ELSEVIER

Decision Support Systems 29 (2000) 249–268

Decision Support Systems

www.elsevier.com/locate/dsw

An experimental analysis of multi-attribute auctions

Martin Bichler

Department of Information Systems, Vienna University of Economics and Business Administration, Augasse 2-6, A-1090 Austria

Accepted 3 April 2000

Abstract

Auctions are a fundamental mechanism to automating negotiations in electronic commerce. We investigate multi-attribute auctions, an economic mechanism, which automates negotiation on multiple attributes of a deal. In this approach, we combine decision analysis techniques and single-sided auction mechanisms in order to procure goods and services. The paper describes an experimental analysis of multi-attribute auctions. First, we will provide an overview of the preliminary game-theoretical and simulation results. Then we will introduce a Web-based implementation of the mechanism and describe the design and the results of an experiment analyzing the economic behavior of multi-attribute auction formats. In the experiment, the utility scores achieved in multi-attribute auctions were significantly higher than those of single-attribute auctions. The efficiency was similar in single-attribute and multi-attribute auctions, and we did not find evidence for revenue equivalence between the multi-attribute auction formats. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Multi-attribute auction; Experimental economics; Decision analysis; Electronic market

1. Introduction

Auctions account for an enormous volume of transactions on the Internet and are a fundamental mechanism to automating price negotiations in electronic commerce. However, in many situations, it is crucial to have negotiations on multiple attributes of a deal such as quality, delivery time or terms of payment. In this paper, we focus on *multi-attribute auctions*, a class of market mechanisms, which enables automated negotiation on multiple attributes of a deal. This feature is especially useful in procurement auctions where buyers negotiate with multiple suppliers over heterogeneous goods. In Ref. [5], we proposed a mechanism and described an implementation system for multi-attribute auctions.

So far there is little knowledge about the economic behavior of multi-attribute auctions and only a small number of game-theoretical articles have been published in this field. In order to deploy multi-attribute auctions in a real-world setting, one has to answer a range of questions about optimal bidding strategies and equilibrium values achieved in these new auction formats. Game-theoretical analysis and computational analysis are one way of finding answers to these questions. Laboratory experimentation is an important methodology not only to test the outcomes of analytical analyses but also to learn more about the applicability and the deployment of a new institution in a real-world setting. This is particularly interesting in the case of multi-attribute auctions where the actual implementation is of consider-

E-mail address: martin.bichler@wu-wien.ac.at (M. Bichler).

PII: S0167-9236(00)00075-0

able importance to the outcome of the mechanism. This paper describes an experimental analysis of multi-attribute auctions. First, we will give an overview of preliminary game-theoretical and simulation results. We will then introduce a Web-based implementation system of the institution and describe the design and the results of an experiment analyzing the economic behavior of multi-attribute auction formats in the context of OTC trading with financial derivatives.

1.1. Experimental economics

Laboratory experiments have become an important source of data for economists. For economic mechanism design, it is useful to study new institutions in the laboratory before introducing them in the field. Laboratory experimentation can facilitate the interplay between the evolution and modification of proposed new exchange institutions. Experimenters can repeat testing to understand and improve the features of new market mechanisms.

Although the results of laboratory experiments are interesting, there is still a question of whether the findings of experimental tests can be generalized. Experimental sciences use "induction" as a basic principle and assume that regularities observed will persist as long as the relevant underlying conditions remain substantially unchanged. What makes experiments so different from other methods microeconomists use is the presence of human subjects. Smith [35] refers to this question as the "parallelism precept": "Propositions about the behavior of individuals and the performance of institutions that have been tested in laboratory micro-economies apply also to non-laboratory micro-economies where similar ceteris paribus conditions hold." Meanwhile, experiments are commonplace in game theory, finance, electronic commerce and many other fields.

When analyzing institutions such as auctions or one-on-one bargaining, the experimental literature is particularly large (see also Ref. [16], p. 8 ff). Many experimental observations of the outcomes of various types of auctions examine game-theoretical hypothesis such as the revenue equivalence theorem (see Ref. [17] for a survey of auction experiments). Others, like McCabe et al. [27] compare the properties of several new market institutions whose theoretical properties are as yet poorly understood.

Laboratory results are joint outcomes of the characteristics of individual subjects, the laboratory institution and the environment. Some institutions, such as the double auction, powerfully influence individual behavior so that the final outcomes are relatively insensitive to the characteristics and behavior of individuals (see Ref. [16], p. 57). We believe that the final outcomes of multi-attribute auctions are much more sensitive to the subject's personal characteristics and the environment (e.g., user interface design issues, decision support for bidders, etc.). We therefore believe that the results of experimental studies will provide valuable results for the design of efficient and stable multi-attribute auction markets.

1.2. Statement of issues and overview of the paper

Multi-attribute auctions pose several practical and theoretical problems. On the practical side, we want to learn about the product and the market characteristics that make multi-attribute auctions an applicable market mechanism. In contrast to sales auctions, the bids submitted in tenders often comprise heterogeneous goods or services and the bid taker has the burden to select the "best" bid. When a bid taker procures a standardized item or she can define all the features of the product at the outset, conventional single-attribute auctions are a powerful way to automating the negotiation, because price is essentially the only negotiable attribute. The negotiation situations we investigate in our experiments describe heterogeneous monopsonies with a single buyer and multiple sellers where the traded goods have several negotiable attributes and the buyer has certain preferences about these attributes. For the implementation process, we had to solve several conceptual issues such as "how can we elicit a buyer's utility function" or "how can we implement an incentive compatible multi-attribute auction?". These issues will be addressed in Section 3.2.

As well as the practical aspects, there are numerous unsolved theoretical questions associated with multi-attribute auctions. In this paper, we will describe an economic experiment and attempt to answer some of these questions. First, we want to learn more about the equilibrium values achieved in multi-attribute auctions and compare the outcomes to conventional single-attribute auctions. From the data we have gathered, we have learned that even with two negotiable attributes, on average, multi-attribute auctions delivered higher utility (on average 4.27%) for the buyer than single-attribute auctions in all of the three auction schemes tested. We also measured the efficiency and found that in 74.47% of the multi-attribute auctions the bidder with the highest valuation won. This is slightly lower than the efficiency measured in single-attribute auctions (79.17%) which is also a consequence of the increased complexity of bidding in a multi-attribute auction. In our experiment, we have not found any evidence for the revenue equivalence theorem and could not show any strategic equivalence of English and second-score auctions.

In Section 2, we will begin with a survey of the relevant literature on auctions and the economic experiments conducted so far. The section will conclude with a description of multi-attribute auctions and a set of questions we want to answer. Section 3 will describe the laboratory procedures, namely the experimental design, the subject pool, the reward mechanism and the software we have used. In contrast to existing literature in this field, we have not deployed the mechanism for governmental or corporate procurement but for the trading with financial derivatives. Thus, the section will also give an overview of trading carried out with over-the-counter (OTC) derivatives which has been used as a scenario and will describe details of the software we have used for our experiments. Section 4 will analyze the results of the experimental sessions we have conducted and finally Section 5 will conclude with a summary of the findings and an outlook on our future research.

2. Relevant theory

In the following sections, a short introduction to auction theory will be given as well as an overview of the huge amount of economic experiments conducted in this field including a summary of multi-attribute auction theory. This should give the reader a general understanding of the relevant questions in auction design. We will conclude with a description of the analyzed mechanisms and a set of questions we have addressed in this context.

2.1. Basics of auction theory

Auction theory is a complex economic subject and can only be briefly discussed here. We shall focus on conventional single-sided auctions. A detailed characterization of double auctions or multi-unit auction theory is omitted. For a more rigorous discussion, see the seminal paper by Vickrey [37], introductions and overviews by Milgrom [30], McAfee and McMillan [26], a more theoretical treatment by Wolfstetter [40] and a strategic analysis with a game-theoretical perspective by Wilson [38].

Auctions have been defined as “a market institution with an explicit set of rules determining resource allocation and prices on the basis of bids from the market participants” [26]. There is a solid theoretical foundation for single-sided auctions. In a single-sided auction, a bid taker offers an object to two or more potential bidders who send bids indicating willingness to pay for the object [30]. Oral or open-cry auctions reveal price quotes and require public and adjustable bids. After a certain time elapse the auction clears, meaning it matches buyers and sellers and determines the price. In the case of an *English auction*, the winner is the remaining participant bidding the highest price. In a *Dutch auction*, the price at which an item is offered for sale starts from a high level and declines steadily until one of the buyers stops the clock and buys the item at that price. Sealed-bid auctions do not reveal price quotes and require private, committed bids. The highest bidder acquires the object and pays the seller her own bid price in a *first-price sealed-bid auction* and pays the second highest bid price in a *second-price* or *Vickrey auction*. There have also been several authors who have investigated these single-sided auctions in the context of procurement and sourcing. We will analyze procurement auctions in more detail in Section 2.3.

The most thoroughly researched auction model is the *symmetric independent private values (SIPV)* model. In this model, all bidders are symmetric/indistinguishable and all bidders have a private valuation for the item which is independent and identically distributed. The bidders are risk neutral and so is the seller. Under these assumptions, the bidders' behavior can be modeled as a non-cooperative game under incomplete information. In keeping with standard terminology, we define the seller's reservation price as the lowest acceptable sale price for the item. The buyer's reservation price (or *valuation*) is the maximum price she is willing to pay for the item. It is interesting to see whether the auctions achieve the same equilibrium price or if we can rank the different auction formats in any order. The surprising outcome of the SIPV model is that with risk neutral bidders, all four auction formats are payoff equivalent. This is also known as the *revenue equivalence theorem*.

There are several cases that violate the prerequisites of revenue equivalence and thus fail to be payoff equivalent. For example, if one assumes risk aversion of bidders, first-price and Dutch auctions generate greater expected revenue than English or Vickrey auction (see Ref. [26], p. 719). Risk aversion does not affect the equilibrium strategy under the English auction. Revenue equivalence also breaks down if one removes the independence assumption or the symmetry assumption of the SIPV model. In all these cases, the outcome of the English auction is relatively stable (see Ref. [40], p. 372 ff for a detailed analysis of the SIPV model).

Another approach to analyzing auctions is the *common value model*. It states that the customer's valuation of the item depends (additionally) on at least one common objective variable, such as resale value or amount of oil in the tract in the case of oil lease auctions. The common variable introduces statistical dependence to bidders' valuations of the object which, in turn, allows a bidder to infer information from other bidders' bids. Many auctions involve some common value element where the value of the item is not known or unsure during the auction. A frequently observed phenomenon in these auctions is the so-called winner's curse where the winner bids more than the good's true value and suffers a loss. The main lesson learned from the common value model is that bidders should shade their bids, as the auction always selects the bidder who has received the most optimistic estimate of the item's value as the winner.

Over the past number of years, much research has been done to extend the framework of auctions, such as combinatorial auctions or multi-stage auctions. For example, combinatorial auctions are an approach to achieving efficient allocations in the cases where bidders place bids on combinations of goods. These bids allow us to express dependencies and complementarities between goods [31,32]. Multi-stage auctions describe an auction process that is divided into several stages [13].

2.2. Experimental results

Experimental economists conducted a considerable number of auction experiments testing the results of auction theory described in the previous section. Kagel and Roth [19] give a good overview of these experiments. A large part of the conducted experiments concentrates on the revenue equivalence theorem. Some test the strategic equivalence of first-price and Dutch auctions and of second-price and English auctions, a precondition for revenue equivalence. In many auction experiments, subjects do not behave in strategically equivalent ways in first-price and Dutch auctions or in English and second-price auctions. Some experiments report higher prices in first-price compared to Dutch auctions, with these higher prices holding across auctions with different numbers of bidders. Kagel et al. [18] also report failures of strategic equivalence in second-price and English auctions. Bidding above the dominant strategy in second-price auctions is relatively widespread whereas in English auctions market prices rapidly converge to the dominant strategy price.

Efficiency in private value auctions can be measured in terms of the percentage of auctions where the high value holder wins the item. First-price and Dutch auctions varied systematically with respect to efficiency, ranging from 88% of the first-price sealed-bid auctions to 80% of the Dutch auctions being efficient [9]. The frequency of efficient outcomes in second-price auctions was shown to be quite comparable to first-price auctions. Theory also predicts that if bidders have constant or decreasing absolute risk aversion and there is uncertainty about the number of bidders this uncertainty raises revenue on average, compared to revealing information about the number of bidders. Behavior changed in most experiments in the way theory predicts. Moreover, experiments showed that an increased number of bidders almost always results in higher bidding in first-price auctions.

2.3. Multi-attribute auctions

In this section, we introduce multi-attribute auctions, a special case of procurement auctions. Several authors investigated tenders and procurement auctions (see Ref. 11,37]). These auctions are mostly deployed in monopsony situations such as governmental or corporate procurement where a bid taker auctions off goods or services she wants to buy. Laffont and Tirole [25] describe many of the issues involved in procurement negotiations ranging from the costs of setting up a tender to evaluating the bids in such a process. They also mention the need that auction theory must be generalized to "multidimensional bidding". In procurement auctions, bidders often provide very different kinds of goods and services in their bids. A good example is the procurement of large food retailers [5]. The suppliers in the market consist of large companies as well as a large number of small- and medium-sized enterprises such as bakeries and breweries. The buyers are a small number of large food retailers who aggregate the demand and distribute it to the end consumer. Purchasing managers have their own preferences for product quality, price, terms of payment and delivery and they are looking for the offer that best satisfies these preferences. The overall utility of a deal for the buyer not only contains the price of the item but a combination of the different attributes.

Conventional procurement auctions only automate negotiation on the price. It would be comfortable to have a mechanism that takes multiple attributes of a deal into account when allocating it to a bidder. In other words, the mechanism should automate multilateral negotiations on multiple attributes of a deal (some authors call this also a multidimensional auction, however, this term is also used in other contexts). Only very little theoretical work has been done in this field up to now (Koppius [23] gives an overview of the existing approaches). Cripps and Ireland [10] introduced a model where bid takers set threshold levels for quality and price and analyze three evaluation schemes. First, they assume a scheme where price bids are only accepted after quality plans have been submitted and approved. In a second scheme, they consider a price auction first and then quality plans are submitted in the order of the level of price bids. The first plan to qualify on quality is then accepted. Finally, they investigate joint submission of a price and a quality plan where the allocation is made to the highest priced plan satisfying the quality threshold. The quality test is conducted by the buyer under the assumption that objectively better projects have a greater probability of being accepted. The schemes essentially produce the same results. The paper, however, leaves open a complete characterization of optimal mechanisms.

A thorough analysis of the design of multi-attribute auctions has been provided by Che [7]. Che studied design competition in government procurement by a model of two-dimensional auctions where firms bid on price and quality. He focuses on an optimal mechanism in cases where bids are evaluated by a scoring rule designed by the procurer. Each bid contains a quality, q, and a price, p, and quantity in this model is normalized to one. The buyer in this model derives a utility from a contract comprising q and p:

$$U(q, p) = V(q) - p, \tag{1}$$

where V is the individual utility function of quality. On the other hand, a winning firm earns profits from a contract (q, p):

$$\pi_i(q, p) = p - c(q, \theta_i) \tag{2}$$

In the cost function c the unit cost is expressed as θ which is private information. θ is assumed to be independently and identically distributed. Losing firms earn zero profits and trade always takes place, even with a very high θ. Che investigates what an optimal auction in this case should look like. In Che's model, an optimal multi-attribute auction selects the firm with the lowest θ. The winning firm is induced to choose quality q which maximizes $V(q)$ considering the costs.

Che considers three auction rules: In a so-called "first-score" auction — a simple generalization of the first-price auction, each firm submits a sealed bid and, upon winning, produces the offered quality at the offered price. In other auction rules, labeled "second-score" and "second-preferred-offer" auctions, the winner is required to match the highest rejected score in the contract. The second-score auction differs from the second-preferred-offer auction in that the latter requires the winner to match the exact quality–price combination of the highest rejected bid while the former has no such constraint. A contract is awarded to the firm whose bid achieves the highest score in a scoring rule $S = S(q, p)$.

In the model, it can be shown that the equilibrium in the first-score auction is reduced to the equilibrium in the first-price auction if the quality is fixed. The Vickrey auction intuition also applies to the second-score auction: If a firm with type θ bids a higher score than the one based on θ, it would risk winning at negative profits without increasing its profit conditional on winning. If the firm bids a lower score, it would forgo some opportunity of winning at positive profits. Similarly, one can show that in a second-preferred-offer auction, each firm will bid a score that will earn the firm zero profit. However, in the second-preferred-offer auction, a winning firm has no control over the quality in the final contract.

Another important question analyzed by Che [7] tries to discover the *optimal scoring rule* for the buyer. He showed that if the scoring function under-rewards quality compared to her utility function, first- and second-score auctions implement an *optimal mechanism*. This is true, because the true utility function fails to internalize the informational costs associated with increasing quality. Che also shows that if the buyer's scoring function reflects the buyer's preference ordering, i.e., equals her utility function, all three auction schemes yield the same expected utility to the buyer — a *two-dimensional extension of the revenue equivalence theorem*.

The costs in Che's model are assumed to be independent across firms. In the context of procurement auctions, one might expect the costs of several bidders not to be independent. Branco [6] derives an optimal auction mechanism for the case when the bidding firms' costs are correlated, but the initial information of firms is independent. This is somehow equivalent to the common value approach in classic auction theory. He shows that when the quality of the item is an issue, the existence of correlations among the costs has significant effects on the design of optimal multi-attribute auctions. Under these conditions, the multi-attribute auctions analyzed by Che are not optimal. Branco then discusses the implementation of the optimal outcome through multi-attribute mechanisms. In his model, the optimal quality is a function of the firm's efficiency, which depends on parameters not known to the firm at the time of the auction. This is true, because in the correlated-costs model, the optimal quality to be provided is a function of all the bidders' parameters. Therefore, unlike in Che's independent-cost model, optimal quality cannot be achieved just through the bidding process. As a result, the procurer has to use a two-stage mechanism: a first-score or second-score auction, followed by a stage of bargaining over quality between the procurer and the winner of the first stage. Branco shows that the two-stage first-score auction and the two-stage second-score auction implement the optimal mechanism.

2.4. Simulation of multi-attribute auctions

The difficulty of multi-attribute auctions is the variety of different scoring functions and parameter settings one can deploy. This is a reason why the basic assumptions of game-theoretical models are kept relatively simple. The models describe two-dimensional negotiations (price and quality) and also the bidders' behavior is modeled in a rather simple way. Nevertheless, the analytic complexity of these models poses tight constraints for the modeler. Computer simulations allow exploring the outcomes of more complex behavior. Bichler and Klimesch [4] developed a general simulation model of multi-attribute auctions. In this model, they assume a generic good, which can be described by its price and a certain number of qualitative attributes. The basic question they try to tackle is whether multi-attribute auctions achieve higher utility than conventional single-attribute auctions. In the buyers' scoring function, they use weights in order to express the importance a buyer puts on the various attributes. Bichler and Klimesch use different weights and a different number of attributes in their simulation.

They show that in cases where all attributes are of equal importance (the weights are uniformly distributed), negotiating on many attributes achieves a higher utility than negotiating only on price (and fixing all other attribute values in advance). That is, if the scoring function correctly mirrors the buyer's utility function, she can expect to be better off at the end. If the buyer's scoring function puts a very high emphasis (i.e., weight) on the price and only very little weight on all other attributes, multi-attribute and single-attribute auctions tend to achieve a very similar outcome. The simulation is sensitive to changes in the basic assumptions such as the initial distributions of attribute values or the cost parameters. However, in all settings there is a positive correlation between the achieved utility scores and the number of negotiable attributes in the auction.

Multi-attribute auctions are more difficult to analyze analytically than conventional auctions. On the one hand, the user's behavior is less determined than, for example, in a Vickrey auction. On the other hand, the multitude of attributes and parameter settings (e.g., weights) makes it hard to define a comprehensive analytical model. In such an environment, economic experiments are more important not only in terms of testing the theory but in terms of development testing, i.e., improvement of the proposed mechanisms and negotiation protocols. In the following section, a detailed description of the mechanisms which were analyzed in the experiments will be given.

2.5. Description of the analyzed mechanisms

We have taken a heuristic approach and have proposed a set of multi-attribute auction mechanisms. In contrast to previous game-theoretical analyses, we do not only analyze first-score and second-score sealed-bid auctions, but also multi-attribute open-cry auctions (e.g., a generalization of the English auction). All three auction schemes follow the same principle. The buyer first has to define her preferences for a certain product in the form of a scoring function (negotiation literature [15] also recommends defining preferences and valuation before every negotiation). The buyer has to reveal this scoring function to suppliers whereas the suppliers do not have to disclose their private values. Then the mechanism designates the contract to the supplier who maximizes the buyer's preferences, i.e., who provides the highest overall score for the buyer.

A basic question in this context is how a buyer determines the scoring function. This is relatively easy in the case of two attributes, but it gets much more complicated if a buyer has to determine the trade-off between three or more attributes. We assume that the scoring is based on the buyer's utility function. Micro-economists were the first to analyze consumers' preferences and utilities. Later researchers in the field of operations research and business administration attempted to utilize utility theory in order to actively make decisions, We therefore use the concepts of classic utility theory and decision analysis in order to determine the buyer's utility function.

A basic question is whether preferences can be mapped into coherent utility functions. In particular, the problems of "expected" utility theory have a long and controversial history [29]. A crucial step was undertaken in a fundamental paper of Wold [39]. He has shown that if preferences are complete, reflexive, transitive, continuous, and weakly monotonic, there then exists a continuous utility function U: $R_+^k \to R$ which represents those preferences. Further progress has been made by Debreu [12], who also emphasizes the distinction between preference and utility. Decision analysis can be seen as an approach to making utility theory operational in actively making decisions in a multi-criteria environment. The research carried out in decision analysis has led to a considerable amount of literature on understanding and improving decision making of individuals, groups and organizations. It is generally considered a branch of the engineering discipline of operations research but also has links to economics, mathematics and psychology.

In this context, we concentrate on multi-objective decisions under certainty which is widely used in business and government decision-making. Here, the basic assumptions about preferences and utility are relaxed and therefore less controversial. Multi-objective decision analysis in general prescribes theories for quantitatively analyzing decisions involving multiple, interdependent objectives. The essence of multi-objective decision analysis is to break decisions down into small pieces that a user can deal with individually and then recombine logically. For the evaluation of offers, the price of a product will be important, but so could its delivery time, quality, and payment conditions. It is important to determine the relationship and trade-off between them. Decision analysis techniques such as the Multi-attribute Utility Theory (MAUT) [8,20], the Analytic Hierarchy Process (AHP) [33] and conjoint analysis [1] are used in a broad range of software packages for decision-makers. They can also be used to determine the utility function of a buyer (see also Refs. [2,21]). These ordinal utility functions are only required to rank-order certain outcomes in a way that is consistent with the decision-maker's preferences for those outcomes. The model used in MAUT, AHP and conjoint analysis is basically an additive utility function where some kind of subjective judgment forms the basis of the weights, and yet the interpretation of the weights is not always clear. Therefore, extreme care must be exercised in making the judgments on which the additive utility function is based.

A precondition for the deployment of an additive utility function under certainty is the existence of mutual *preferential independence* between the attributes. An attribute x^i is said to be preferentially independent of x^j if preferences for specific outcomes of x^i do not depend on the level of attribute x^j. It is fair to say that preferential independence holds for many situations and in the following we will concentrate on cases with preferential independence. Keeney and Raiffa [20] provide a more detailed discussion of multi-attribute utility modeling.

Based on these preconditions, we can provide a more formal description of the mechanisms we analyze. A bid received by the auctioneer can be described as a vector $\boldsymbol{Q}$ of n relevant attributes indexed by i. We have a set B of bids and index the m bids by j. A vector $\boldsymbol{x}_j = (x_j^1 \dots x_j^n)$ can be specified where x_j^i is the level of attribute i in bid b_j. In the case of an additive scoring function $S(x_j)$, the buyer evaluates each relevant attribute x_j^i through a scoring function $S_i(x_j^i)$. Suppose that the additive scoring function corresponds to the buyer's true utility function $U(x_j)$. Then the individual scoring function, S: $Q \to R$, translates the value of an attribute into "utility units". The overall utility $S(x_j)$ for a bid b_j is then given by the sum of all individual scorings of the attributes. For a bid b_j that has values $x_j^1 \dots x_j^{n''}$ and weights $w_1 \dots w_n$ on the n relevant attributes, the overall utility for a bid is then given by:

$$S(x_j) = \sum_{i=1}^{n} w_i S_i\left(x_j^i\right) \tag{3}$$

A reasonable objective in allocating the deal to the suppliers is to allocate them in a way that maximizes the utility for the buyer, i.e., to the supplier providing the bid with the highest overall utility score for the buyer. The function max $S(x_j)$ (with $1 \leq j \leq m$) gives us the utility score of the winning bid and can be determined through various auction schemes. In a *first-score sealed-bid auction*, the winner gets a contract awarded containing the attributes x_j of the winning bid. Alternatives with the same overall utility are indifferent. In these cases, the first bid is the winning bid. The *multi-attribute English auction* (also first-score open-cry auction) works in the same way, however, all bids are made available to the participants during an auction period. In a *second-score sealed-bid auction*, we take the overall utility achieved by the second highest bid $S_{\max-1}$ and transform the gap to the highest overall utility $(S_{\max} - S_{\max-1})$ into implied volatility. Consequently, the winning bidder can charge a higher option price. This is similar to the procedure described as second-score auction by Che [7]. In the first-score and second-score sealed-bid schemes, the auction closes after a certain pre-announced deadline. In a multi-attribute English auction, bids are made public and the auction closes after a certain time elapse in which nobody submits a bid. For several practical reasons, we have not implemented a multi-attribute generalization of the Dutch auction.

2.6. Research questions

We implemented the three multi-attribute auction schemes as a service of an electronic brokerage system on the Internet. Section 3.2 describes several details of this implementation such as the elicitation of the users' preferences or the calculation of the winner. We used the electronic brokerage system to get empirical results to a number of questions:

• *Are the equilibrium values achieved in a multi-attribute auction higher compared to single-attribute auctions with respect to the underlying scoring function of the bid taker*?

A basic question of auction design is which auction format maximizes the bid takers profit. In a multi-attribute auction, the bidder has several possibilities to improve the value of a bid for the bid taker, sometimes even without increasing her costs and thereby creating joint gains for all parties. In corporate procurement, for example, a bidder might have more possibilities for bringing in her own strength in terms of quality, delivery time or payment conditions.

• *Are multi-attribute and single-attribute auctions efficient*?

Pareto efficiency can be measured in terms of the percentage of auctions where the high value holder wins the item. We want to learn about the efficiency of multi-attribute auctions compared to single-attribute auctions.

• *Are all multi-attribute auction formats payoff equivalent*?

Experimental tests of single-attribute auctions could not prove the revenue equivalence theorem or strategic equivalence of the English and Vickrey auction. We will investigate the revenue equivalence of multi-attribute English, first score, and second-score auctions.

Besides these basic questions, there are several other issues which we did not tackle in our first set of experiments. For example, we did not test the impact of bidder experience on the outcome of multi-attribute auctions. We feel that due to the inherent complexity of bidding on multiple attributes, there is a correlation between bidder experience and outcome, however, we also think that decision support on the bidders' side can help a bidder in finding her "best" bid. Other questions would be how the number of bidders influence the equilibrium values achieved in various multi-attribute auction schemes and how different types of information revelation can influence the outcome. Bidders in a multi-attribute English auction can be informed about the overall utility of a rival's bid or they can be informed about the detailed attribute combination a rival has bid. We leave these questions open for future experiments.

3. Laboratory procedures

In this section, we will describe a set of experiments which we have conducted to find answers to our questions. We also wanted to learn about the applicability of the mechanisms in "real-world situations". Bidding on multiple attributes can be quite complex and we have been interested in the participants' ability to bid reasonably on multiple attributes and their trust in the mechanism. The experiments we have conducted were based on a real-world scenario. On the one hand, this should demonstrate the applicability of the mechanism in this domain. On the other hand, it is necessary for the bidder to know about dependencies among the different attributes and the context of a negotiation in order to provide reasonable bids.

3.1. Basic environment and scenario

The literature on multi-attribute auctions mostly investigates procurement scenarios in a governmental or corporate environment. We have been using the trading with non-standardized financial products as a scenario for our experiment, in particular, we have focused on OTC financial derivatives. In the OTC market for options there are several well-known attributes (premium, strike price, duration, etc.) which are negotiated. As participants have different market expectations and different risk attitudes, they value combinations of attribute values differently. In the following we will give a short introduction to OTC derivatives trading. This should provide the necessary background and depict the scenario (see also Ref. [28]).

An option is the right to buy (call) or sell (put) an underlying instrument at a fixed point in time at a strike price. It is bought by paying the option premium/price upon conclusion of the contract and restricts the risk of the buyer to this premium. For example, the holder of a call purchases from the seller (writer) of the call the right to demand delivery of the underlying contract at the agreed price any time upon (American style) or exactly at (European style) expiration of the option contract. The strategies of market participants as well as their valuations depend on the investor's market expectations, the investor's objective and risk tolerance and the chosen market. Whereas standardized options are traded on an exchange, OTC options are traded off-floor. Today, there are standardized options and futures for shares, stock indices, bonds, currencies, wheat and tulip bulbs, to name just a few. In general, option contracts are based on a number of preset terms and criteria:

- Type of option (call or put)
- Style (American or European)
- Underlying instrument and price
- Contract size or number of underlying instruments
- Maturity
- Strike price

All these criteria influence the option premium no matter whether the options are traded on an exchange or over-the-counter. For example, every change in the price of the underlying is reflected by a change in the option premium. The measure reflecting the size and frequency of fluctuations in the price of the underlying item over a specific period of time is called volatility. If the market anticipates great volatility of the underlying item, the option premium will be high. In order to set a certain option premium in the context of its strike price and other parameters, traders often use the so-called *implied volatility*, which indicates the volatility implied by a certain market price. Thus, the value of a certain price can be measured independent of the strike price. The lower the implied volatility, the better it is for the buyer of the call. For many traders it has become common practice to quote an option's market price in terms of implied volatility [24].

On an exchange, all of these attributes are specified in advance and the only negotiable attribute is the price. This makes trading much easier, however, it reduces the number of derivatives traded to a small set of possible products. As a result, over the past few years, the volume of OTC contracts has been growing enormously. Trading OTC options is not bound to an organizational structure in that supply and demand are concentrated on a centralized trading floor. Potential buyers of OTC options bargain with a number of investment brokers or banks on attributes such as the strike price, the style, the maturity and the premium of an option. Terms and conditions are usually not determined by auction, but by way of bargaining. Financial engineers created a whole bunch of different financial OTC products tailored for specific purposes, ranging from plain vanilla options where all important attributes are negotiated during the bargaining process, to exotic derivatives with certain predefined properties (see Ref. [22] for different types of options and details of option pricing). The volume of OTC transactions has increased enormously over the past decade. In particular, institutional investors have a requirement for special derivative products.

Traditional OTC transactions, however, have several disadvantages. Bilateral negotiations with banks or investment brokers are conducted over the phone, leading to high transaction costs for a deal. In contrast to electronic exchanges, investors lose their anonymity and also have to bear the contracting risk. New approaches are trying to establish efficient, electronic trading systems for OTC derivatives. For example, in 1993 U.S. options exchanges began with the development of Flexible EXchange® Options (FLEX Options), an electronic trading system for the trading of index options. Equity FLEX Options (E-FLEX Options) has broadened the concept to encompass listed equity options. Equity FLEXible Exchange Options provide the opportunity to customize key contract terms, including the expiration date, exercise style and strike price of an equity option. Then prices are determined anonymously, using single-sided auction mechanisms. Thus, options and futures can be designed to fit an investor's own investment strategies and goals. Both systems have been designed to extend investor access to customized derivative products. They are used on the NASDAQ, the American Stock Exchange (AMEX), the Chicago Board Options Exchange (CBOE) and the Pacific Stock Exchange (PSE). While systems like this combine many advantages of OTC trading and electronic exchanges, the deployed auction mechanisms only automate negotiations on the price. In contrast, participants on an OTC market have the possibility to also bargain on several attributes of a contract such as strike price, style, contract maturity or contract size. This gives a participant many more degrees of freedom during the negotiation and has the potential to achieve a better deal for both parties.

In our experiments, a buyer has tried to buy a call on a certain index or share traded on the Vienna Stock Exchange and bidders provided bids in a number of different auction schemes. Every session comprised six different trials where we deployed three single-attribute and three multi-attribute auctions (FPSB, English, Vickrey). During an auction period, the buyer specifies her preferences in the form of a scoring function. Bidders provide bids containing values for strike price and implied volatility. The winner is then computed based on these bids and the rules of the deployed auction scheme. In the single-attribute trials, bidders bid only the implied volatility as a substitute for the option premium. In multi-attribute trials, the buyer additionally provides a list of strike prices the bidder could choose from and bidders can consequently provide bids on both attributes. We have chosen implied volatility instead of the option premium, as strike price and implied volatility are mutually independent preferences — a prerequisite for an additive utility function (see Section 2.5). In some attributes, we also tested bidding on three attributes, namely strike price, implied volatility and style. Section 3.2 describes details of the institution we deployed and of the software tools we have used to conduct the experiments.

3.2. Institution

For our experiment, we implemented an electronic brokerage service on the Internet. The electronic broker provides several single- and multi-attribute auction mechanisms as well as a buyer client and a bidder client in the form of a Java applet.

3.2.1. The electronic broker

The server-side application logic, i.e., registration of participants, forwarding of requests for bids, determination of winners, etc., is implemented in Perl and PL/SQL on top of an Oracle8 database. Messages between the buyer client and server such as bids and requests for bids are exchanged in a predefined XML format using HTTP. There are numerous ways to implementing client/server applications of this type, based on middleware such as OMG/CORBA or Java RMI (see Ref. [34] for an overview). Fan et al. [14] give a good example of an Internet marketplace based on distributed object infrastructures. We have chosen this approach because of its simplicity and flexibility and the fact that we can build on the ubiquitous Web infrastructure. The client-side runs in most Web browsers, the server-side implementation does not require additional infrastructure such as an Object Request Broker, and can be implemented using a variety of programming languages. Therefore, it is possible to conduct the experiments in different computer labs without the installation of additional client-side software. All the transaction data is stored in the relational database and can be evaluated using standard spreadsheet and statistics packages. To our knowledge, this is the first implementation of multi-attribute auctions on the Internet (see Ref. [3] for a detailed description of the implementation). In addition to the bidder client in our implementation, we provided bidders with a decision aid in the form of an MS Excel spreadsheet helping them to determine values for strike price and implied volatility based on their market expectations and risk attitude. The electronic brokerage service provides a very general framework and can easily be customized for different types of products.

3.2.2. The buyer client

During an auction period, a buyer specifies her utility function using the Java applet which can be downloaded over the Web (see Fig. 2). Eliciting the buyers' preferences is one of the key problems that need to be addressed by the graphical user interface of the applet. The need to get "true" data in an appropriate form from the user poses both psychological and theoretical problems in creating a suitable interface. In particular, the buyer may not have thought about what her fundamental preferences are in sufficient detail and must therefore be prompted with a suitable set of questions. We then need to map the buyer's preferences, as input by the applet, into coherent utility functions. In our current implementation, we use MAUT and deploy an additive utility function. An additive utility function is both easy to implement and intuitive to the user. The additivity assumption implies that attributes are preferentially independent and there are no interaction effects. It depends largely on the type of item traded whether this requirement is satisfied. In our case, we assume preferential independence of the strike price as well as implied volatility and style of an OTC derivative.

The additive utility function is composed of two different kinds of elements, scores on individual attribute scales (also called "individual utility functions") and weights for these attributes. For the evaluation of bids, the implied volatility is important, but so is its strike price. It is important to determine the relationship and trade-off between them. The assessment of the individual utility functions and weights is a core issue, when using MAUT. The buyer evaluates each relevant attribute value of a bid x_j^i through an individual utility function $S_i(x_j^i)$ and indicates its relative importance value by a weight w_i (see Section 2.5). All the weights are positive and add up to 1. A reasonable way to determine the weights for the various attributes is to determine the marginal rate of substitution between one particular attribute and any other attribute. Several other methods are described in Ref. [8].

Individual utility functions provide a means to measuring the accomplishment of the fundamental objectives. Some scoring functions are easily defined. If minimizing implied volatility is an objective, then we can, for example, define a linear function of the implied volatility that assigns values between 1 and 0, with some highest acceptable value for the implied volatility providing zero utility and zero implied volatility providing a utility of 1. We call this type of proportional scores a "continuous attribute". The individual utility $S(x_j^i)$ of a continuous attribute can then be computed by:

$$S\left(x_j^i\right) = \frac{x_j^i - \text{worst value}}{\text{best value} - \text{worst value}} \tag{4}$$

Another way to assess utilities, particularly appropriate for attributes that are not naturally quantitative, is to assess them on the basis of some ratio comparison. We call this type of attribute "discrete attribute". Suppose that the buyer in our scenario provides four different strike prices to choose from (1300, 1320, 1340, 1360) for the bidder. Clearly, this is not something that is readily measurable on a meaningful proportional scale. Using a ratio scale, the buyer might conclude that 1320 is twice as good as a strike price of 1360 and so on. In the applet the buyer assigns some number of points between 0 and 100 to each possible alternative. In this way, for example, the buyer might assign 94 points to 1300, 86 points to 1320, 70 points to 1340 and 44 points to a strike price of 1360 (Fig. 1).

Now, we scale these assessments so that they range from 0 for the worst alternative to 1 for the best alternative. This can be achieved by solving two equations simultaneously to compute the constants a and b:

$$0 = a + b(44)$$

$$1 = a + b(94)$$

This results in utilities of 1 for 1300, 0.84 for 1320, 0.52 for 1340 and 0 for 1360. Using these two types of individual utility functions, the overall utility for a bid is given by the sum of all weighted utilities of the attribute values. For a bid that has values $x_j^1 \dots x_j^n$ on the n attributes, the overall utility for a bid is again given by Eq. (3). The bid with the highest overall score is the most desirable under this rule. Similar procedures are used for procurement decisions or to choose between different products in electronic catalogs [36].

Fig. 2 shows a screenshot of the Java applet we use in our implementation on the buyer side. The user interface consists of several areas. In the upper left area, the buyer supplies a unique identifier, which she gets upon registration through a WWW form. In the text field above, the buyer can specify a parameter file for a certain product traded on the marketplace. Thus, it is very easy to adapt the applet to any other kind of product by simply changing the

Fig. 1. Attribute scales.

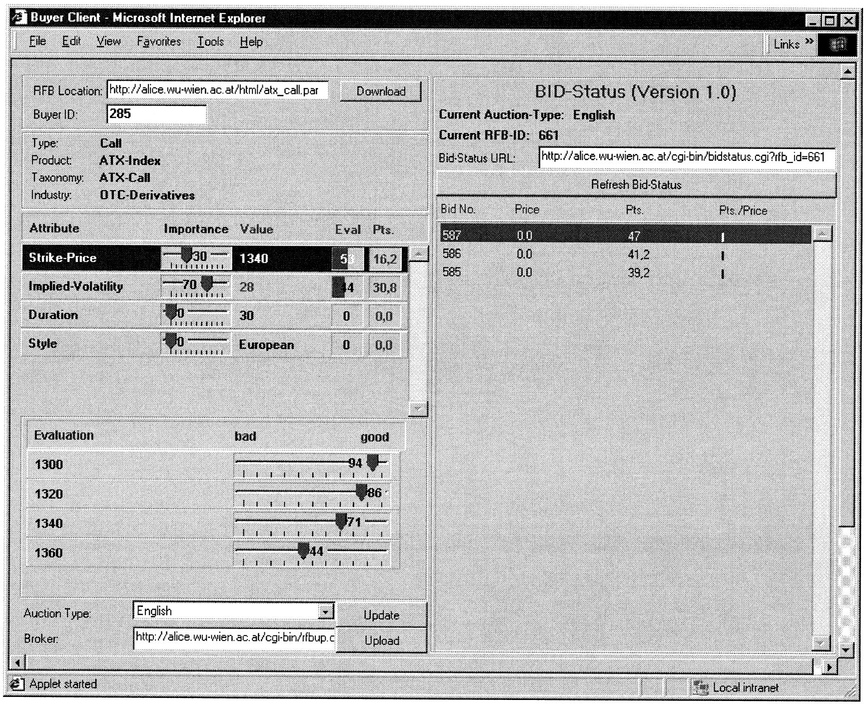

Fig. 2. Buyer client.

parameter file. In this case, we trade calls on the Austrian Trading Index (ATX). Below we find a list of relevant attributes for the auction. The negotiable attributes in this case are the strike price and the implied volatility. "Duration", i.e., maturity and "Style" are fixed in advance. In the lower left panel, users can define the individual utility functions for the negotiable attributes, which can be either continuous or discrete functions as described above. The utility function of the strike price (as shown in the screenshot) is determined in a discrete form. From the input of the buyer, the applet compiles a Request for Bids (RFB) in XML format and sends the RFB via HTTP to the electronic brokerage service. The RFB contains the bidder ID, the product description and the parameters of the additive utility function. The brokerage service parses the RFB, retains all the relevant data in the database and informs potential bidders via e-mail.

After the auction has begun, the buyer can query a list of bids submitted on the right-hand side of the applet, ranked by overall utility (third column). By clicking on a certain bid, the buyer can see the details of every bid in the form of green numbers on the left-hand side of the applet.

3.2.3. The bidder client

Bidders download the RFB from the URL they have received via e-mail to their bidder client (see Fig. 3). This Java applet allows parameters for all negotiable attributes to be entered and to upload an XML-formatted bid via HTTP to the brokerage service. Bidders also have to register via a Web form, in order to get a bidder ID. The applet shows important parameters contained in the RFB and allows values for the negotiable attributes to be entered. In the case of a discrete attribute (e.g., the strike price), the bidder can select a value from a drop down listbox. The numbers in brackets give information about the utility of each value. In the case of continuous attributes, the bidder can enter a number in a text field. The numbers must be within a certain range depicted right beside the text field. " + –" means that the individual utility function is down-

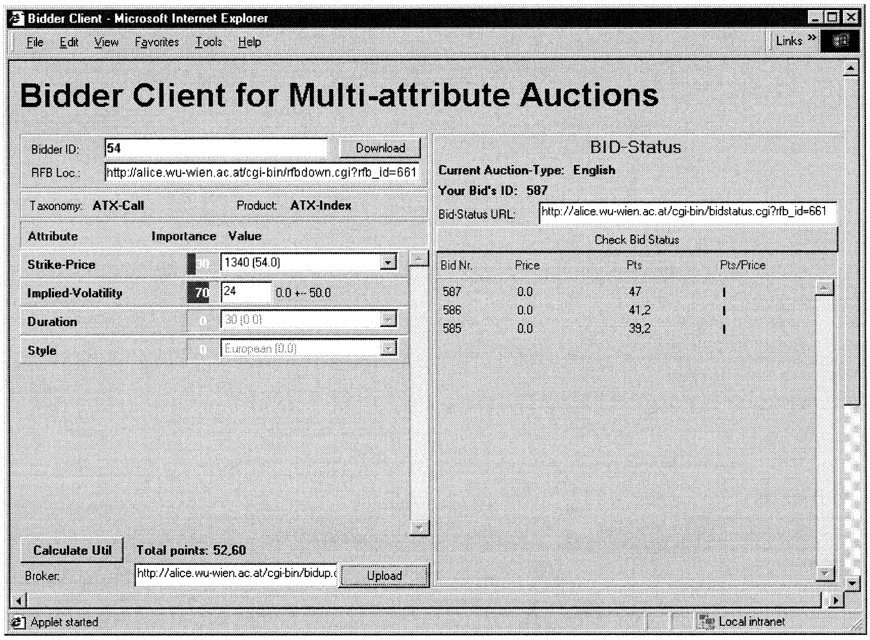

Fig. 3. Bidder client.

ward sloping or in other words, the lower the implied volatility is in a bid, the more utility points are achieved. The "Calculate Util" button on the lower left corner of the applet can be used by the bidder to compute the utilities achieved with different attribute values. In the case of an open-cry auction (e.g., an English auction), the brokerage service reveals information about the bids submitted so far on the right-hand side of the applet.

Bidders in the experiment calculate the implied volatility for a certain strike price depending on the minimum profit they want to make given a certain market expectation. The calculation of implied volatility is not trivial and thus we provided the bidders with a decision aid in the form of an Excel spreadsheet. Thus, the bidders only have to enter values for their market expectation and the profit they want to make and the spreadsheet calculates implied volatility values for the various strike prices according to the Newton–Raphson method [24]. This auxiliary tool reduced the complexity of bidding for the subjects and also less experienced subjects could provide reasonable bids after one or two dry runs.

So far, we have implemented three types of auctions, namely English, Vickrey and first-price sealed-bid auction, each of which computes the winning bid in a different way. In a first-price sealed-bid auction, the winning bid is the one providing the highest overall utility. A Vickrey auction computes the winning bid in the same way, however, the winner can charge a higher price for the call. Here, we transform the gap between the utility points of the highest bid and the second highest bid into implied volatility, meaning the amount of money by which the winner can increase the premium for her call. This should lead to incentive compatibility, similar to a conventional Vickrey auction. In both types of mechanisms, the auction closes after a certain pre-announced deadline. Finally, in an English auction, bids are made public and the auction closes after a certain elapse time, in which nobody submits a bid.

3.3. Experimental design, subject pool and reward mechanism

In this section, we describe the experimental design as well as the subject pool and the reward mechanism we used in our experiment. For this purpose, we would like to introduce some basic terminology. An *experiment* is a set of observations gathered in a controlled environment. A *trial* is an indivisible unit of observation in an experiment, e.g., a single auction or a single round of market trading. Finally, a *session* is a group of trials conducted on the same day, usually with the same set of subjects.

In May and in October 1999, we conducted 16 sessions with MBA students at the Vienna University of Economics and Business Administration. The subjects were students of an introductory IS class. All subjects had a basic knowledge of economics and financial derivatives. In every session, a group of subjects conducted six different trials, namely a first-score sealed-bid auction, a second-score sealed-bid, and an English auction, all of them in their single-attribute and their multi-attribute form. Conducting single-attribute and multi-attribute trials allowed us to compare the outcomes of both types of auction formats. A basic question in this context is how many subjects should participate in a single session. Some experimental economists suggest that the number of subjects should be identical to the number observed in real-world settings. According to Friedman and Sunder [16], two or three subjects in identical situations are sufficient to attain competitive results in a laboratory experiment. In our experiment, we had four subjects bidding in every trial.

An experimental design specifies how variables are controlled within and across a block of trials. We focused in particular on the equilibrium values achieved under the different auction types. Besides the auction mechanism used in a certain trial, the equilibrium values can be influenced by several variables such as the number of participants in the auction, bidder experience, risk attitudes as well as other individual or group idiosyncrasies. Ideally, all these variables are controlled either directly or indirectly through appropriate randomization procedures that ensure independence from the focus variables. By keeping the number of subjects constant across all trials, we controlled the influence of the number of bidders on the outcome. As we had different subjects in all trials, we tried to randomize individual and group idiosyncrasies such as the individual risk attitudes or learning curves.

Before the experiment, we introduced the scenario in a 40-min lecture for all subjects. Subjects were provided with examples of valuations and bids along with profit calculations to illustrate how the auction works. Before each session (approximately 1 1/2 h), we conducted two dry runs in order to familiarize the subjects with multi-attribute bidding and the bidder applet. This was perceived to be sufficient by the participants. In all trials of a session we used the same scenario. Before a session began, we repeated the scenario and asked all participants to provide us with a list of *valuations*, i.e., a minimum implicit volatility value for each strike price. These valuations were used afterwards to analyze the efficiency and strategic equivalence of the different auction schemes. In all sessions, we kept each subject's valuations and rewards as private information not available to other subjects.

In order to give the MBA students an incentive to bid reasonably during all auction periods, we had to introduce a reward mechanism. Induced-value theory gives guidelines for establishing control of a subject's preferences in a laboratory economy through appropriate reward mechanisms (see Ref. [16], p. 12 ff). A reward mechanism should be dominant, in that it is the only significant motivation for each subject and is the determinant of her actions. The extent of similarities between laboratory and field environments that permits generalizing laboratory findings to field environments is also important. In our trials, we wanted the subjects to bid reasonably given their risk attitude and market expectations. After maturity of the option expired (after a month), we took the market prices of the Vienna Stock Exchange and computed the profits and losses for all winners of an auction. To participate in a session, every subject gained a certain amount of credit points in the final exam. We ranked the subjects by their profits and gave them additional credit points depending on their profit. If a subject made a loss she also lost part or all of the credit points for the final exam.

4. Data and results

As mentioned in Section 2.3, the experiments were centered around three different topics: Are the equilibrium values achieved in a multi-attribute auction higher compared to single-attribute auctions? Are all multi-attribute auction formats payoff equivalent? And finally, are multi-attribute and single-attribute auctions efficient? These questions will be addressed in the following subsections.

4.1. Comparison of multi-attribute and single-attribute auction formats

A crucial question is whether multi-attribute auctions achieve higher utility scores than single-attribute auctions with respect to the underlying utility function of the bid taker. We assumed that the scoring function in multi-attribute auctions correctly mirrors the buyer's utility function and used this function to calculate a utility score for the winning bid in single-attribute auctions and in multi-attribute auctions. This allowed us to compare the three auction schemes in their single-attribute and multi-attribute format. Over time, the ATX changed and different subjects faced slightly different conditions in their session. We therefore computed the utility score of the winning bid as a percentage of the highest valuation given by the participants at the beginning of each session. This allowed us to compare different trials. As we compare the payoff of different auction formats, we have used the utility scores of the second best bid in the case of a second-score auction.

In our experiment, the utility scores achieved in multi-attribute auctions were significantly above those of single-attribute auctions for groups of size $n = 4$. Using a t-test with a significance level of $\alpha = 0.05$, we had to reject the null hypothesis of revenue equivalence between single-attribute and multi-attribute auction formats, but we had to accept the hypothesis that multi-attribute auctions achieve higher utility scores than single-attribute auctions. Multi-attribute auctions in our experiment achieved, on average, 4.27% higher utility than single-attribute formats. In 72.92% of all trials, the overall utility achieved in multi-attribute auctions was higher than in single-attribute auctions. In more detail, the utility achieved in multi-attribute first-score auctions was on average 3.18% higher than in single-attribute first-price sealed-bid auctions, the utility achieved in multi-attribute second-score auctions was 4.25% higher than its single-attribute counterpart and finally, the utility achieved in multi-attribute English auctions was 5.39% higher than in the single-attribute English auction (Fig. 4).

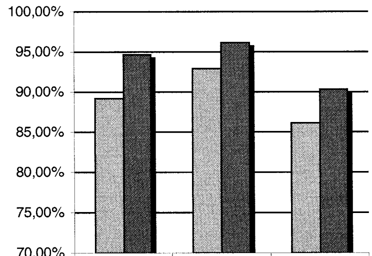

	English	First score	Second score
Single-attribute	89,25%	92,97%	86,08%
Multi-attribute	94,64%	96,15%	90,32%

Fig. 4. Difference between multi-attribute and single-attribute auctions.

One explanation for this result is that in a multi-attribute auction, a bidder has more possibilities for improving the value of a bid for the bid taker, sometimes even without increasing her own costs. We conducted most trials with two negotiable attributes, but also some test trials with three negotiable attributes. Here, the difference in achieved overall utility was even higher. However, what we have learned is that in the case of a higher number of negotiable attributes, there is a need for bidders to have more sophisticated decision support. It is simply not that easy anymore to find the combination of attributes that achieves the highest utility for the buyer, i.e., the best bid. In most cases, subjects used all negotiable attributes to improve their bid.

Besides, we had a control group with which we conducted six sessions within the period of 3 weeks. The results of this group were similar to the ones described above and we could not find evidence of the impact of bidder experience in the two-dimensional scenario we used. However, the impact of bidder experience on multi-attribute auctions has to be investigated more thoroughly.

4.2. Revenue equivalence of multi-attribute auction formats

Using a t-test ($\alpha = 0.05$), we also had to reject the null hypothesis of revenue equivalence between various pairs of multi-attribute auction formats. The utility scores achieved in multi-attribute first-score auctions were significantly higher than the ones in multi-attribute English or second-score auctions. Here

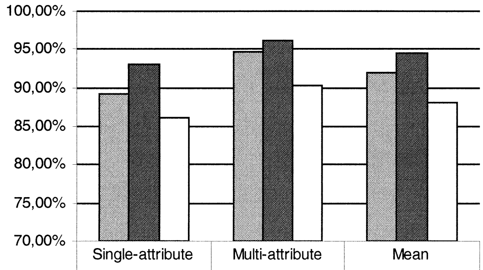

	Single-attribute	Multi-attribute	Mean
English	89.25%	94.64%	91.94%
First score	92.97%	96.15%	94.56%
Second score	86.08%	90.32%	88.20%

Fig. 5. Average equilibrium values below the dominant strategy price.

we have also used the utility scores of the second best bid in the case of a second-score auction, i.e., the payoff for the bid taker. In all trials, the utility scores achieved in multi-attribute first-score auctions were, on average, 1.52% higher than those achieved in multi-attribute English auctions and 5.83% higher than those achieved in second-score auctions. Fig. 5 shows that bidders in our experiment used similar strategies in single-attribute as well as multi-attribute auctions, as both histograms follow the same pattern. In the single-attribute trials, the equilibrium values also achieved in first-price sealed-bid auctions were, on average, 3.72% higher than those achieved in English auctions and 6.9% higher than those achieved in second-price sealed-bid auctions.

Bidders in our experiment defined their market expectation and risk premium independently of each other and calculated their bids based on these valuations. Therefore, in the sealed-bid auctions, we have many features of a private value model. Of course, our scenario also had a common value element. Because English auctions reveal a lot of information about the other bidders' valuations, we conducted the English trials at the end of each session. In fact, several bidders revised their initial market expectations in these trials and bid above their initial valuation. On the other hand, the subjects' market expectations in our experiments differed substantially. One reason for this is that during the first set of experiments, the ATX dropped by nearly 100 points. Consequently, the winning bidders in our English auctions did not often have to bid up to their valuation. The higher bidding in first-score auctions is consistent with the assumption of risk aversion in the IPV model. For a bidder, the increment in wealth associated with winning the auction at a reduced bid weighs less than the possible loss of not winning due to such a bid. The results of the second-score auctions depicted in Fig. 5 are the ones of the second best bid and are therefore lower than the other auctions.

4.3. Strategic equivalence between second-score and English auctions

A basic assumption of the IPV model is the strategic equivalence of the English and the second-score auction. In both cases, the dominant strategy is to bid up to one's true valuation. In the experiment, the utility scores of the highest bid in multi-attribute second-score auctions were, on average, 0.4% below the dominant strategy score. In multi-attribute English auctions, they were on average 5.36% below the highest valuation. The single-attribute trials were quite similar. In single-attribute second-price auctions, bids were 2.57% below the dominant strategy price, in single-attribute English auctions they were 10.75% below. Using a t-test ($\alpha = 0.05$), we had to reject the null hypothesis of strategic equivalence between English and second-score auction in both the single-attribute and the multi-attribute case (Fig. 6).

Bidders in our experiment had a good understanding of bidding strategies in both English and second-score auctions. This is a reason why the best bid in a second-score auction is close to the dominant strategy score. The lower scores in English

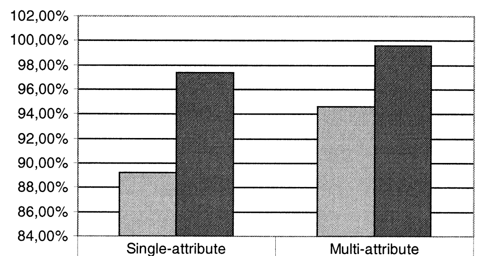

	Single-attribute	Multi-attribute
English	89.25%	94.64%
Second score	97.41%	99.60%

Fig. 6. Comparison of the winning bids in English and second-score auctions.

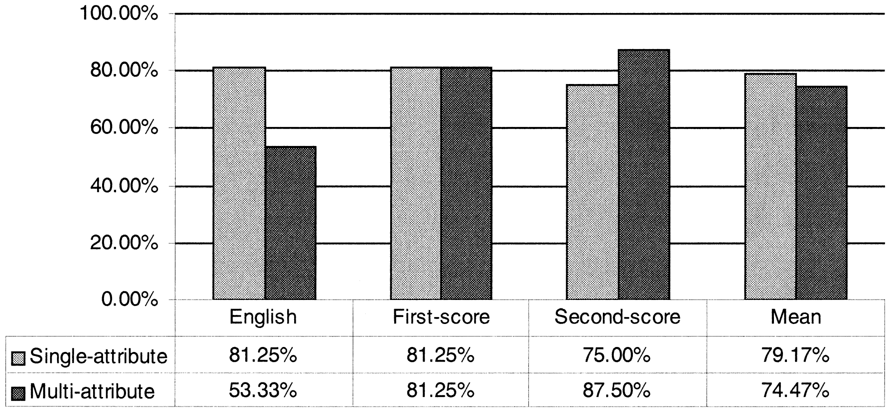

Fig. 7. Efficiency of different auction formats.

auctions can again be attributed to the fact that bidders' market expectations and risk attitudes were quite heterogeneous and the bidders with the highest valuation did not have to bid up to their valuation in order to win the auction. Another explanation for the behavioral breakdown of the strategic equivalence of English and second-score auctions is the preference reversal phenomenon, where theoretically equivalent ways of eliciting individual preferences do not produce the same preference ordering.

4.4. Efficiency of multi-attribute auctions

In single-attribute private value auctions, efficiency is measured in terms of the percentage of auctions where the high value holder wins the item. Efficiency has to be computed slightly different in the case of multi-attribute auctions. Here, the high value holder is the one where one of her valuations (containing strike price and volatility) provides the highest overall utility score for the buyer. Subjects had to report these valuations before each session to the experimenters, based on their market expectations. In all trials, 79.17% of the single-attribute auctions and 74.47% of the multi-attribute auctions were efficient. Using a Chi-square test ($\alpha = 0.05$), you would not find a difference between the efficiency of single-attribute and two-attribute auctions (Fig. 7).

A possible explanation for the high efficiency in all auction schemes is that the market expectations in most sessions varied substantially. Kagel et al. (see Ref. [19], p. 573) reports that the larger the difference between the highest valuation and the lowest valuation and the smaller the number of bidders, the higher the average efficiency levels reported. The slightly lower efficiency achieved in multi-attribute auctions is a possible consequence of the difficulty for the bidder to determine the "best" bid, meaning the combination of values providing the highest utility for the buyer. It can happen that bidders bid below their true valuation per accident. The low efficiency in multi-attribute English auctions is due to the fact that several bidders revised their initial valuations during the course of an English auction. As we already mentioned, in the case of more than two negotiable variables, it is crucial to provide bidders with a decision support tool helping them to determine the "best" bid easily. This can help to maintain high efficiency in multi-attribute auctions.

5. Conclusion and discussion

We believe that multi-attribute auctions are an interesting extension to the set of market mechanisms already in use. The experiments we conducted were an important first step and helped us learn about the many issues one has to consider when implementing a multi-attribute auction. A basic question of auction design theory is which auction format maximizes the bid-takers profit. We were able to show that in the experiment, the overall utility achieved in multi-attribute auction formats was significantly higher than in single attribute auctions. We also found that financial derivatives are a suitable application domain for multi-attribute auction schemes. These are promising results and we think that multi-attribute auction mechanisms have the potential to also automate negotiations in many other real-world exchanges, e.g., in tourism market places, corporate and governmental procurement.

It is important to consider a few issues when applying multi-attribute auctions. Bidding is more complex in multi-attribute auctions, as it is not obvious for the bidder right from the start which combination of attributes provides the highest overall utility for the bid-taker. Appropriate decision support tools for the bidder play a crucial role in the case of three or more negotiable attributes. If multi-attribute auctions are deployed in another application domain, it is also important that preferential independence of the negotiable attributes is given. If this is not the case, the auctioneer has to use a more complex utility function covering the interdependencies between the attributes.

Besides, buyers also have to get used to the new tool and have to learn about the consequences of different parameter settings in their utility function. We think that in a professional environment such as corporate procurement buyers will adapt quickly to the new tool. Less experienced buyers, however, do not know market conditions that well and face the danger of the parameters in their utility function not corresponding to their preferences. For these cases, we developed an additional feature in our electronic brokerage service allowing buyers to adapt certain parameters during the course of an English auction. For example, if a buyer recognizes that all bids achieve very high scores on a certain attribute, she can decrease the importance of this attribute. Bidders are informed of a change in the buyer's utility function via e-mail and can adapt to these changes. This enables a more dynamic bidding procedure.

References

[1] K. Backhaus, B. Erichson, W. Plinke, R. Weiber, Multivariate Analysemethoden, 8th edn., Springer, Berlin, Germany, 1996.

[2] M. Bichler, Decision analysis — a critical enabler for multi-attribute auctions, Proceedings of 12th Electronic Commerce Conference, Bled, Slovenia, 1998.

[3] M. Bichler, M. Kaukal, Design and implementation of a brokerage service for electronic procurement, Proceedings of DEXA '99 Workshop on Electronic Commerce, Florence, Italy, 1999.

[4] M. Bichler, R. Klimesch, Simulation multivariater Auktionen- Eine Analyse des Handels mit Finanzderivaten, Wirtschaftsinformatik (1999) in press.

[5] M. Bichler, M. Kaukal, A. Segev, Multi-attribute auctions for electronic procurement, Proceedings of First IBM IAC Workshop on Internet Based Negotiation Technologies, Yorktown Heights, NY, 1999.

[6] F. Branco, The design of multidimensional auctions, RAND Journal of Economics 28 (1) (1997) 63–81.

[7] Y.-K. Che, Design competition through multidimensional auctions, RAND Journal of Economics 24 (4) (1993) 668–680.

[8] R.T. Clemen, Making Hard Decisions: An Introduction to Decision Analysis, Wadsworth Publishing, Belmont, CA, 1996.

[9] J.C. Cox, B. Roberson, V.L. Smith, Theory and behavior of single object auctions, in: V.L. Smith (Ed.), Research in Experimental Economics, JAI Press, Greenwich, CT, 1982.

[10] M. Cripps, N. Ireland, The design of auctions and tenders with quality thresholds: the symmetric case, The Economic Journal 104 (1994) 316–326.

[11] S. Dasgupta, D.F. Spulber, Managing procurement auctions, Information Economics and Policy 4 (1989) 5–29.

[12] G. Debreu, Representation of a Preference Ordering by a Numerical Function, in: R. Thrall, R. Davis (Eds.), Decision Processes, Wiley, New York, 1954, pp. 159–166.

[13] R. Engelbrecht-Wiggans, On a possible benefit to bid takers from using multi-stage auctions, Management Science 34 (9) (1989) 1109–1120.

[14] M. Fan, J. Stallaert, A.B. Whinston, A web-based financial trading system, IEEE Computer 32 (4) (1999) 64–70.

[15] R. Fisher, D. Ertel, Getting Ready to Negotiate, Penguin Books, New York, 1995.

[16] D. Friedman, S. Sunder, Experimental Methods: A Primer for Economists, Cambridge Univ. Press, Cambridge, UK, 1994.

[17] J.H. Kagel, Auctions: a survey of experimental research, in: J.H. Kagel, A.E. Roth (Eds.), The Handbook of Experimental Economics, Princeton Univ. Press, Princeton, 1995, pp. 501–587.

[18] J.H. Kagel, A.E. Roth, The Handbook of Experimental Economics, Princeton Univ. Press, Princeton, NJ, 1995.

[19] J.H. Kagel, R.M. Marstad, D. Levin, Information impact and allocation rules in auctions with affiliated private values: a laboratory study, Econometrica 55 (1987) 1275–1304.

[20] R.L. Keeny, H. Raiffa, Decision Making with Multiple Objectives: Preferences and Value Tradeoffs, Cambridge Univ. Press, Cambridge, UK, 1993.

[21] G.E. Kersten, S.J. Noronha, Negotiation and the Internet:

Users' Expectations and Acceptance vol. 1998 Carleton University, Canada, 1998.

[22] R.W. Kolb, Financial derivatives, New York Institute of Finance, Englewood Cliffs, NJ, USA, 1993.

[23] O.R. Koppius, Electronic multidimensional auctions: trading mechansims and applications, Proceedings of Edispuut Workshop '98: E-Commerce — Crossing Boundaries, Rotterdam, The Netherlands, 1998.

[24] Y.-K. Kwok, Mathematical Models of Financial Derivatives, Springer, Singapore, 1998.

[25] J.-J. Laffont, J. Tirole, A Theory of Incentives in Procurement and Regulation, The MIT Press, Cambridge, MA, 1993.

[26] R. McAfee, P.J. McMillan, Auctions and bidding, Journal of Economic Literature 25 (1987) 699–738.

[27] K.A. McCabe, S.J. Rassenti, V.L. Smith, Designing a uniform price double auction: an experimental evaluation, in: D. Friedman, J. Rust (Eds.), The Double Auction Market: Theories and Evidence vol. 15 Addison-Wesley, Reading, MA, 1993.

[28] R.M. McLaughlin et al., Over-the-Counter Derivative Products, McGraw-Hill, New York, 1998.

[29] G.B. Metha, Preference and utility, in: S. Barbera, P.J. Hammond, C. Seidl (Eds.), Handbook of Utility Theory vol. 1 Kluwer Academic Publishing, Dordrecht, 1998, pp. 2–40.

[30] P.R. Milgrom, R.J. Weber, A theory of auctions and competitive bidding, Econometrica 50 (1982) 1089–1122.

[31] S. Rassenti, V.L. Smith, R.L. Bulfin, A combinatorial auction mechanism for airport time slot allocations, Bell Journal of Economics 13 (1982) 402–417.

[32] M.H. Rothkopf, A. Pekec, Computationally manageable combinatorial auctions, Proceedings of Maryland Auction Conference, Maryland, 1998.

[33] T.L. Saaty, The Analytic Hierarchy Process, McGraw-Hill, New York, USA, 1980.

[34] A. Segev, M. Bichler et al., Component-based electronic commerce, in: M. Shaw, R. Blanning, T. Strader, A. Whinston (Eds.), Handbook on Electronic Commerce, Springer, New York, 1999.

[35] V.L. Smith, Microeconomic systems as an experimental science, American Economic Review 72 (1982) 923–955.

[36] M. Stolze, Soft navigation in product catalogs, Proceedings of European Conference on Digital Libraries, Heraklion, Crete, Greece, 1998.

[37] W. Vickrey, Counterspeculation, auctions, and competitive sealed tenders, Journal of Finance 3 (1961) 8–37.

[38] R. Wilson, Strategic analysis of auctions, in: R.J. Aumann, S. Hart (Eds.), Handbook of Game Theory vol. I North-Holland, Amsterdam, 1992, pp. 227–271.

[39] H. Wold et al., A synthesis of pure demand analysis, Skandinavisk Actuarietidskrift 26 (1943) 85–118.

[40] E. Wolfstetter, Auctions: an introduction, Journal of Economic Surveys 10 (4) (1996) 367–420.

Martin Bichler is an Assistant Professor of Information Systems at the Vienna University of Economics and Business Administration. He received his MSc from the University of Vienna, Austria and his PhD in Information Systems from the Vienna University of Economics and Business Administration. He is also an affiliate research fellow at UC Berkeley's Fisher Center for Management and Information Technology. His research focuses on electronic commerce infrastructures and advanced auction schemes. Dr. Bichler has published papers in journals such as, e.g., Journal of Distributed and Parallel Databases, International Journal of Electronic Markets, Wirtschaftsinformatik, International Journal of Cooperative Information Systems, and Journal of End User Computing.