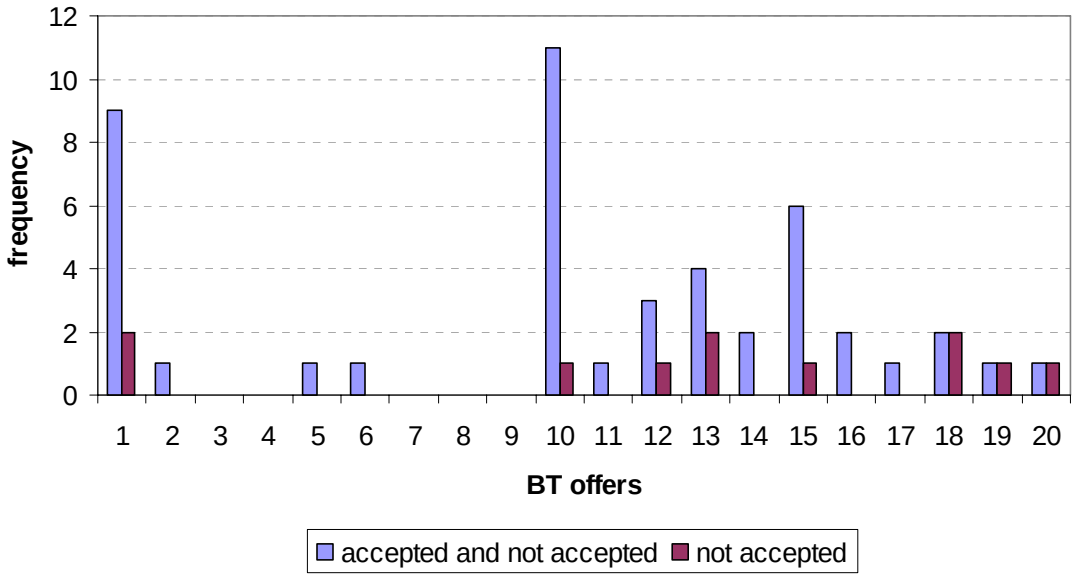

Figure 5 helps to understand the large variance in final price. It displays the pattern of offers accepted and not accepted in bilateral trading. An offer is defined as accepted when the buyer accepted the BIN price or submitted a winning bid in the auction, respectively. An offer is defined as not accepted 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. In some cases, it might be because a buyer was not available or because (s)he forgot to submit a bid, or a last-minute bid failed to be successfully submitted. Thus, after the experiment was completed, we asked all BT-buyers to briefly explain their strategy.

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. [31] While rejections of price offers below the equal split are sometimes observed also in laboratory research and are not necessarily

[31] 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.

inconsistent with social utility models,[32] 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, to the extent the buyers' explanation is accurate, 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 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, both qualitatively and quantitatively. 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 'greedy' offers starting at 18 € were not acceptable to buyers. There is a significant, negative 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).

In comparison, in the Roth et al. (1991) laboratory ultimatum bargaining game, 26.4% of the offers were rejected (there was no significant second mover experience effect), very close to what is observed here. This suggests that the behavior we observe in the bilateral bargaining treatment is consistent with the quantitative, as well as qualitative, influence of social preferences observed in the laboratory.[33]

The attentive 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

[32] Several other experimental studies including Bornstein and Yaniv (1998), Güth et al. (forthcomimg), Gehrig et al. (forthcoming), 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.

[33] Calibrations of social utility functions are presented in Fehr and Schmidt (1999) and DeBruyn and Bolton (2006).

was higher than the auction start price.[34] 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 buyers who take the higher BIN price is 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.[35] 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 buyer from these risks. 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, $p = 0.023$).[36]

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.[37] Leaving feedback on the trading partner can be an(other) 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 (2006), 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

[34] The average difference between BIN and auction start price in these 5 cases is 3.30 €.
[35] 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.
[36] It is sometimes argued that buyers prefer BIN-offers when they are impatient. 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.
[37] See Resnick and Zeckhauser (2002) for an early description and analysis of eBay's mechanism.

of receiving negative feedback (Bolton et al. 2007). 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.[38]

Overall, 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 downloaded 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 are left with 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;[39] 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" 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.[40]

One might at first glance suspect that a concern for a good reputation is sufficient to explain the social behavior we observe in BT. That is, the prospect of being rewarded with a good feedback may be a good incentive for sellers to make a fair offer. Yet such reputational concerns cannot explain the key phenomenon on the buyer side: Rejections in the BT-treatment do not only abolish the buyer's profit but also make it impossible for buyers to improve their own reputation and to spoil the seller's reputation.

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,

[38] 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.

[39] The neutral one occurred in the bargaining with shill bidding; the buyer left a comment "I suspect shill bidding. Otherwise ok." (translation from German).

[40] 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.

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 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 *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 $p = 0.035$). However, 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 sensible to the concerns of the trading partner.

We finally ask whether and how buyers' acceptance behavior is correlated with reputation scores. There are competing null hypotheses. Are buyers who do not accept an offer in our bilateral bargaining environment more or less experienced than accepting buyers? 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

Starting with a seminal paper in 1962, Vernon Smith showed in a series of classroom experiments that certain market institutions converge rapidly and robustly to the competitive equilibrium. In particular, it has been shown that the double auction mechanism works with very little learning, with incomplete information about other traders' (induced) costs and valuations, and with only a modest number of experienced traders (see also Plott 1982). Our paper complements this line of literature by showing that competitive behavior can be surprisingly robust to the noise and various complexities inherent to real world markets.

The first online auction experiment 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 put 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." In contrast, in our eBay experiment we took great care to gain control for the critical assumptions of the models, because in social preference models institutions matter greatly. Suppose, e.g., we had instead auctioned off Magic Cards on eBay to all interested traders, just as Lucking-Reiley did. Then, understanding prices would require estimating private and common valuation components of the Magic Cards, beliefs over the distribution of values, beliefs over the number of potential buyers, and risk preferences. Because there are obviously many degrees of freedom for standard bidding models, the residual contribution of social preferences to explaining prices can be made almost arbitrarily small. However, when testing the comparative 'practical' performance of auction formats, as Lucking-Reiley did, none of these issues seem terribly relevant, if only because they can reasonably be taken as constant across treatments.

List's (2006) paper bridges lab and field behavior in carefully controlled lab and field gift-exchange environments, also using real traders in a real (sports card) market. His data suggest that while the laboratory patterns appear to be rather stable against changes in subject pools, communication channels, framing and information scenarios, moving from lab markets to field markets significantly reduced the influence of social preferences. However, based on findings such as Falk's (forthcoming) study of charitable giving, List remarks that there may be other

domains in which gift exchange and social preferences might be significant. Our approach complements List's approach in several respects. Methodologically, we choose a design that can probably be interpreted as being somewhere between List's comparison of lab and field experiments; by gaining control over the degree of competitive pressure, opportunity costs and repeated interaction effects, we can study how social behavior in a natural trading environment responds to factors, which have been identified as important in laboratory research and social preference theory. Also, our study supports List's observation that reputational concerns may play an important role in market transactions in that they can trigger behavior that is consistent with social preferences. Our BT-traders use their influence on the trading partner's reputation to reward fair and punish relatively unfair behavior, which may have a disciplining effect on sellers. However, the reputation effects cannot explain the social behavior we observe on the buyer side (as explained in section III.5). So, our field experiment suggests that the impact of social preferences is neither restricted to off-market domains nor to laboratory contexts. One might speculate that this disagreement of the data has something to do with the methodological differences and the degree of control inserted in the different studies, with the particular markets studied, or with the fact that List studied positive reciprocity, while our work is mainly concerned with negative reciprocity (e.g., a buyer can reject if the seller's price-offer if perceived as unfair).[41]

Gneezy and List (2006) investigated social preferences in a field experiment of a labor market, which resembles a laboratory gift-exchange game environment. They found that a higher wage was reciprocated by more effort on the part of the employees during the early hours of the task – which is in line with what is typically found in the laboratory. After a few hours, however, effort levels in the gift treatment mirrored those in their non-gift treatment. As a consequence, the authors question the idea that experimental results gathered in the span of an hour or two can always be used to make inference on tasks that are inherently much longer lived. They concluded from their data: "We interpret our findings as suggesting that great care should be taken before making inference from laboratory experiments, which might be deemed as "hot" decision making, to field environments, which typically revolve around "cold" decision making."[42] Our study

[41] For the latter argument, observe that experimental results suggest 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).

[42] Interestingly, Lucking-Reiley's study also found that differences of duration of laboratory and field interaction may have important implications. He found that Dutch auctions earn approximately 30 percent more revenue than

complements Gneezy and List's work in at least two respects. First, we identify social behavior in an extremely adverse environment, where interaction is, unlike in labor markets, among anonymous traders with no common history or common future, and with relatively restricted communication channels and strategic options. Second, our study speaks to the effect of hot vs. cold decision making. In the jargon of Gneezy and List, the traders in our experiment make "cold" decisions. Sellers had about a week to think about their offer, and buyers typically had a week to think about their responses. Still, we observe strong evidence for social behavior in BT. It seems that cold decision making in the sense of having substantially more time to think about one's response to an unfair price offer does not necessarily diminish the tendency to reject.

There is also a related and fast-growing literature dealing with the economics of the eBay market (Bajari and Hortacsu 2004, and Ockenfels et al. 2006 provide 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 preferences on the seller side may also contribute to the success of eBay's BIN-format. A few other papers used the actual eBay platform for controlled experiments. Jin and Kato (2004, 2005) conducted a field experiment searching for winner's curse on eBay. They bid on eBay auctions for ungraded baseball cards and then let a professional grading service evaluate the cards.

Another related literature deals with eBay's online feedback system (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). The paper, which is probably closest to ours, in

---

first-price auctions, which was inconsistent with previous laboratory studies. He also observed that his descending-price auctions took much longer than laboratory experiments and speculated that the higher revenues are because bidders might have been impatient to complete their purchase (see Katok and Kwasnica 2005, who found laboratory evidence that clock speed matters and a simple model of impatient bidders).

methodology, is Resnick et al.'s (2006) field experiment on eBay, investigating the impact of (manipulated) differential feedback reputation scores on revenue. However, none of these online market studies investigated the impact of social preferences on bidding and trading behavior, and none were conducted to identify trading patterns that connect controlled laboratory and fully uncontrolled market behavior.

Many behaviors in our experiment are consistent with the social utility hypothesis. Yet consistency does not necessarily imply that our observations are *explained* by social preferences. One may argue that trading on eBay involves more complicated trade-offs than trading in the laboratory. For instance, our eBay bidders may face higher opportunity costs of time when bidding or higher costs associated with winning than bidders in the laboratory, where everybody has basically already committed to stay until the end of the experiment and gets paid in cash before leaving. Indeed, if buyers differ with respect to *transaction costs* associated with bidding and winning, 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.

It is, however, unlikely that transactions costs capture the essence of what we are seeing in our experiment.[43] For one, we did everything to minimize transaction costs. There is no shipping and handling involved, all payoffs were wired, and all pecuniary transactions costs (eBay and bank fees) were paid by the experimenter. Also, the modal offer is the equal split, which is easily explained by a concern for equity on the buyer and/or on the seller side but not by transaction costs arguments. And some 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 most important argument against transaction costs arguments, however, does not come from our field data but from lab experiments. Our bargaining patterns mirror those observed in laboratory ultimatum games. Because there are only negligible transaction costs in the laboratory, if at all, the laboratory results demonstrate that they are not necessary to produce the bargaining patterns observed on eBay – although transaction costs could, of course, reinforce the behavior. That is, laboratory and field experiments produce complementary evidence. If we had only laboratory

[43] There are, of course, other possible complexities. One hypothesis is that eBay bidders enjoy bidding on eBay more than bidding in the laboratory. But there is no clear evidence for this; no price has ever risen above 20 €, the induced pecuniary value. The same evidence speaks against a utility for winning.

results, one may be inclined not to believe that the behavior in simple, artificial laboratory games may generalize to behavior in more complex, naturally occurring markets. If we had only the field results, one may be inclined to believe that the explanation must lie in the institutional, social or motivational complexities that underlie most real world interactions. Laboratory and our field evidence together, however, make for a surprisingly coherent picture.

V. Conclusions

The data from the field experiment demonstrate that the social preference behavior observed under lab conditions, both in quality and in magnitude, extend to experienced traders operating in their natural trading system. More specifically, the data confirm the dichotomy between equitable bargaining and competitive bidding predicted by social preference equilibrium and suggested by laboratory evidence. Buyers in the bilateral trading environment were prepared to reject unfair price offers, and the strength of social preference we observe is quite similar in magnitude to what laboratory experiments report. 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. Also, supporting a hypothesis derived from social preference theory not previously tested in either environment, sellers prefer to choose the auction format in the competitive trading environment, while they tend to be indifferent between available formats in the bilateral trading environment.

One of the principle lessons the present work has for the vetting of laboratory results in the field is the need for experiments that lose control in a 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, noise and malfeasance, 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.

Maybe the most notable out-of-equilibrium phenomenon in our field study is that sellers often 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 in the field tends to strengthen the dichotomy observed in the laboratory. Because of some very low price offers in the bilateral trading environment, bargaining revenues are even

more dispersed than in the laboratory. Competitive trading, on the other hand, not only masks social preferences but is also robust against the increased strategic and cognitive complexities on eBay. There are miscalculations, noise and malfeasance behind otherwise competitive looking behavior.

We also find that the behavior in our controlled field experiment is intertwined with uncontrolled field behavior. A priori experience with eBay predicts and affects 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. 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.

An important part of social interaction takes place *after* the price discovery process. Once the price has been determined, the transaction boils down to a bilateral one – independent of the degree of buyer competition in the price discovery phase. In this contracting stage, the seller and buyer deal exclusively with one another concerning the non-price terms of the transaction (i.e, shipping, quality expectations, etc). On eBay, this phase ends with a reputation building stage. The feedback on the transaction partner given (or not given) on this stage will typically depend on overall trading behavior. While our experiment suppressed all non-price terms, it nevertheless demonstrates that this stage is used as an opportunity to bring social concerns into play. This way, preferences and market institutions interact to promote social behavior.

Our study is open to relatively easy replication, as well as extension in various directions. Many of these involve a changing of control or a 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, etc. With the emergence of online trading platforms, a highly controlled and flexible analysis of field behavior becomes 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.

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*, Erhlbaum: 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 (2007). "Designing EBay's New Feedback Mechanism: Lab and Field Evidence." Mimeo.

Bolton, G. E., E. Katok and A. Ockenfels (2004). "Trust among Internet Traders: A Behavioral Economics Approach." *Analyse und Kritik*, 26(1), 185-202.

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), pp. 166-193.

Bornstein, G., and I. Yaniv (1998). "Individual and group behavior in the ultimatum game: are groups more rational players?" Experimental Economics, 1, 101–8.

Camerer, C.F. (2003). "*Behavioral Game Theory: Experiments in Strategic Interaction.*" Princeton University Press: Princeton.

Cohen, A. (2002). "The Perfect Store: Inside Ebay." Little, Brown and Company.

DeBruyn, A., and G.E. Bolton (2006). "Estimating the Influence of Fairness on Bargaining Behavior." Working paper.

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. Woody (2006). "The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias." University of Maryland, Working Paper.

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.

Falk, A. (forthcoming). “Gift-exchange in the Field.” *Econometrica*.

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 (forthcoming). „Buying a pig in a poke: An experimental study of unconditional veto power.” *Journal of Economic Psychology*.

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 (forthcoming). “Bargaining outside the lab - a newspaper experiment of a three-person ultimatum game.” *Economic Journal*.

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 (2004). "Consumer Frauds and the Uninformed: Evidence from An Online Field Experiment." University of Maryland, Working Paper.

Jin, G. Z., Kato, A. (2005). "That’s News to Me! Information Revelation in Professional Certification Markets." University of Maryland, Working Paper.

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.

Katok, E. and A.M. Kwasnica (2005). “Time is Money: The effect of clock speed on seller’s revenue in Dutch auctions.” ISBM Working Paper 11-2002.

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.

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. (2006). “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.

Mitzkewitz, M., and R. Nagel (1993). "Envy, Greed and Anticipation in Ultimatum Games with Incomplete Information." *International Journal of Game Theory*, *22*(2), 171-198.

Namazi, A. and A. Schadschneider (2006). “Statistical properties of online auctions.” *International Journal of Physics, C17, 1485*.

Kagel, J.H., C. Kim and D. Moser (1996). “Fairness in Ultimatum Games with Asymmetric Information and Asymmetric Payoffs.” *Games and Economic Behavior*, 13, 100-110.

Klein, T.J., C. Lambertz, G. Spagnolo, K.O. Stahl (2007). “Reputation Building in Anonymous Marlets: Evidence from eBay.” Working Paper, University of Mannheim.

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.

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.

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. (1995). "Bargaining Experiments." In John H. Kagel and Alvin E. Roth, eds., *Handbook of Experimental Economics*. Princeton: Princeton University Press, 253-348.

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.

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 ca 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 to 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:

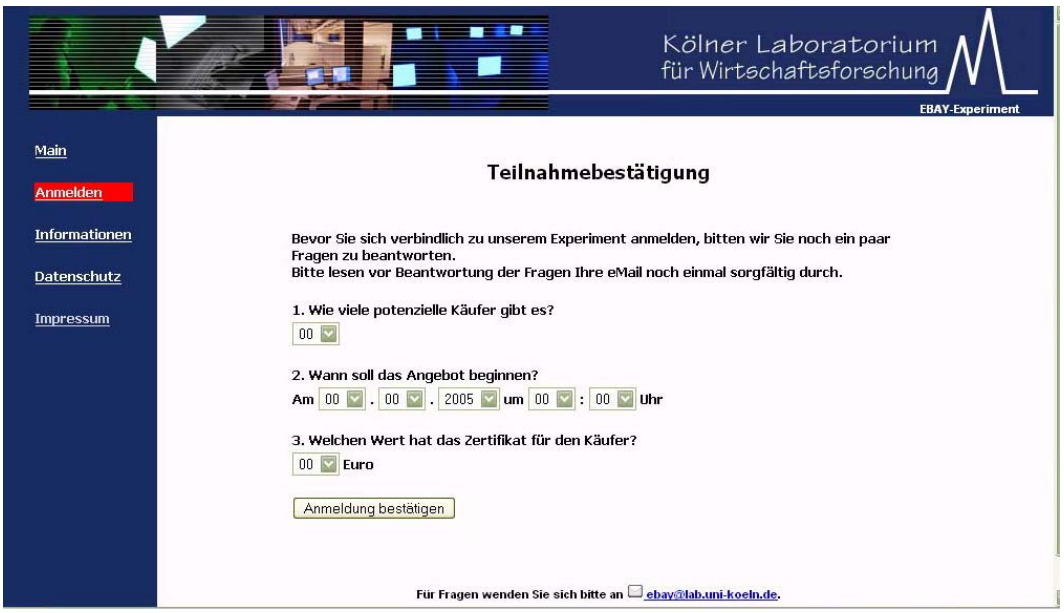

CESifo Working Paper Series

for full list see www.cesifo-group.org/wp
(address: Poschingerstr. 5, 81679 Munich, Germany, office@cesifo.de)

2176 Alexander Chudik and M. Hashem Pesaran, Infinite Dimensional VARs and Factor Models, December 2007

2177 Christoph Moser and Axel Dreher, Do Markets Care about Central Bank Governor Changes? Evidence from Emerging Markets, December 2007

2178 Alessandra Sgobbi and Carlo Carraro, A Stochastic Multiple Players Multi-Issues Bargaining Model for the Piave River Basin, December 2007

2179 Christa Hainz, Creditor Passivity: The Effects of Bank Competition and Institutions on the Strategic Use of Bankruptcy Filings, December 2007

2180 Emilia Del Bono, Andrea Weber and Rudolf Winter-Ebmer, Clash of Career and Family: Fertility Decisions after Job Displacement, January 2008

2181 Harald Badinger and Peter Egger, Intra- and Inter-Industry Productivity Spillovers in OECD Manufacturing: A Spatial Econometric Perspective, January 2008

2182 María del Carmen Boado-Penas, Salvador Valdés-Prieto and Carlos Vidal-Meliá, the Actuarial Balance Sheet for Pay-As-You-Go Finance: Solvency Indicators for Spain and Sweden, January 2008

2183 Assar Lindbeck, Economic-Social Interaction in China, January 2008

2184 Pierre Dubois, Bruno Jullien and Thierry Magnac, Formal and Informal Risk Sharing in LDCs: Theory and Empirical Evidence, January 2008

2185 Roel M. W. J. Beetsma, Ward E. Romp and Siert J. Vos, Intergenerational Risk Sharing, Pensions and Endogenous Labor Supply in General Equilibrium, January 2008

2186 Lans Bovenberg and Coen Teulings, Rhineland Exit?, January 2008

2187 Wolfgang Leininger and Axel Ockenfels, The Penalty-Duel and Institutional Design: Is there a Neeskens-Effect?, January 2008

2188 Sándor Csengődi and Dieter M. Urban, Foreign Takeovers and Wage Dispersion in Hungary, January 2008

2189 Joerg Baten and Andreas Böhm, Trends of Children's Height and Parental Unemployment: A Large-Scale Anthropometric Study on Eastern Germany, 1994 – 2006, January 2008

2190 Chris van Klaveren, Bernard van Praag and Henriette Maassen van den Brink, A Public Good Version of the Collective Household Model: An Empirical Approach with an Application to British Household Data, January 2008

2191 Harry Garretsen and Jolanda Peeters, FDI and the Relevance of Spatial Linkages: Do third Country Effects Matter for Dutch FDI?, January 2008

2192 Jan Bouckaert, Hans Degryse and Theon van Dijk, Price Discrimination Bans on Dominant Firms, January 2008

2193 M. Hashem Pesaran, L. Vanessa Smith and Takashi Yamagata, Panel Unit Root Tests in the Presence of a Multifactor Error Structure, January 2008

2194 Tomer Blumkin, Bradley J. Ruffle and Yosef Ganun, Are Income and Consumption Taxes ever really Equivalent? Evidence from a Real-Effort Experiment with Real Goods, January 2008

2195 Mika Widgrén, The Impact of Council's Internal Decision-Making Rules on the Future EU, January 2008

2196 Antonis Adam, Margarita Katsimi and Thomas Moutos, Inequality and the Import Demand Function, January 2008

2197 Helmut Seitz, Democratic Participation and the Size of Regions: An Empirical Study Using Data on German Counties, January 2008

2198 Theresa Fahrenberger and Hans Gersbach, Minority Voting and Long-term Decisions, January 2008

2199 Chiara Dalle Nogare and Roberto Ricciuti, Term Limits: Do they really Affect Fiscal Policy Choices?, January 2008

2200 Andreas Bühn and Friedrich Schneider, MIMIC Models, Cointegration and Error Correction: An Application to the French Shadow Economy, January 2008

2201 Seppo Kari, Hanna Karikallio and Jukka Pirttilä, Anticipating Tax Change: Evidence from the Finnish Corporate Income Tax Reform of 2005, January 2008

2202 Walter Krämer and André Güttler, On Comparing the Accuracy of Default Predictions in the Rating Industry, January 2008

2203 Syed M. Ahsan and Panagiotis Tsigaris, The Efficiency Loss of Capital Income Taxation under Imperfect Loss Offset Provisions, January 2008

2204 P. Mohnen, F. C. Palm, S. Schim van der Loeff and A. Tiwari, Financial Constraints and other Obstacles: Are they a Threat to Innovation Activity?, January 2008

2205 Sascha O. Becker and Mathias Hoffmann, Equity Fund Ownership and the Cross-Regional Diversification of Household Risk, January 2008

2206 Pedro R. D. Bom and Jenny E. Ligthart, How Productive is Public Capital? A Meta-Analysis, January 2008

2207 Martin Werding, Ageing and Productivity Growth: Are there Macro-level Cohort Effects of Human Capital?, January 2008

2208 Frederick van der Ploeg and Steven Poelhekke, Globalization and the Rise of Mega-Cities in the Developing World, February 2008

2209 Sara Biancini, Regulating National Firms in a Common Market, February 2008

2210 Jin Cao and Gerhard Illing, Liquidity Shortages and Monetary Policy, February 2008

2211 Mathias Kifmann, The Design of Pension Pay Out Options when the Health Status during Retirement is Uncertain, February 2008

2212 Laszlo Goerke, Tax Overpayments, Tax Evasion, and Book-Tax Differences, February 2008

2213 Jun-ichi Itaya and Heinrich W. Ursprung, Price and Death, February 2008

2214 Valentina Bosetti, Carlo Carraro and Emanuele Massetti, Banking Permits: Economic Efficiency and Distributional Effects, February 2008

2215 Assar Lindbeck, Mårten Palme and Mats Persson, Social Interaction and Sickness Absence, February 2008

2216 Gary E. Bolton and Axel Ockenfels, The Limits of Trust in Economic Transactions - Investigations of Perfect Reputation Systems, February 2008

2217 Hartmut Egger and Peter Egger, The Trade and Welfare Effects of Mergers in Space, February 2008

2218 Dorothee Crayen and Joerg Baten, Global Trends in Numeracy 1820-1949 and its Implications for Long-Run Growth, February 2008

2219 Stephane Dees, M. Hashem Pesaran, L. Vanessa Smith and Ron P. Smith, Identification of New Keynesian Phillips Curves from a Global Perspective, February 2008

2220 Jerome L. Stein, A Tale of Two Debt Crises: A Stochastic Optimal Control Analysis, February 2008

2221 Michael Melvin, Lukas Menkhoff and Maik Schmeling, Automating Exchange Rate Target Zones: Intervention via an Electronic Limit Order Book, February 2008

2222 Raymond Riezman and Ping Wang, Preference Bias and Outsourcing to Market: A Steady-State Analysis, February 2008

2223 Lars-Erik Borge and Jørn Rattsø, Young and Old Competing for Public Welfare Services, February 2008

2224 Jose Apesteguia, Steffen Huck, Jörg Oechssler and Simon Weidenholzer, Imitation and the Evolution of Walrasian Behavior: Theoretically Fragile but Behaviorally Robust, February 2008

2225 Walter Krämer, Long Memory with Markov-Switching GARCH, February 2008

2226 António Afonso and Christophe Rault, What do we really Know about Fiscal Sustainability in the EU? A Panel Data Diagnostic, February 2008

2227 Sergey M. Kadochnikov and Igor M. Drapkin, Market Structure, Technological Gap and Vertical Linkage Effects from Foreign Direct Investment, February 2008

2228 Guglielmo Maria Caporale, Davide Ciferri and Alessandro Girardi, Fiscal Shocks and Real Exchange Rate Dynamics: Some Evidence for Latin America, February 2008

2229 Scott Alan Carson, Geography and Insolation in 19th Century US African-American and White Statures, February 2008

2230 Wolfgang Buchholz and Jan Schumacher, Discounting and Welfare Analysis Over Time: Choosing the η, February 2008

2231 M. Hashem Pesaran, Christoph Schleicher and Paolo Zaffaroni, Model Averaging in Risk Management with an Application to Futures Markets, February 2008

2232 Wilhelm Kohler, Offshoring: Why Do Stories Differ?, February 2008

2233 Stefan Bach, Giacomo Corneo and Viktor Steiner, Effective Taxation of Top Incomes in Germany, 1992-2002, February 2008

2234 Robert S. Chirinko, σ: The Long And Short Of It, February 2008

2235 Volker Grossmann and Holger Strulik, Should Continued Family Firms Face Lower Taxes than other Estates?, February 2008

2236 Guido Tabellini, The Scope of Cooperation: Values and Incentives, February 2008

2237 Heinrich W. Ursprung and Christian Wiermann, Reputation, Price, and Death: An Empirical Analysis of Art Price Formation, March 2008

2238 Hans Fehr and Christian Habermann, Private Retirement Savings in Germany: The Structure of Tax Incentives and Annuitization, March 2008

2239 Joseph Francois and Ian Wooton, Market Structure and Market Access, March 2008

2240 Hiroyuki Kasahara and Beverly Lapham, Productivity and the Decision to Import and Export: Theory and Evidence, March 2008

2241 Gary E. Bolton and Axel Ockenfels, Does Laboratory Trading Mirror Behavior in Real World Markets? Fair Bargaining and Competitive Bidding on EBay, March 2008