Journal of Economic Behavior & Organization 67 (2008) 832–843

JOURNAL OF Economic Behavior & Organization

www.elsevier.com/locate/econbase

Price formation in a sequential selling mechanism☆

Radosveta Ivanova-Stenzel[a], Sabine Kröger[b,*]

[a] Humboldt University Berlin, School of Business and Economics, Institute of Economic Theory I, Spandauer Str. 1, D - 10178 Berlin, Germany

[b] Université Laval, Département d'économique, and CIRPÉE, Pavillon J.A.DeSève, Québec City, Québec, G1K 7P4 Canada

Received 19 May 2006; received in revised form 4 February 2008; accepted 7 February 2008
Available online 15 February 2008

Abstract

This paper studies a dynamic trade institution, where an auction is combined with a "Buy-It-Now" option. This option presents a take-it-or-leave-it price offered by the seller to a potential buyer before the auction. If the buyer rejects this buyout price, the object is auctioned off. In equilibrium, sales should take place only in the auction. An experimental test reveals that average buyout prices and profits are well captured by the theoretical prediction. However, a substantial amount of sales takes place before the auction. This is caused by offering (too) low or accepting (too) high buyout prices. We discuss alternative explanations such as risk preferences and (wrong) formation of beliefs that might account for agents' behavior.

JEL classification: C72, C91, D44, D82

Keywords: Auction; Bargaining; Sequential mechanism; Risk preferences; Experiment

1. Introduction

Traditionally, bargaining and fixed price trading have been analyzed separately from auctions. More recently, attention has been directed to the analysis of strategic interaction in hybrid environments combining auctions and negotiations, for example, a seller who negotiates the sale of a good with a small number of buyers, before, in case of disagreement, auctioning off the good to a broader set of agents. Such a sequential mechanism is the focus of the current paper.

☆ A previous version of this paper circulated under the title "Behavior in Combined Mechanisms: Auctions with a Pre-Negotiation Stage—An Experimental Investigation." We are grateful to Charles Bellemare and Tim Grebe for detailed discussions and extended comments. We are thankful to Thomas Giebe, René Kirkegaard, Wieland Müller, Abhinay Muthoo, Stanley Reynolds, Vjollca Sadiraj, Stefan Seifert, two anonymous referees and an associate editor of this journal, seminar participants at the Economics Department at Laval University, Québec, Canada, at the Max-Planck-Insitut Jena, Germany, and at the Economics Department at University of Arizona, Tucson, USA, as well as participants at the ESA North American meeting, Tucson, USA, and IMEBE Cordoba, Spain. Financial support from the Deutsche Forschungsgemeinschaft, SFB Transregio 15, ("Governance and Efficiency as Economic Systems"), Humboldt University Berlin is gratefully acknowledged. The second author is thankful for the kind hospitality at the Humboldt University Berlin and the Economic Science Laboratory at the University of Arizona where part of this work was completed.

* Corresponding author. Tel.: +1 418 6562131x12896; fax: +1 418 6562707.
E-mail addresses: ivanova@wiwi.hu-berlin.de (R. Ivanova-Stenzel), skroger@ecn.ulaval.ca (S. Kröger).

0167-2681/$ – see front matter
doi:10.1016/j.jebo.2008.02.003

One of the reasons for the growing interest in combined mechanisms is their increasing use by successful internet auction sites such as eBay and Yahoo.[1] Those internet platforms employ a selling mechanism where a call for bids in an auction is combined with a take-it-or-leave-it price offer.[2] For example, at Yahoo sellers can offer their products throughout the auction at a permanent buyout price. At eBay, a seller may announce a temporary buyout price, the so-called Buy-It-Now price, additionally to the call for bids. Once a buyer accepts the Buy-It-Now offer, the sale is concluded at this price. Otherwise, a buyer can start the auction by submitting a bid, in which case the buyout price disappears and the final price is determined in the auction. Shortly after the introduction of the Buy-It-Now option in November 2000, eBay reported that 30 percent, 35 percent and even 45 percent (eBay Q1, Q2 and Q4 2001) of all listings in eBay auctions included a Buy-It-Now option. Similar figures have been reported by Reynolds and Wooders (in press) who observe 40 percent of auctions with a Buy-It-Now offer. The economic relevance of such hybrid mechanisms is reflected in the volume of transactions. eBay, for example, reported that fixed price trades on their sites comprised 4.4 billion USD in the first half of 2004 translating to 28 percent of the total Gross Merchandize Volume (eBay Inc., 2004).[3]

Early theoretical analysis fails to explain these observations. For example, Bulow and Klemperer (1996, p. 182) hypothesize that “if a seller could negotiate with *N* bidders while maintaining the right to subsequently hold an English auction without a reserve price and with an additional bidder, the seller would always do better to proceed directly to the auction,” a conjecture in clear contrast to the behavior observed in reality.[4]

More recently, Kirkegaard (2004) reconsiders the results by Bulow and Klemperer and shows that the English auction can be improved by negotiations prior to the auction if buyers are asymmetric or the marginal revenue is non-monotonic. Mathews (2004) shows that the probability of successful settlements using the buyout option increases with the impatience of either sellers or buyers. eBay, for example, indicates that the average auction duration decreased by almost 10 percent, which, according to eBay, is due to an increased use of the Buy-It-Now option (eBay, Q4 2002). This provides some evidence, although not conclusive, supporting the impatience hypothesis.

Risk preferences can also account for the success of the buyout option. For example, Reynolds and Wooders compare eBay and Yahoo auctions where a buyout price is offered to all buyers simultaneously. They show that the riskiness of the auction as an outside option can induce risk averse buyers to accept buyout prices that are higher than the expected price from the auction. Alternatively, Mathews and Katzman (2006) demonstrate that risk averse sellers prefer an agreement at the buy price by asking for lower prices than those expected from the auction. Budish and Takeyama (2001) study English and first price auctions with a buyout option and show that risk neutral sellers can increase their profits by offering a permanent buy price during the auction to risk averse buyers. Likewise, Hidvégi et al. (2006) consider permanent buy prices, allowing for risk aversion of both sellers and buyers. They show that properly set buy prices increase social welfare and expected utility of all agents.

In this study, we focus on auctions with a temporary buyout price such as used on eBay. In those auctions, buyers arrive at different points in time, and there is always a “decisive buyer” who will either accept the buyout price or start the auction by submitting a bid. In the latter case, the buyout price disappears; hence buyers arriving later are not informed about the rejected price offer. They can only participate in the auction.[5] Therefore, in contrast to the literature mentioned above, we propose an asymmetric treatment of buyers. More precisely, we consider a two-stage mechanism, where the seller offers a buyout price to only one of the buyers. If this buyer rejects, a second-price sealed-bid auction with additional buyers takes place. We test experimentally whether actual behavior in a controlled environment exhibits qualitative properties similar to the theoretical prediction: all sales will take place in the auction since sellers will post buyout prices that are too high to be accepted.

[1] http://www.ebay.com; http://www.yahoo.com.

[2] A take-it-or-leave-it price offer, also known as ultimatum bargaining, constitutes one of the simplest of all bargaining procedures (Binmore, 1992, p. 197). In the literature on combined mechanisms such a price offer is often referred to as “Buy-It-Now price” (as on eBay), “buyout price” or “buy price.” We will use these terms interchangeably.

[3] Two selling formats on eBay lead to a fixed price trade: an auction with a Buy-It-Now price or a pure fixed price (the item is listed at a set price and bidding is not possible). According to eBay, acceptance of the Buy-It-Now price before an auction is the primary contributor to the fixed price trades.

[4] See Kirkegaard (2006) for a general proof.

[5] Note that this combined procedure resembles the situation considered by Bulow and Klemperer: a seller can make a take-it-or-leave-it offer to a buyer who might either accept or reject the offer. In the latter case, a second buyer joins and a second-price auction takes place (Section II.C “Negotiations followed by an auction,” p. 189).

There are some empirical and experimental papers studying the Buy-It-Now option on eBay. However, these papers examine this issue in contexts quite different than we are concerned with here. For example, Anderson et al. (2004) collect transactions data from auctions conducted on eBay and study the impact of the characteristics of the seller, the good, and the transaction on seller's choice of selling format. Using a field experiment, Durham et al. (2004) look at the impact of buyer reputation, seller reputation, and the magnitude of the Buy-It-Now price on buyer behavior. Our point of concern is to understand why such dynamic trade forms exist at the first place. Therefore, we test the theoretical prediction in an environment controlling for valuations of buyers and sellers as well as for the private value character of the good for sale.

Our experimental results show that average prices and profits are well described by the theoretical prediction. However, we observe a substantial amount of sales at the buyout price and a substantial portion of buyout prices well below the theoretical prediction. Our experimental design excludes time preferences as an explanation for the observed behavior. Therefore, we investigate whether and to what extent risk preferences can help to resolve the puzzle. The literature on buyout options thus far has investigated the role of risk aversion either for permanent buyout prices (Hidvégi et al.) or in the case of temporary buyout prices only for one market side (Mathews and Katzman; Reynolds and Wooders). We study outcomes of auctions with temporary buyout prices allowing for risk aversion on both market sides.

We show that by relaxing the assumption of risk neutrality for both sellers and buyers, successful sales at the buyout price can occur in equilibrium. Using existing population estimates of risk preferences we provide quantitative predictions for the distribution of sellers' buyout prices. We find that including risk preferences can only partly account for agents' behavior: it improves the fit for buyers, but is not sufficient to explain sellers' deviations from equilibrium buyout prices. More than one-third of all observed buyout prices lie outside the predicted price range, indicating systematically under- and overpricing. Therefore, we discuss the role of belief formation and learning in order to understand and to explain agents' decisions.

The remainder of the paper is organized as follows. Section 2 introduces the model and benchmark predictions. The experimental design and procedure are presented in Section 3. Section 4 contains the analysis of the experimental results. In Section 5 we investigate whether and to what extent incorporating risk attitudes into the model can explain our experimental data. In Section 6 we discuss alternative behavioral concepts. Section 7 concludes.

2. Theoretical background: model and predictions

We consider a situation where a seller offers a single indivisible object for sale to n potential buyers.[6] The seller's valuation is common knowledge and, for simplicity, normalized to zero. Buyers' valuations for the good, v_i with $i = 1, \dots, n$, are independent random variables and private information. The seller announces a take-it-or-leave-it offer (i.e., a buyout price p) to one randomly selected buyer whom we will refer to as "buyer 1." If the offer is accepted, the sale takes place at this price. If it is rejected, a second-price sealed-bid auction without a reserve price and with all n buyers is conducted. In this case, buyers place their bids simultaneously. The bidder who submits the highest bid is awarded the object and pays a price equal to the second highest bid. The other bidders have zero payoffs.

The question arises whether the seller might benefit from the buyout option in the combined mechanism. In case of symmetric bidders and continuous valuations, a pure second-price auction without reserve price is optimal (i.e., maximizes the seller's revenue) subject to the constraint that the good is sold with probability one. In a second-price auction, every bidder truthfully reveals his private valuation. Hence, the object will always be allocated to the buyer with the highest value. If the combined mechanism leads to the same allocation of the good and agents are risk neutral, it will yield the same expected revenue to the seller (Myerson, 1981). In the combined mechanism, however, there is a positive probability that buyer 1 receives the object regardless of whether he has the highest value or not. This implies that the second-price auction without reserve price and the combined mechanism are not revenue equivalent. Since the second-price auction is optimal, the seller prefers in any case to sell in the auction.

More precisely, the interaction between seller and buyer 1 before the auction is limited to offering a take-it-or-leave-it price and the reaction to it. Buyer 1 will accept the buyout price if his payoff is at least as high as the expected payoff

[6] We thank René Kirkegaard and an anonymous referee for suggestions in this section.

from participating in the auction. That is

$$v_1 - p \geq (v_1 - \mathbf{E}[V^{(2)}|\max_{j\neq 1} V_j < v_1]) Pr(\max_{j\neq 1} V_j < v_1) \tag{1}$$

where j denotes one of the $\{2, \ldots, n\}$ bidders, v_1 buyer 1's valuation, $V^{(2)}$ the second highest valuation of all bidders, $Pr(\max_{j\neq 1} V_j < v_1)$ the probability that buyer 1 has the highest valuation of all bidders and $\mathbf{E}[\cdot]$ is the expectation operator. This means that buyer 1 will only accept prices p that are smaller or equal to a threshold price $\tilde{p}(v_1)$ satisfying

$$p \leq \tilde{p}(v_1) = \mathbf{E}[V^{(2)}|\max_{j\neq 1} V_j < v_1] Pr(\max_{j\neq 1} V_j < v_1) + v_1 \cdot (1 - Pr(\max_{j\neq 1} V_j < v_1)). \tag{2}$$

Therefore, at a given buyout price p, only the types of buyer 1 satisfying inequality (2) will accept this price. In the auction, however, those types generate an expected profit for the seller of

$$\mathbf{E}[V^{(2)}|\max_{j\neq 1} V_j < v_1] Pr(\max_{j\neq 1} V_j < v_1) + \mathbf{E}[V^{(2)}|\max_{j\neq 1} V_j \geq v_1](1 - Pr(\max_{j\neq 1} V_j < v_1)), \tag{3}$$

which is higher than the buyout price p because $\mathbf{E}[V^{(2)}|\max_{j\neq 1} V_j \geq v_1] \geq v_1$. Consequently, there exists no buyout price at which seller and buyer 1 could reach an agreement.

In the following, we derive the closed form solution for the parameters used in the experiment, that serves as our benchmark prediction. In the experiment, valuations of $n = 2$ buyers were drawn from the uniform distribution with support [0,1], where we denote by $F(x) = x$ and by $f(x) = 1$ the cumulative distribution function of each buyers' values and its probability distribution function, respectively. Buyer 1 will accept a buyout price if

$$p \leq \tilde{p}(v_1) = \frac{(1 - (1 - v_1)^2)}{2}. \tag{4}$$

In order to ensure not to receive less than the expected profit in the auction, the seller should ask at least for the threshold price of the bidder with the highest valuation. The optimal buyout prices are therefore $p^* \geq \tilde{p}(1) = 1/2$, which will never be accepted. If the seller asked for a buyout price below 1/2, buyers with a value above a threshold $\tilde{v}(p)$ (i.e., $v_1 \geq \tilde{v}(p) = 1 - \sqrt{1 - 2p}$) would accept this offer. Therefore, the fact that a buyout price below 1/2 has been rejected or accepted reveals information about buyer 1's type, allowing the seller to update his information about the support of the valuation of buyer 1 and therefore his expected profit from the auction.

More precisely, the seller's expected profit from the combined mechanism is

$$\Pi^S(p) = p \int_{\tilde{v}(p)}^{1} f(x)\,\mathrm{d}x + \int_0^{\tilde{v}(p)} y g^{(1)}(y, \tilde{v}(p))\,\mathrm{d}y \int_0^{\tilde{v}(p)} f(x)\,\mathrm{d}x. \tag{5}$$

The second term of Eq. (5) denotes the expected profit in the auction that takes the adverse selection effect into account (i.e., that either one or both bidders' valuations lay in the interval $[0, \tilde{v}(p)]$ (with $\tilde{v}(p > 1/2) = 1$) and where $g^{(1)}(y, \tilde{v}(p))$ denotes the density function of the first order statistic with one or two random variables in this interval).[7]

Since sales will always take place in the auction, ex ante expected earnings are 1/3 for the seller and 1/6 for each buyer.

3. Experiment

In the experiment, each participant was either a buyer or a seller. One seller and two buyers constituted a trading group. The composition of the trading groups was changed between periods: each period sellers and buyers were rematched randomly. An experimental session consisted of 32 periods, which were divided into four cycles of eight trading periods. Each buyer was in the role of buyer 1 either in the odd or the even cycles (i.e., in 16 out of the 32 periods). In all other periods a buyer joined the auction only if the price offer had been rejected. Buyers' private reselling values for the product were randomly and independently drawn from the set $V = \{0, 1, 2, \ldots, 99, 100\}$ with all $v_i \in V$ being equally likely. Participants could choose integer buyout prices and bids between 0 and 100. All values were denoted in a fictitious currency termed ECU for Experimental Currency Unit.[8]

[7] See Rohatgi (1987) for distributions of order statistics with random sample size. For the problem here, $g^{(1)}(y, \tilde{v}(p)) = (1 + \tilde{v}(p) - 2y)/\tilde{v}(p)$.

[8] See Appendix A (available on the JEBO website) for a shortened and translated version of the instructions. Complete sets of the original instructions (in German) are available upon request.

Table 1
Descriptive statistics for all periods and per cycle (1 cycle = 8 periods) and theoretical predictions

Cycle	Number of observations	(1) Buyout price		(2) Acceptance rate (in percent)	(3) Profits				(4) Efficiency (in percent)
					Seller		Buyer [a]		
		Mean	S.D.		Mean	S.D.	Mean	S.D.	
1	240	0.53	0.17	35	0.34	0.19	0.14	0.21	80.4
2	240	0.51	0.18	35	0.32	0.20	0.16	0.23	85.0
3	240	0.52	0.18	28	0.33	0.19	0.16	0.23	87.9
4	240	0.50	0.16	36	0.32	0.18	0.16	0.23	86.7
All	960	0.51	0.17	33	0.33	0.19	0.15	0.22	85.0
Theory		$0.5 \leq p \leq 1$		0	0.33		0.17		100

[a] Number of observations = 480 per cycle, since there were two buyers per trading round.

In the beginning of a single period the trading groups were formed and sellers were asked to submit their buyout price offer. The buyer who was in the role of buyer 1 was informed about his private value and the buyout price. After buyer 1 had accepted or rejected the buyout price, each group member was told whether or not a sale had taken place. If the price offer was accepted, the sale occured and there was no auction. If the buyout price was rejected, the other buyer was informed about his private value and a second-price sealed-bid auction with both buyers took place. All group members were informed about the outcome of the auction: who won the auction, the price paid by the winner, and their own payoff in the current period. Buyers who attended only the auction were not informed about the buyout price. At the end of a period, each participant received an account of the own total profit up to this period. In a post-experimental questionnaire participants answered standard demographic questions and were asked to comment briefly on their reasoning during the experiment.

The experiment comprised 10 sessions with a total of 90 participants. Half of the sessions were with 6 and the other half with 12 participants each. We pooled the data from the sessions of different sizes since we did not find significant differences between them.[9] All experimental sessions were computerized, and the software system was created with z-Tree (Fischbacher, 2007). The experiment was conducted at Humboldt-University Berlin, Germany, and most participants were students of economics, business administration, law, and physics. One session lasted on average 90 min. The conversion rate of ECU earned by each participant into cash was 1 ECU = 0.0125 EUR. Participants' total earnings ranged between 8.05 EUR and 16.86 EUR with a mean of 11.82 EUR (as a seller: 13.14 EUR, as a buyer: 11.15 EUR).[10]

4. Results

Table 1 reports descriptive statistics of the experimental outcomes for all periods and, in order to illustrate the changes over time, for all cycles. For ease of comparison to the theoretical benchmark, we report our results for normalized valuations (i.e., all experimental outcomes are transformed from the [0,100] to the [0,1]—range). In total we observe 960 trades. Average buyout prices (column 1) with $p = 0.51$ as well as average earnings for sellers and buyers with 0.33 and 0.15 (columns 3) are rather stable over time and in line with the theoretical prediction.[11] The object was allocated to the buyer with the highest valuation in 85 percent of all sales (column 4). Since in the combined mechanism buyer 1 receives the object with positive probability, regardless of whether he has the highest value or not, one might expect efficiency to be lower than in a pure second-price auction. However, the observed share of efficient allocations is quite comparable to those reported in other second-price sealed-bid auction experiments with two bidders (e.g., 88 percent in Güth et al., 2005, and 91 percent in Pezanis-Christou, 2002).

[9] The p-value of Mann–Whitney U-tests (henceforth MWU) on session averages concerning price offers, acceptance rates, final sale prices and efficiencies ranges between 0.22 and 1.

[10] These numbers include a starting capital for buyers of 5 EUR.

[11] Comparing the session averages of those variables and the benchmark prediction reveals no significant difference (MWU-test: $p = 0.45$ for seller as well as buyer earnings; sign-test: $p = 0.62$ for buyout prices). The same holds for the comparison between the first and the second halves of the experiment (Wilcoxon signed rank test: $p > 0.28$ for all variables).

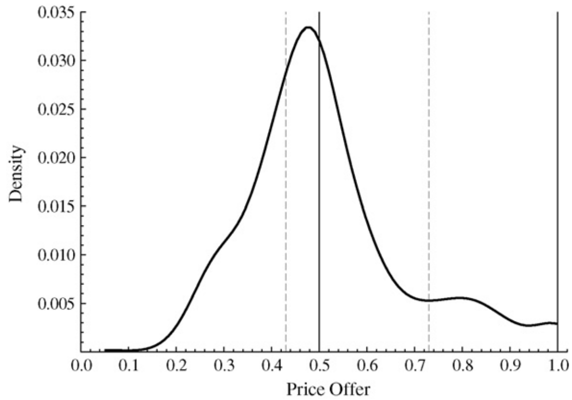

Fig. 1. Buyout price density estimation, gaussian kernel. Solid lines represent the interval of optimal buyout prices assuming risk neutral agents. Dashed lines represent the interval of buyout prices predicted by the simulation for agents with heterogeneous risk preferences

These results show that experimental outcomes are on average well captured by the benchmark prediction. However, contrary to the theoretical prediction, we observe one-third of all sales taking place at the buyout price before the auction (see column 2 of Table 1).[12] In order to understand these findings better, we proceed with a more detailed analysis of the data. More precisely, we first investigate buyers' bidding behavior in the auction and continue with their acceptance behavior of the buyout price before we turn our attention to sellers' price setting behavior.

Bidding should (theoretically) not be influenced by the buyout option since bidding the own value is a weakly dominant strategy. Indeed, experiencing the buyout option before the auction seems not to change bidding behavior in the auction. Despite some slight over-and underbidding, we do not find significant differences between bids and the predicted equilibrium bidding strategy. In fact, half of all observed bids are equal to subjects' valuations. When comparing the session average relative bid deviations to a zero vector, we cannot reject the hypothesis that the relative bid deviation is equal to zero.[13] Moreover, bid deviations do not significantly change over time.[14] Given these observations, truthful bidding turns out to be a reasonably good prediction.

Buyer 1 had to choose between accepting the buyout price or going to the auction. In the experiment, each buyer was confronted half of the time with buyout prices, which leaves us with information about the acceptance behavior of each buyer for 16 periods. We find that only five out of 60 buyers (8 percent) behaved according to the theoretical benchmark (i.e., accepted buyout prices that were below their threshold price, $\tilde{p}(v_1)$ as defined in Eq. (4), and rejected otherwise). The majority of buyers either accepted buyout prices that were above their threshold price (57 percent) or rejected buyout prices that were below their threshold price (18 percent), which obviously violates the theoretical prediction.[15]

Sellers had to choose the buyout price. Fig. 1 presents the density estimate of buyout prices ranging from 0.05 to 1. The vertical lines present visual aids to compare our data to different theoretical benchmarks. The interval of buyout prices included by the two solid lines present the (risk neutral) benchmark. The two dashed lines describe the interval of buyout prices based on a model relaxing risk neutrality that we present below. Half of the buyout prices (48.4 percent) conform with the risk neutral theoretical prediction (i.e., are greater or equal to 0.50). However, the other half (51.6 percent) of observed buyout prices is below 0.50, which is clearly not in line with the theoretical benchmark.

All these observations indicate that the game theoretic model assuming homogeneous and risk-neutral agents seems to be a good predictor for average buyout prices and profits. Nevertheless, it cannot very well rationalize observed individual decisions that appear to drive the finding of successful sales at the buyout price.

[12] Comparing the acceptance rates on session level between the first and second halves of the experiment reveals that there are no significant changes over time (MWU-test: $p = 0.729$). Considering only buyout prices that were below the valuation of buyers, 63 percent of all sales occur before the auction.

[13] MWU-test: $p = 0.114$ for both the first and second part of the experiment.

[14] Wilcoxon signed rank test: $p = 0.625$ comparing session averages of first and second part of the experiment.

[15] The remaining 10 buyers (17 percent) could not be classified since the deviation of their behavior was not consistent (i.e., the same person accepted buyout prices a risk neutral person would have rejected and would reject other buyout prices a risk neutral person would have accepted).

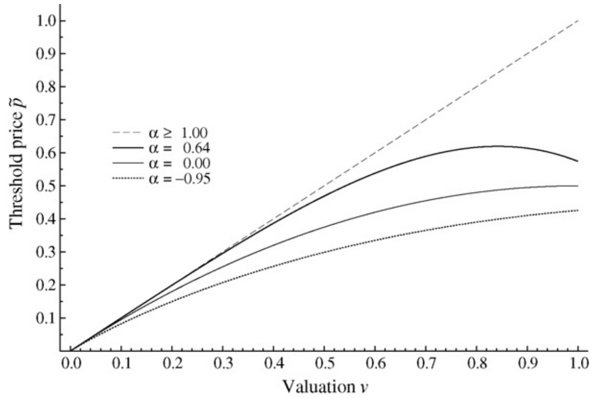

Fig. 2. Relation between threshold price and valuation for different levels of risk aversion.

This leaves us with the possibility that behavior might be based upon unobserved heterogeneity in subjects' preferences. One possible explanation for the observed deviations might be heterogeneity in risk preferences of both sellers and buyers. Risk aversion might explain, for instance, why the majority of buyers accepted buyout prices, which yield a lower profit then the expected profit from the auction and which a risk neutral buyer would have rejected. On the other hand, buyout prices below the theoretical prediction might be driven by risk aversion of sellers.

5. Risk preferences

In this section we investigate whether and to what extent incorporating risk preferences into the model can explain our experimental data. Therefore, all agents are assumed to be expected utility maximizers and have preferences that can be represented by a utility function $u(\cdot)$ that is twice differentiable, strictly increasing, and satisfies $u''(\cdot) \leq 0$ everywhere on its support. More precisely, for our analysis we restrict risk preferences to belong to the class of constant relative risk aversion (hereafter CRRA), $u(x) = x^{(1-\alpha)}/(1-\alpha)$, where α is the Arrow-Pratt measure of relative risk preferences.[16,17]

5.1. Buyer

Allowing buyers to be risk averse might explain acceptance of buyout prices that lay above the expected price of the auction. In the following, we are going to investigate whether the level of risk aversion needed to account for the observed acceptance behavior is reasonable.

Given our distributional assumptions and functional form, the threshold price of a buyer with risk preferences of α and valuation v is defined by

$$\tilde{p}(v) = v - \left(\frac{v^{(2-\alpha)}}{2-\alpha}\right)^{(1/(1-\alpha))}. \tag{6}$$

Fig. 2 plots the relation between threshold price and valuation for buyers with different levels of risk preferences. A buyer who is risk neutral ($\alpha = 0$) will never accept a buyout price above 0.50, whereas a buyer who has risk preferences of $\alpha = 0.64$ might accept buyout prices up to 0.62. Note that the relation between threshold price and valuation might be non-monotone. For example, a buyer with $\alpha = 0.64$ will not only reject a buyout price of 0.60 if his valuation is lower than 0.73 but also if his valuation is higher than 0.95, implying that a buyer with a very high valuation is more

[16] This specification implies risk loving behavior for $\alpha < 0$, risk neutrality for $\alpha = 0$ and risk aversion for $\alpha > 0$. When $\alpha = 1$, the natural logarithm, $u(x) = \ln(x)$, is used.

[17] We experimented also with an exponential utility function, $u(x) = 1 - \exp^{-\alpha x}$, a specification that exhibits constant absolute risk aversion (hereafter CARA). We found that the CRRA form of the utility function fits our data better than the CARA form (see footnote 22).

likely to go for the risky outcome from the auction rather than to accept the certain outcome from the buyout option.[18] If a buyer is sufficiently risk averse (i.e., $\alpha \geq 1$), he would be willing to accept a buyout price equal to his valuation.[19] The threshold is shifted upwards with increasing levels of α implying that higher buyout prices are more likely to be accepted when a buyer is more risk averse.

We can estimate the average risk preference of the buyer population directly from the experimental data.[20] In the model, a buyer accepts a buyout price if

$$u(v-p)+\varepsilon_1 \geq \int_0^v u(v-x)f(x)\,\mathrm{d}x+\varepsilon_2, \tag{7}$$

where we assume the unobservable error terms, ε_i with $i \in \{1, 2\}$, to follow a normal distribution $\varepsilon_i \sim N(0, \sigma^2)$. Assuming that risk preferences can be represented by $u(x) = x^{1-\alpha}/(1-\alpha)$, given the distributional assumptions, and the decision rule in Eq. (7), we estimate the risk preference parameter in the buyers' population by maximum likelihood.[21]

The parameter estimate for buyers' risk preferences is $\hat{\alpha} = 0.64$ with a loglikelihood function value of −288.61. The model assuming risk neutral buyers (fixing $\alpha = 0$) results in a likelihood function value of −319.56. A likelihood ratio-test with a test value of 61.89 (5 percent χ^2 critical value of 3.84) corroborates that allowing for risk preferences improves the fit of the data significantly.[22] Overall, buyers seem to have reacted in a risk averse manner, which could explain the acceptance of high buyout prices.

The estimated level of buyers' risk attitudes is in line with estimates reported in the literature. Numerous studies also based on responses of student subjects estimated average levels of risk preferences to be around 0.3–0.7 and to be quite robust to different decision environments (e.g., gambles, other individual decision tasks, games, and auctions). For example, Cox and Oaxaca (1996), Goeree et al. (2002), Chen and Plott (1998) and Ivanova-Stenzel and Salmon (2004) estimate relative risk preferences from bids observed in private value auction experiments to be $\alpha = 0.67, 0.52, 0.48$, and 0.34, respectively. Goeree et al. (2003) and Holt and Laury (2002) use experimental data from individual decision tasks to estimate risk parameters for each person. They find the average risk attitudes to be around 0.28 and 0.32, respectively.[23]

5.2. Seller

Contrary to the risk neutral case, we could not derive an explicit solution for the case of general risk preferences for sellers. Nevertheless, we can solve numerically for optimal buyout prices and derive the distribution of buyout prices for any given distribution of agents' risk attitudes. This also implies that we can only indirectly investigate whether risk preferences improve the fit of the model for sellers. We will do this by comparing predicted buyout price distributions to the buyout prices observed in the experiment.

To derive quantitative buyout price predictions using "reasonable" levels of risk preferences, we rely on estimated risk preference distributions in the literature. In order to check for the robustness of the predictions, we will use four

[18] This non-monotonicity is driven by the choice of the utility function. With exponential utility, for example, monotonicity of the threshold price–value relation is maintained.

[19] This is somewhat counterintuitive since in this case such a buyer would get a zero payoff whereas by rejecting he would have a chance of a positive net payoff. On the other hand, given the assumption of CRRA and the specification of the utility function, a buyer with $\alpha > 1$ and a valuation between 0 and 1 will always have negative net utility from entering the auction.

[20] Since we do not have enough data points to estimate the risk preference parameter of each individual buyer, the analysis is based on pooled data and reflects the behavior of the entire population assuming homogeneity and stable preferences among buyers.

[21] For $v \geq p$ ($N = 482$) the choice probabilities are given by

$$Pr\left(\varepsilon > -\frac{1}{1-\alpha}\left((v-p)^{(1-\alpha)} - \frac{v^{(2-\alpha)}}{2-\alpha}\right)\right)$$

with $\varepsilon = \varepsilon_1 - \varepsilon_2$ and $\varepsilon \sim N(0, 1)$.

[22] Modeling the decision using the exponential utility function results in a loglikelihood function value of −300.58. A comparison with the model applying the power utility function results in a likelihood ratio test value of 23.94, indicating a significantly better fit for the power function model.

[23] Holt and Laury compare different treatments with normal and extremely high payments. We concentrate on their treatment with comparable real monetary incentives.

Table 2
Summary of different studies

(1) Study	(2) Sample size	(3) Distribution of α-estimates quantiles			(4a) Prediction of buyout prices (p) quantiles			(4b) Acceptance rate (in percent)
		10 percent	50 percent	90 percent	10 percent	50 percent	90 percent	
G,H&P	42	−0.38	0.14	0.89	0.48	0.55	0.68	17.6
H&L	175	0.00	0.28	0.83	0.48	0.53	0.56	21.2
C&O	40	0.37	0.72	0.92	0.53	0.54	0.56	36.9
I&S	55	−0.17	0.46	0.72	0.51	0.51	0.56	25.8

Col. (1): "G,H&P"—Goeree et al. (2003), "H&L"—Holt and Laury (2002), "C&O"—Cox and Oaxaca (1996), "I&S"—Ivanova-Stenzel and Salmon (2004); Col. (3): reported distribution of individual risk preference estimates; Col. (4a) and (4b): predicted buyout price distributions and acceptance rates based on those estimates

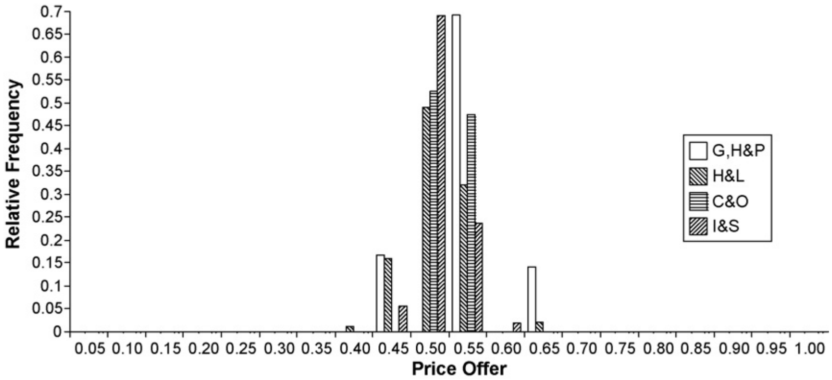

Fig. 3. Predicted buyout prices based on risk preference distribution estimates of four studies. "G,H&P"—Goeree et al. (2003), "H&L"—Holt and Laury (2002), "C&O"—Cox and Oaxaca (1996), "I&S"—Ivanova-Stenzel and Salmon (2004)

different frequency distributions of estimated individual risk preferences provided by the studies mentioned above, Cox and Oaxaca (henceforth C&O) and Ivanova-Stenzel and Salmon (henceforth I&S) for auctions, and Goeree et al. (2003, henceforth G,H&P) and Holt and Laury (henceforth H&L) for individual decision tasks. Estimated individual risk attitudes of those four studies lay in the interval $-1.48 \leq \alpha \leq 1.37$, and characteristics of the distributions (10 percent, 50 percent and 90 percent quantiles) are reported in column 3 of Table 2.[24] Based on these distributions we simulate outcomes allowing for varying risk preferences of both buyers and sellers.[25] We investigate the quantitative change in the decision variables and the expected amount of accepted buyout prices.

Buyers will accept buyout prices up to 1 when they are extremely risk averse (i.e., $\alpha \geq 1$) and only up to 0.45 when they are risk loving (i.e., $\alpha = -1.48$), which is the lowest level of the preference parameter estimates reported by the four studies. This translates into an increase (decrease) of the maximum buyout price a risk neutral agent would accept by 100 percent (− 10 percent). In order to determine the impact of risk on sellers' price setting behavior, we simulate buyout prices offered by sellers with different risk attitudes. Sellers calculate their optimal buyout price given their own risk attitude and the distribution of risk preferences within the buyer population. Fig. 3 presents predicted buyout price frequencies for each of the estimated population distributions of risk preferences. The range predicted by all simulations is $p \in [0.43, 0.69]$. Compared to the benchmark case with risk neutral sellers and risk neutral buyers, this range is smaller in size and is shifted downwards. However, the lowest buyout price is 0.43 and thus still close to the prediction for risk neutral agents of 0.5. The stronger decrease of the upper bound from 1.0 to 0.69 can be explained by the fact that buyers who are risk averse might accept even high buyout prices. The seller can therefore increase his profit by keeping buyout prices high yet affordable for risk averse buyers.

[24] H&L classify their participants in nine categories. We assign all subjects within a category the mean of this category as individual risk parameter. Subjects in the outer categories, $\alpha < -0.95$ and $\alpha > 1.37$, were assigned $\alpha = -0.95$ and $\alpha = 1.37$. We do the same for G,H&P who distinguish between seven risk categories, with $\alpha < -0.56$ and $\alpha > 0.93$ as lower and upper bound.

[25] We assume common knowledge about the population distribution of risk attitudes, which is assumed to be the same for both buyers and sellers.

Medians and quantiles of simulated buyout prices are reported in Table 2. Medians of simulated buyout prices are similar between studies (0.55 (G,H,&P), 0.53 (H&L), 0.54 (C&O) and 0.51 (I&S)), however, simulated buy prices based on parameter estimates elicited via lottery choices (G,H&P and H&L) are spread more widely than those based on auction data (C&O and I&S). Column 4b of Table 2 reports acceptance rates given the simulated price distributions and the distribution of risk preferences among buyers. Predicted acceptance rates of buyout prices vary between 17.6 percent and 36.9 percent.

Thus, allowing for general risk preferences opens a price floor within which agreements are possible at the buyout price. The exact magnitude depends, however, on the particular preference distribution.

The acceptance rate of 33 percent observed in our experiment lies within the range of predicted acceptance rates with heterogeneous agents. The median buyout price observed in the experiment is 0.49 and thus close to the predicted median buyout prices (0.51–0.55). However, allowing for heterogeneous risk preferences seems only slightly to improve the explanatory power of the model for sellers. The interval of buyout prices $p \in [0.43, 0.69]$ including all predictions covers more than half of observed buyout prices (56 percent) (see Fig. 1), which is an improvement by 8 percent with regard to the (risk neutral) benchmark prediction. Still, almost half of the buyout prices remain unexplained. Buyout prices in lower and higher ranges are much more dispersed than predicted. A Pearson Goodness of fit test strongly rejects the prediction of the model for each of the estimated distributions.[26]

Furthermore, whereas buyout prices above 0.69 could still be explained by the risk neutral benchmark, buyout prices below 0.43 are not captured by any of the two models. What is puzzling is that those low offers comprise 29 percent of all buyout prices. This fraction is not only stable over time but is generally caused by the same sellers:[27] One-third of all sellers offer buyout prices below 0.43 more than half of the time.

6. Discussion

In the previous section we have shown that allowing for risk preferences can theoretically explain the existence of successful sales at the buyout price (i.e., before the auction). Participants' comments in the post-experimental questionnaire confirm that risk attitudes might indeed drive the offer of low or the acceptance of high buyout prices. For example, some participants in the seller's role were concerned about the auction because it generated "too volatile prices" and mentioned that they favored an agreement at the buyout price. Participants in the role of buyers emphasized that they preferred to exercise the buyout option as "the chance of buying the item seemed to be higher." The analysis of the experimental data shows, however, that behavior can only partly be explained by agents' risk attitudes. Even though relaxing the assumption of risk neutrality improves significantly the fit for buyers' behavior, it is not sufficient to explain sellers' behavior. Observed buyout prices vary much more than predicted with a substantial number of buyout prices outside the predicted interval.

One might be tempted to explain the excess variance in buyout prices with noise in behavior. There are two problems with this argument. First, the observed distribution of buyout prices is far from uniform, as would be required by a model where sellers randomly choose across all possible price offers. Second, noise would presumably decline over time as participants have the opportunity to learn and adjust their buyout prices during the 32 periods of the experiment. Nevertheless, the distribution of observed buyout prices remains stable over time regarding systematical under- and overpricing compared to the theoretical predictions.

Another possible way to look at our results is in light of the bargaining literature investigating environments with asymmetric information. For example, Samuelson and Bazerman (1985) show that subjects systematically deviate from the predicted behavior and fall prey to the "winner's curse," in the sense that they either enter into loss-making purchases or forgo profit-making opportunities. The latter might apply also to our experimental situation where the seller (the uninformed party) has to condition his behavior on the strategic reaction of buyer 1 (his informed opponent).

In our experiment a seller faced a cognitively very demanding decision task. First, he had to consider buyers' acceptance threshold values. Second, conditional on these threshold values, he had to calculate his own expected utility

[26] The Pearson Goodness of fit test, also known as χ^2-test, requires independent observations. As different price offers of an individual seller might not fulfil this requirement, we use mean price offers of individual sellers for the test.

[27] A comparison of the fraction of low offers between the first and second part of the experiment reveals no significant changes (Wilcoxon signed rank test: $p = 0.977$).

for different buyout prices and had to choose the price offer that maximizes this utility. Such reasoning requires that sellers not only optimized but also correctly conditioned their buyout price on the buyer's reaction.

Let us consider, for example, a risk neutral seller who does not condition on the buyer's reaction. From the point of view of such a (boundedly rational) seller, any accepted buyout price, p, that is above the expected profit from the auction will be preferable. Such a seller has to ensure that his price offer is above the expected auction profit and will be accepted with positive probability such that $E(\pi_A) \leq p \leq E(v)$. For example, with $E(\pi_A) = 1/3$ and $E(v) = 1/2$, a non-conditioning seller might offer prices between $1/3 \leq p \leq 1/2$.

Following this argument we can explain not only the existence of low buyout prices, but also the fact that the observed underpricing remains stable over time. Suppose a seller neglects the strategic reaction of a buyer towards his own buyout price and uses solely his experience to build his expectations. A seller who offers low buyout prices in order to avoid low (expected) outcomes in the auction is more likely to experience low profits in the auction. Buyers with relatively high values will accept low buyout prices, but buyers who cannot even afford those low price offers will reject them and go to the auction. This leads to the selection of low value buyers into the auction and consequently to low auction prices, reinforcing the seller's expectation about the low prospects of the auction. In order to form correct experience-based beliefs about auction profits, sellers would need to experiment with different buyout prices. Experimentation might be costly in the short run (because of forgone profit opportunities) and might therefore deliberately not be chosen by decision makers (Einhorn and Hoghart, 1981). However, if a seller is not willing to "invest" in learning, he might form wrong expectations about the relation between demanded buyout price and final profits.

In the answers to the post-experimental questionnaire we found some evidence supporting the possibility that subjects in our experiment fell prey to such a "seller's curse." Some sellers indeed argued that the "auction generated too low prices" and that this was the reason why they preferred to reach an agreement before the auction. However, our data does not allow us to test this conjecture directly. Experimental investigations of this issue and other behavioral explanations might yield promising answers and further insights in behavior in such hybrid mechanisms.[28]

7. Summary

In this paper, we studied a dynamic trade institution that combines an auction with a take-it-or-leave-it price offer, the so called buyout option. Under standard assumptions, sales should always take place in the auction, rendering the existence of the buyout option obsolete. An experimental test of the theory suggested that the theoretical benchmark can very well predict average buyout prices and profits in such combined mechanism. However, persistent departure from the theoretical prediction at the individual level (i.e., offering too low or accepting too high buyout prices) resulted in a substantial volume of transactions before the auction. This finding led us to explore other behavioral explanations in order to account for the observed deviations.

First, we showed that allowing for individual heterogeneity in risk preferences for both market sides might lead to sales before the auction. This result seems to be driven rather by risk aversion of buyers than of sellers as risk averse buyers should accept higher buyout prices more frequently. Risk averse sellers, on the other side, lower their buyout price only marginally below the lowest price set by risk neutral sellers. By using existing population estimates of risk preference parameters, we were able to make quantitative predictions about the distribution of sellers' price offers and acceptance behavior of buyers. When we compared these predictions to the experimental data, however, we found that behavior can be explained only partially by agents' risk preferences. Relaxing the assumption of risk neutrality improved the fit of the model for buyers, but could not account for a substantial part of individual sellers' decisions. This was especially striking for low price offers. Even though participants gained experience during the experiment, they did not increase their buyout prices, thus their profit opportunities.

We therefore discussed an alternative explanation of individual seller behavior. An important issue seems to be whether sellers are aware of the information asymmetry. If sellers fail to anticipate buyers' strategic reaction to their buyout price, they might choose buyout prices that are too low and forgo profit-making opportunities, a fallacy that resembles the winner's curse in negotiations.

[28] We plan to pursue such investigations in our future research.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jebo.2008.02.003.

References

Anderson, S.T., Friedman, D., Milam, G.H., Singh, N., 2004. Seller strategies on eBay. UC Santa Cruz Economics Working Paper No. 564.
Binmore, K., 1992. Fun and Games. D.C. Heath and Company, Lexington, MA.
Budish, E.B., Takeyama, L.N., 2001. Buy prices in online auctions: irrationality on the Internet? Economics Letters 72, 325–333.
Bulow, J., Klemperer, P., 1996. Auctions versus negotiations. American Economic Review 86, 180–194.
Chen, K.-Y., Plott, C.R., 1998. Nonlinear behavior in sealed bid first-price auctions. Games and Economic Behavior 25, 34–78.
Cox, J.C., Oaxaca, R.L., 1996. Is bidding behavior consistent with bidding theory for private value auctions? In: Isaac, R.M. (Ed.), Research in Experimental Economics, vol. 6. JAI Press Inc., pp. 131–148.
Durham, Y., Roelofs, M.R., Standifird, S.S., 2004. eBay's Buy-It-Now function: who, when, and how. Topics in Economic Analysis & Policy 4, 1–22.
Einhorn, H.J., Hoghart, R.M., 1981. Behavioral decision theory: processes of judgment and choice. Journal of Accounting Research 19, 1–31.
eBay Inc., 2004. Financial results: First, Second and Fourth Quarter 2001. Fourth Quarter 2002. First and Second Quarter 2004. eBay Inc.
Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. Experimental Economics 10, 171–178.
Goeree, J.K., Holt, C.A., Palfrey, T., 2002. Quantal response equilibrium and overbidding in private value auctions. Journal of Economic Theory 104, 247–272.
Goeree, J.K., Holt, C.A., Palfrey, T., 2003. Risk averse behavior in generalized matching pennies game. Games and Economic Behavior 45, 97–113.
Güth, W., Ivanova-Stenzel, R., Wolfstetter, E., 2005. Bidding behavior in asymmetric auctions: an experimental study. European Economic Review 4, 1891–1913.
Hidvégi, Z., Wang, W., Whinston, A.B., 2006. Buy-Price English auction. Journal of Economic Theory 129, 31–56.
Holt, C.A., Laury, S.K., 2002. Risk aversion and incentive effects. American Economic Review 92, 1644–1655.
Ivanova-Stenzel, R., Salmon, T.C., 2004. Bidder preferences among auction institutions. Economic Inquiry 42, 223–236.
Kirkegaard, R., 2004. Auctions versus negotiations revisited. Working Paper. University of Aarhus.
Kirkegaard, R., 2006. A short proof of the Bulow-Klemperer auctions vs. negotiations result. Economic Theory 28, 449–452.
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 (Zeitschrift für Nationalökonomie) 81, 25–52.
Mathews, T., Katzman, B., 2006. The role of varying risk attitudes in an auction with a buyout option. Economic Theory 6, 597–613.
Myerson, R.B., 1981. Optimal auction design. Mathematics of Operations Research 6, 58–73.
Pezanis-Christou, P., 2002. On the impact of low-balling: experimental results in asymmetric auctions. International Journal of Game Theory 31, 69–89.
Reynolds, S., Wooders, J., in press. Auctions with a buy Price. Economic Theory, doi:10.1007/s00199-006-0182-7.
Rohatgi, V.K., 1987. Distribution of order statistics with random sample size. Communications in Statistics Theory and Mehods 16, 3739–3743.
Samuelson, W.F., Bazerman, M.H., 1985. The winner's curse in bilateral negotiations. In: Smith, V.L. (Ed.), Research in Experimental Economics, vol. 3. JAI Press Inc., pp. 105–137.