MERIT-Infonomics Research Memorandum series

Auctions - the Big Winner Among Trading Mechanisms for the Internet Economy

Rudolf Müller

2001-016

MERIT – Maastricht Economic Research Institute on Innovation and Technology
PO Box 616
6200 MD Maastricht
The Netherlands
T: +31 43 3883875
F: +31 43 3884905

http://meritbbs.unimaas.nl
e-mail:secr-merit@merit.unimaas.nl

International Institute of Infonomics

PO Box 2606
6401 DC Heerlen
The Netherlands
T: +31 45 5707690
F: +31 45 5706262

http://www.infonomics.nl
e-mail: secr@infonomics.nl

Auctions - the Big Winner Among Trading Mechanisms for the Internet Economy

Rudolf Müller

May 2001

Abstract

Auctions are probably the most important mechanism for dynamic pricing in electronic commerce. Although they constitute a very old mechanism as well, the new popularity has raisen a lot of questions on the appropriate design of an auction mechanism for a particular situation. This chapter describes reasons for auction popularity by setting them into the context of trends in electronic commerce. We then illustrate the main issues in auction design. Our analysis starts with simple single-item auctions, as we can see them in many B2C markets. We then look at the more complex auction designs, which are necessary for B2B markets. For the latter design has to take into account that buyers want to purchase collections of items and services, and that the valuation for winning collections is not simply equal to the sum of valuations of single items. We show how multi-item auction mechanisms can benefit from a synthesis of microeconomic and mathematical optimization models.

Rudolf Müller
International Institute of Infonomics and Department Quantitative Economics
University of Maastricht, PO Box 616, 6200 MD Maastricht, Netherlands
Phone +31-43-3883799, e-mail: r.muller@ke.unimaas.nl
http://www.personeel.unimaas.nl/r.muller

Introduction

An auction is a mechanism to re-allocate a set of goods to a set of market participants on the basis of bids and asks. In its classical form one seller, the auctioneer, wants to find a buyer for a single, indivisible item among a group of bidders. The best known auction design for this case is the increasing bid auction, or *English auction*. In this auction the auctioneer receives bids until she decides to terminate the auction, at which point the bidder with the highest bid receives the item, unless this bid is below the auctioneer's reservation price. The price that the bidder has to pay is equal to her last bid. An auction design is specified through three elements:

1. the *bidding rules* define what bidders may bid for and when.
2. the *market clearing rules* define when and how the allocation of items to bidders is decided and what bidders have to pay.
3. the *information disclosure rules* define when and which information is disclosed to whom.

The English auction has the simple bidding rule that every bidder can make a bid at every time, the market clearing rule that the highest bid wins and has to pay this bid, and the information disclosure rule that the highest bid is known. If the auction takes place in a room where all bidders are present, then bids are known to all bidders. In English auctions done on the Internet, only the price of the current high bid may be announced.

Auctions may not only be used to sell items or services but also to purchase them. We call them in this case a *reverse auction*. For purchasing *first price, sealed bid auctions* are the common format. Here bidders submit sealed bids. The bidding rules define who is allowed to participate. The market clearing rule assigns the bid to the bidder with the lowest price for the item. The price that is paid by the auctioneer is equal to this bid price. This format is typical for contracting complex services in the construction industry. The advantage of a reverse auction is that the buyer can specify in detail the service, and due to sealed bids is not required to take the cheapest offer, but to apply other criteria than price in the market clearing phase.

As another example for a classical auction we mention the *Dutch auction*. It uses a price clock that starts at a high price, which steadily decreases until a bidder decides to buy at the current price. By shouting or other means she stops the bidding process and receives the item at the current price. The Dutch auction is typically used for selling agricultural products. It is fast since it ends with the first bid from a bidder, and thus in particular appropriate for markets that have to negotiate a high volume of transactions in a short period of time.

Finally, there are so-called *double auctions* in which buyers make bids and sellers make asks, where a participant may be buyer and seller at the same time. Bids and asks are displayed in an order book. As soon as a bid for an item is higher than the current ask, the item is traded for a price in between of the two values. Double auctions are used at stock exchanges, and now frequently at spot markets for commodities in B2B marketplaces.

The Internet has become an important platform for trading, which is illustrated through all chapters of this book. Auctions got within Internet trading a much more prominent role than they had in offline trading. Reasons for that are discussed in section 2. Section 3 is dedicated to the fundamental choices one has in auction design. We illustrate them along single item auctions. In Section 4 we introduce the more complex setting, where several indivisible items are auctioned at the same time. Such auctions are of particular relevance in B2B e-commerce. Section 5 illustrates possible designs for such multi-item auctions, as well as the challenges in implementing such design. We summarize in Section 6.

The goal of this article about auctions and the Internet is to provide an introduction to the topic. We summarize those features of auctions that have to be understood for users and providers of Internet auctions. We try to omit mathematical details where possible, but do nevertheless try to explain the mathematical problems that have to be solved in auction design. We hope to stimulate the reader to start from this article for a more detailed tour through the recent auction literature. Last but not least we try to combine quite different aspects of auction theory in an interdisciplinary approach. A reader who would like to learn at this point more about the history of auctions, or read about the different formats in more detail, should have a look at (Agorics Inc., 1996). For diving deeper into the topic we recommend (Klemperer, 1999).

Auctions on the Internet

Why are auctions such a popular trading mechanism on the Internet? Broadly spoken, the reason for this popularity is that Internet enables a wide range of organizations and people to use auctions for a wide range of items and services at rather convenient transaction costs. Let us elaborate on this in some more detail, without repeating however the many possible explanations that can be found in (Herschlag & Zwick, 2000).

Before Internet could be used as communication platform for auctions, the range of potential participants in auctions, either as a bidder or an auctioneer, was rather limited. Now with the Internet, and with auction platforms accessible via the Internet, everybody can participate and even create an auction. At first place this is a phenomenon of tremendously decreased transaction costs. Physical presence at the auction is not necessary anymore, when product information as well as bids and asks can be communicated electronically. Many companies have seen this opportunity in an early stage of E-commerce penetration on the Web and set-up *private-to-private* auction platforms. Trading of collectibles has soon been complemented by using auctions as trading mechanism in retailing. Auction sites functioned as an "electronic catalogue with dynamic pricing", selling new products in *business-to-private* auctions. Finally, the business-to-business applications entered the stage. Meanwhile the Internet hosts a huge collection of auctions. The best way of getting an overview is probably to consult listings like http://www.internetauctionlist.com/ or http://www.auctionguide.com/.

The popularity of Internet auctions goes hand in hand with other business trends observed for the Internet. Three of these should be mentioned here.

Firstly, we see *changing roles of intermediaries* and *new forms of intermediation* (Sarkar, Butler, & Steinfield, 1998, Scott, 2000,). Already in October 1998, the Keenan report stated "The power of instant communication destroys the power of middleman to hide the real price from buyers and sellers, creating new intermediaries who will control the distribution of basic goods. Distribution channels that are inherently inefficient, such as wholesale-retail chains, may be re-intermediated by a new middleman equipped with Internet Exchange technology" (Kennanvision.com, 1998).

Auctions on the large scale are only feasible due to the Internet. They are thus examples of *cybermediaries*. To some extend they constitute dis-intermediation as they are likely to replace a significant part of the business of brick-and-mortar auction houses. Whether the Internet auctions will be sustainable is likely to dependon whether or not they can provide a reliable service to their customers at a reasonable price. The first generation Internet auctions seem to require just a software platform to create a successful auction business. This rather simple business model for consumer-to-consumer auctions, which lets the participants take care of all other phases of the transaction, seemed to be sufficient for success. This is best explained best by enormous *first mover advantages*, as an auction is a perfect example for strong network externalities. Sellers like to use sites that have many visitors, since a large number of bidders increases the expected revenue. Bidders again prefer these sites because they can chose from a large number of different offers. They might hope that the segmentation of the market has the effect of less competitors in a specific auction, ignoring to some extent that they are looking for precisely that product and value it therefore higher.

With more and more auction sites entering the stage, and with the trend that portals or companies that started with traditional retailing use now auctions as one of many versions of trading, the very simple model of trading platform can however risk to become a commodity. The trading itself is only a small part of the total change of ownership that, at the end has to include financial and physical settlement. Herschlag and Zwick (Herschlag & Zwick, 2000) give the example of Teletrade as an auction site that adds services for settlement for buyers and sellers, while ebay likes to classify itself as a “person-to-person” trading community. Whether the latter is a sustainable business model depends very much on the experiences that consumers make with their private trading partners. It could well be that online auction sites develop their own netiquete, maybe established through ratings of the auctioneers and bidders. But it could also be that on the long term auction sites which take care of logistics and settlement have an advantage against the platform only solutions. When more customer service is required, the traditional auction houses may even become the strongest competitors as they can rely on an established brand name as well as on their experiences in all phases of the settlement. However this brand name has to be protected and thus these sites will have to provide an above average service. It seems likely that we will see a range of service versions, where higher quality of service might have to be paid for by participation fees. Such fees could also serve partly to provide customers an insurance against bad fulfilment of the contract.

The second business trend is *customising*, hand-in-hand with *personal pricing*. The principal idea of an auction is to find among a set of potential buyers those with the highest willingness to pay, where the second highest valuation is about what can be expected as revenue (due to a fundamental theorem on auctions this is true for all major single-item auctions with private, independently distributed valuations of bidders, see the next section). Auctions make therefore sense in cases where the valuations are diverse, and where identification of the customer with highest valuation is difficult. However this concept seems to be in contrast to the expectation that consumers of an online auction have. Typically, they would like to make a bargain. Auctions would certainly become unpopular if they were used by the high-end customers, driving prices to the same level or above prices in the store or catalogue. Therefore online auctions which sell new products have to put much attention on proper product selection. There has to be a kind of scarcity and the product has to address a consumer group that observes a discount through the auction. Typical products are thus completely new products which are not available yet, or products that are replaced by a new version. In both cases the auction implements a two-fold segmentation. It selects through the type of products consumers who, in the first case, do not want to wait for until the next generation is in the store, and, in the second case, do not care for the most recent technology. Among those it applies personal pricing by finding the consumer with highest willingness to pay.

Real customizing in such a market is rather limited. There has to be more than one bidder interested in the same copy, since otherwise the second price would be zero. A typical customisation dimension is time. If for example bandwidth in a telecommunication network is required for a certain time-interval at a certain capacity, or if a traveller needs a ticket for a very specific time, an auction can have the strength that searching and negotiating an alternative trade is not feasible, as it would exceed the time limit. On the other hand, auctioning last minute tickets identifies those travellers who are not bound to a certain time.

Auctions are not the only trading mechanism that can realise personalised pricing. There are for example one-to-one price negotiation sides on the Web, in which a customer negotiates individually with a merchant on the price. From the list of frequently asked questions at haggelzone.com one can read (as of 11.12.2000) "Hagglezone.com is the anti-auction**.** At www.hagglezone.com, prices go down, not up. The buyer is in control of negotiating the price down against the chosen Haggler, instead of bidding prices up against hordes of other consumers". A "haggler" means a salesperson the customer tries to negotiate with a price. Note that the underlying mechanism is similar to a Dutch auction, rather than an "anti-auction". Namely, in case the merchant is negotiating at the same time with several customers (which is not observable by the buyer), he can lower the price steadily until the first customer is willing to buy. If there are several items to sell, she can however continue decreasing prices for the other customers. Only if there is exactly one customer interested in the item it differs from a Dutch auction. In this case an auction would not be recommendable to the merchant, so the advantage is on the merchant's side. One-to-one price negotiation is actually most appropriate, if more terms than the price are part of the negotiation, in which case there is likely exactly one customer for a given mix of terms. Furthermore, the price serves as a trade-off between matched and un-matched customer expectations. An example is given by www.tradeaccess.com. Finally, there are sites like www.priceline.com realising B2C reverse auctions, and sites like www.letsbuyit.com that realises a flexible price through bundling of demand.

The third business trend is *customer involvement*. At an auction site it is the customer who determines the price, not the auctioneer. With many items on sale at the same time, the customer gets a huge selection of items. In particular for moderated auctions, as they can be found on various sites (for example at www.ricardo.de) the customer contributes to the entertainment to non-active visitors of the auction site by taking part in an open, observable competition. Finally, customers of private-to-private auctions contribute to the success by setting up own auctions. Customers can even become almost professional traders by the help of the online auction infrastructure. Auctions like www.ricardo.de observe that items in their private-to-private auctions have previously been purchased in the business-to-consumer area. The Keenan report cites www.eBay.com with the information that in 1999, 10.000 customers made most of their personal income from trading goods with eBay (Keenanvision, 2000).

Auction Design

There are three major categories of auctions

1. The classical auction in which an auctioneer sells to a group of bidders. The introduction presented already the main versions: *English auction*, *first price sealed*

bid auction, and *Dutch auction*. A fourth member of this category is the *second-price sealed bid auction*, also called the *Vickrey auction*, in which the winner has to pay only the second highest bid price.

2. The reverse auction in which the auctioneer wants to purchase items from a set of bidders. One may use all of the four formats here, although the first-price sealed-bid auction seems the most used one.
3. The double auction in which a group of buyers and sellers meet each other. The classical model is an *exchange*. Here sellers post asks and buyers post bids.

Inside each of these categories, many parameters have to be chosen to fine-tune the auction procedure. These include features like activity rules, minimum increments, or decrements, as well as information disclosure policy in multi-round auctions. A mix of general format and features can be used to address specifically the goals of an auction, given a certain market situation. Of increasing interest are hereby *multi-item auctions*, where the seller has a set of items to sell, and buyers are interested in purchasing certain subsets of this set. In this section we look at the role of information in auction design for single items, and then see in the next section why multi-item auctions require a special treatment.

Why is it such an important factor in the design of an auction, which information bidders and the auctioneer have about the value of the items on sale, and the information that they have about the other agents' value estimates? We will describe in the following two extremes, in order to give an answer to this question. Again, when compared with the rich literature on auctions, see e.g. (Klemperer, 1999), our treatment has to be rather introductory. Furthermore, information distribution in practice is usually in between the two extremes.

The first extreme is that where every bidder has *private value* of the item. The mathematical formalization is a value v_i of bidder i. To capture the degree of information that other bidders have about *i's* valuation we assume that they observe it as random variable X_i. In case of a discrete random variable they know thus the probability that the valuation of i takes a certain value. The private value model assumes that these random variables X_i are independent random variables. It is save to assume that the random variables are all identical, if bidders have no information about differences in the competitors' values for the item. An example for an auction with a private value model is that of auctioning a collectible, e.g. some painting from an *unknown* painter from the 19th century, where the value for every bidder is a function of personal taste. We assume an unknown painter, because we want to exclude that the purchase is meant as an investment for later sale, in which case we would loose some of the independence of valuations.

The second extreme is that of a *common value*. In this case the value of the item is independent of the bidder, or in other words common to every bidder. Mathematically this situation is modelled by a single random variable. The exact distribution of the random variable is unknown to the bidders, and possibly as well to the auctioneer. However all parties may have a certain degree of information about this random variable, by having done a priori some research. An example is shares on the stock market. Their value in the future, i.e. the price at which the stock market will trade them, is completely independent of bidders' personal taste. The value is a random variable, of which well-informed bidders have a better estimate than less informed bidders.

Auctions in practice are situated in between these two extremes. Buying a painting might well be seen as an investment, in which case the future demand for the painting, in other words the possible price that it might achieve in an auction, plays a significant role for its value. The decision which shares to buy at which price can also be influenced by preferences for a brand name, or from hedging considerations for the own portfolio, which adds a private value component.

Let us see now how the two cases influence the choice for the auction design. Notice that the bidding process may reveal information, in other words it signals other bidders. The auctioneer might want to support these signals, if it reduces the risks of bidders and thus increases the willingness to pay. She might want to mind these signals, if bidders can collude this way and decrease the outcome of the auction.

Consider the ascending price English auction. If bidders have private valuations v_i, and if they behave rationally, they should participate in the auction until the bid reaches v_i. Indeed at that point a bidder's *utility* u_i which is defined by $u_i = v_i - p$, p being the price that the winner has to pay, is becoming *0*, so she becomes indifferent between win or not to win. Observing the bid prices of other bidders is not of relevance for her. The only impact that they have to her strategy is that they determine at which price she is able to win the item. If she has the highest valuation and all bidders participate actively, she will purchase the item at a price slightly higher than the second highest bid from her competitors. If competitors follow the same rationale, this bid is equal to the second highest valuation of bidders. In this case the English auction realizes a so-called *second price auction*.

Three factors play an important role here. Firstly, it is the *activity* of bidders. If activity is not stimulated, then the bidder with the highest valuation risks to make an initial bid that is already strictly higher than the second highest valuation. In this case all others drop out and she doesn't realize the maximum possible utility. It's thus wise for all bidders to start with a careful low bid. It's also advisable for the auctioneer to let the auction start with a low initial bid, even if her reservation price is higher. The *reservation price* is the lowest price at which the auctioneer is willing to sell. In Ebay.com private-to-private auctioneers can for example set a start price *and* a reservation price. The latter is not visible to the bidders.

Secondly, the *termination rule* of the auction plays a role. In the traditional offline auction setting, the auctioneer uses the *going, going, gone* mechanism to finalise the bidding. This enables the auctioneer to evaluate carefully, whether the current highest bidder is also the one with the highest valuation. If, like in many online auctions, the end is a fixed point in time, the auction looses its flavour of a second price auction. Now a bidder with a low non-competitive valuation may succeed to win by making the highest bid just before the auction finishes. Bidders can protect themselves against such competition only if they make a high bid early enough. But this makes it likely that a winning bidder has to pay a price that is very close to her own valuation, which turns the auction mechanism into a *first price auction*. Online auctions are therefore considered typically as a hybrid of a first and second price auction (Ockenfels & Roth, 2000). Some auction sites, e.g. www.amazon.com, avoid this by extending the auction automatically for a couple of minutes after the last bid was made. Quite astonishingly, the bidders can in principle avoid to pay more than the second highest valuation in a first price auction by using their information and doing some calculus. Indeed if they know how the random variables determining other bidders' valuations are distributed they can calculate a strategy that gives them an expected utility that is equal to that in the English auction. Basically, they have to calculate the *expected value of the second highest bid*. For the interested reader we recommend an auction survey by Wolfstetter (Wolfstetter, 1996), or the original paper by Myerson (Myerson, 1981), to learn more about the precise conditions under which the *revenue equivalence theorem* applies.

Thirdly, *minimal bid increments* are of relevance. A minimal bid increment is the amount by which the next bid has to be higher than the current highest bid. Minimum bid increments can have the effect that not necessarily the bidder with highest valuation wins the auction. Say, for example, that the highest valuation among all bidders is 98, the minimal increment is 5 and the auction is at price of 95. If the current highest bid has not been by the bidder with highest valuation, the latter has to drop out, since winning with 100 would have a negative utility.

Fixed end times and minimal increments may thus both lead to the effect that the auction is not *efficient*. Efficiency is defined as the property that the bidder with highest valuation will win the auction. Efficiency is desirable from a welfare perspective, and also from an auctioneer's perspective. If the bidder with highest valuation wins, and if there is a strong competition, meaning that the second highest valuation is close to the highest, an efficient second price auction is close to optimal for the auctioneer. In economic literature an auction design is said to be *optimal* if the expected revenue for the auctioneer is maximized. We will use the attribute optimal in a different way throughout this paper, namely with respect to a specific instance of bidder valuations. We say that the auction is optimal for this instance if the revenue is equal to the maximum valuation.

Let us now turn back to the role of information, and consider the case of a common value, where the final value of the item will be the same for every bidder, but the bidders have different estimates about this value. We look first at the English auction. We observe easily that the bidders should become careful with every bidder who drops from the bidding, since this indicates that her estimate tells that the (common) value of the item will be less than or equal to the current highest bid. The more bidders drop out, the more reliable becomes this signal from the auction process. An ascending price auction thus reveals information. Necessary is however activity by the bidders. We saw above that with a fixed termination date early bidding is discouraged, and thus information revelation is abandoned. The auction becomes a first price auction, with a significant risk of a *winner's curse*. The winner's curse is the effect that a winner in a common value auction pays more than the (later) value of the item. The winner's curse is actually not tied to first price auctions, and can also occur in second price auctions with common value, though it will be less severe of course.

The possibility to exchange information during the auction process makes up for the main difference between open outcry auctions and *sealed bid* auctions, in which the auctioneer collects bids in sealed envelopes and decides who wins the auction. Certainly, sealed bid auctions have the advantage of further lowering transaction costs, but the stronger the common value component is in the valuation of the item, the higher the risk of a winner's curse. Furthermore, sealed bid auctions require trust in the auctioneer that she does not manipulate the bids. The sealed bid auction knows two versions: the first price sealed bid auction, and the Vickrey auction in which the price is that of the second highest bid. The revenue equivalence theorem (Myerson, 1981) tells us that the expected revenue for the auctioneer is the same for both models in the case of private values, since in the first price auction the bidder should adjust his bid to the expected second highest bid. Under common value situations this is not necessarily the case.

The online auction www.ricardo.de experimented in May 2000 with a second price Vickrey auction. They called it an *undercover auction*. Results have been reported on their site at http://www.undercover.ricardo.de/undercover/mid.htm (11.12.2000). Comments to an announcement of that auction at www.zdnet.com showed that bidders have a hard time understanding the principle of the second price. A concern was that a bidder might be tempted to bid very high, because this increases the chance to win the auction, while only the second price has to be payed. However this strategy would lead to a large loss if more than one bidder applies it. Rather than excluding by this argument the strategy, comments at www.zdnet.com brandmarked the auction mechanism as extremely unfair! Online auctions can well offer Vickrey auctions in a less explicit way. So does www.qxl.com, which offers to set a maximum bid upto which a bidding agent will increase the bid on behalf of the bidder whenever her bid is not the highest: “When you bid, we ask you to put in the maximum bid amount you are willing to pay for the item. Remember that bid amount reflect the amount you are willing to pay per item. So, if you bid £10, and you select to buy a quantity of 10, that means your total bill could be £100 for this auction. Once you set your Max. Bid, we place a bid on your behalf to enter you in the auction, and every time you're outbid, to make sure you stay the high bidder, up to the bid you specify.”

The ricardo case indicates that smart auction design should take the user into account, in the sense that it might be too complex for her to behave optimally. If we follow this argumentation in a more formal way, we see that different auction designs have quite different *computational complexity* for the bidders and the auctioneer. The reader should take computational complexity as a measure of mathematical tasks that have to be solved in order to optimize the own strategy. In an English auction this is quite simple: the bidder observes and decides at every point whether she is willing to increase the bid or not. In the common value case the decision has to take the bidding of competitors into account, in the other case not. Two recent papers that take the cognitive costs of bidders into account are (Parkes, Ungar & Foster, 1999), and (Nisan & Ronen, 2000).

Let us finally consider the Dutch auction. In the case of private valuations the bidder should use the same strategy as in the sealed bid first price auction or in the auction with a fixed termination date: forecast the second highest valuation of all bidders, and stop the descending auction clock at this value, in case that the own valuation is higher. Notice once more the computational problem of calculating this termination point! Without computational effort, maximum expected utility is not realized, because the bidder either fails to win, or she wins at a price strictly higher than the second highest valuation among the bidders. In case of a common value auction, the Dutch auction has the disadvantage that it does not reveal any information about other bidders’ estimates. Bidders can simply not observe other bidders’ decision to leave the auction. This delivers again a high risk for a winner’s curse.

What we learn from this discussion is that auction theory is a very complex field of mathematical research. Although there are strong theorems that can be used to reduce the number of cases to be considered, like the revenue equivalence theorem or the revelation principle (Myerson, 1979), these theorems have to make assumptions that are not necessarily applicable in practice. Firstly, they assume *rational decision making*, secondly they assume unbounded computational and mathematical capabilities of decision makers, and thirdly they seem not to reflect that the process of an auction can have a strong psychological influence on the decision behaviour, which is in any case not necessarily completely rational. Every specific auction design requires not only a detailed theoretical foundation, but also empirical and experimental analysis, in order to be able to predict its outcome for bidders and auctioneer. The many auctions on the Internet do not only challenge auction research for that, but also turn out to become a means to improve the understanding of auctions. With millions of auctions going on all the time, auction research gets the empirical data required to fine-tune designs. As an outstanding example of such online research we refer here to (Ockenfels & Roth 2000).

To summarise, we can say that research on single-item auctions has to look in detail on bidders information about the item, on the spread of this information among bidders, the computational complexity of evaluating this information, and the diffusion of information during the auction process.

Complements and Substitutes

So far we considered the case that exactly one item is sold in an auction, and that the bidders valuation of that item is completely independent from other events. This assumption is hardly feasible in practice. Bidders might at the same time be active in many auctions, and the value of an item that they can win might well depend on other items they can purchase. For example, they might succeed to buy a *complement* for an item, giving the union of the two items a higher value than the sum of individual values. They might also face the situation in which they have purchased items that *substitute* each other. We consider in this section auctions in which one auctioneer tries to sell a *set of items* to a group of bidders. Such auctions are called *multi-item* auctions. The items may be different or identical. In the pure case of identical items we talk about *multi-unit* auctions. We assume furthermore that items are not divisible, such that it is not possible, say, that two bidders win both half of an item.

Multi-item auctions have numerous applications, like selling airport time slots (Rassenti, Smith & Bulfin, 1982), railroad segements (Brewer, 1999), and shipping contracts (Caplice, 1996). They are of interest as coordination mechanism in multi-agent systems (Nisan & Ronen, 2000), and have recently been investigated largely from a computer science and operations research perspective (e.g., Fujishima, Leyton-Brown & Shoham, 1998, Leyton-Brown, Pearson & Shoham, 2000, Leyton-Brown, Shoham & Tennenholtz, 2000, Rothkopf, Pekec & Harstad, 1998, Sandholm, 1999, Sandholm & Suri, 2000, Tennenholtz, 2000, Vohra & de Vries, 2000, Wellman, Walsh, Wurman & MacKie-Mason, 1998).

In terms of the broad range of auction features that we discussed in the previous section we have to restrict ourselves. We assume in the following private value. The value is now not anymore related to one single item, but is a function that maps every subset of items to a real, non-negative number. Such a function determines the *type* of a bidder. A type is thus a mapping $v: 2^S \rightarrow \Re$. The outcome of the auction is an allocation of subsets of items to bidders. We denote by I the set of items and by B the set of bidders. Let J be a numbering of the union of all possible bids from all bidders. The set J gives a unique identifier to every subset for every bidder. We use it in the following to simplify our notation.

As in every auction mechanism it has to be decided which are the winning bids. Based on our notation we can use an *0-1* vector x to model this decision, with the interpretation that $x_j = 1$ if and only if bid j is assigned. By w_j we denote the bid price. Efficiency of the auction can now be expressed as follows. The final allocation x should maximise the expression $\sum_{j \in J} v_j x_j$. . If p_j denotes the price which has to be paid for winning bid j, then optimality is achieved if $\sum_{j \in J} p_j x_j$ equals this maximum value.

Already the number of possible allocations of items, it is $(|B| + 1)^{|S|}$, indicates that multi-item auctions can be expected to be far more complex than single-item auctions. They are almost the most general market mechanism for markets with indivisible goods, and have in particular been studied in the literature on resource allocation in multi-agent systems. What they don’t capture are preferences of a bidder on allocations among the other bidders. They are also a special case of multi-attribute auctions. To see this observe that we can represent a subset by an attribute vector, with one component for every item in *I*, and an attribute having value *1*, if the corresponding item is contained in the subset, *0*, otherwise. This way every bid, becomes a bid on a *product*, which satisfies certain attributes. The fact, that only a limited number of copies are available of each item translates into the fact that the auctioneer in the multi-attribute auction can only fulfil a certain mix of contracts.

Multi-item auctions started to gain scientific popularity in terms of the huge frequency auctions organised by state authorities since about 1994. Some failures of the first designs of such auctions in New Zealand, and Australia, caused the FCC in the US to invite auction experts to help in creating an appropriate multi-item auction (McMillan, 1994). The result was a multi-round, parallel single-item auction, with rules for minimum increments and participation. Let us try to illustrate why, aside from the sheer number of possible allocations, multi-item auctions are such a complex issue.

A major challenge in multi-item auction design is to solve the *exposure problem*. Suppose that items *Q* and *R* are strong complements to each other for a bidder *A*, thus she would be willing to pay 100$ if she can purchase both of them, while she values *Q* and *R* alone only at *20$* each. Suppose that competitors *B* and *C* have only interest in *Q* and *R*, respectively, with a valuation of *30$*. If *Q* and *R* are auctioned independently in single-item auctions, then *A* may have to bid on both, *Q* and *R*, higher than *30*, due to the competition by *B* and *C*. At this time she might well win *Q* at, say, *30*, but fail to win *R*, leaving her with a negative utility of –*10 $*. If bidder *A* can announce her high valuation for *{Q,R}* to the auctioneer, *B* and *C* might on the other hand be discouraged by the *threshold problem*. Individually they are not able to compete against the high bid of *A*. Together their value might however add to more than the *100$* by *A*. If the auction design is not capable to reveal this, *A* might win despite the higher revenue obtainable for the auctioneer if she assigns to *B* and *C*. However revealing this information in an ascending price auction might also be tricky for *B* and *C*. If for example *B* makes a high bid for *Q*, then a relative low bid from *C* for *R* is enough to beat *A*. *C* enjoys in this case *free riding* on the high bid from *B*.

A way out of this dilemma is to use a *utilitarian revelation mechanism*, in which bidders tell the auctioneer their type and the auctioneer computes an allocation based on this information. A problem with this approach is that bidders might be better off if they do not tell the truth, but report a wrong type *w* instead. This problem will be discussed below. A second problem of the revelation mechanism relates to the fact that an open, ascending price auction may function to inform other bidders in common value auctions. This issue is not addressed here, as we said to restrict ourselves to private value auctions. One might however argue that a utilitarian revelation mechanism may be repeated in several rounds, allowing bidders to adjust their type by withdrawing or reducing a bid made in a previous round. This kind of mechanism has to our knowledge not been studied in the literature, althoughseveral authors did propose ascending price combinatorial auctions in order to overcome the complexity of winner allocation in an revelation mechanism. This will be the topic of the next section.

Sealed Bid Combinatorial Auctions

In a sealed bid combinatorial auction bidders make bids for subsets of items. We assume for a moment that this gives a bid vector *w*. The auctioneer computes on this basis an allocation *x*. If the bids reveal the true valuations of bidders, i.e. $w = v$, then finding an optimal solution of the following integer linear program yields an efficient allocation:

$$\begin{array}{ll} \max & wx \\ \text{s.t.} & Ax \leq 1 \\ & x \text{ integer} \end{array}$$

In this formulation *A* is a matrix of zeroes and ones with a row for every item *i*, and a *1* at position *j* of row *i*, if and only if item *i* is contained in bid *j*. The right-hand side 1 stands for a vector with as many components as *A* has rows, with a *1* in each component. Two issues have to be considered at this point. The first is whether this is a reasonable way of representing the allocation problem. The second is how bidders can be given incentives to bid $w = v$. This is necessary to make the approach efficient, and due to a result of Monderer and Tennenholtz (Monderer, & Tennenholtz, 1998) it is a basis to design an optimal auction, i.e., an auction with maximum possible revenue for the auctioneer.

Coding and Computational Complexity of Sealed-bid Combinatorial Auctions

Let us look at the integer linear programming model above and discuss how well it is suited to find an allocation of subsets to bidders in a sealed bid combinatorial auction.

Firstly, we observe that several bids from the same bidder may win at the same time, as long as they are disjoint. This happens however only in the case of substitutes, since only in this case the sum of the values of two disjoint bids may be higher than the bid for the union of the two sets. To protect the bidder against winning several bids at the same time we have to add additional constraints to the linear programming model. If the bidder wants to win at most one bid from all her bids, one additional row per bidder suffices. It has a *1* in every column of the matrix that represents a bid by this bidder, and a *1* as right-hand-side, all other coefficients are equal to *0*. If a single bidder wants to express that there are several subsets of bids, where she wants to win one from each subset, but possibly several bids in total, we have to add such a constraint for each of these subsets. Such a bid is called an OR of XOR bids (Nisan, 2000). Notably, other logic constraints, like *if winning this bid I want also to win another bid*, can be expressed by adding linear constraints to the model. Nisan describes in a recent paper how different representations of combinatorial auctions can be converted into each other (Nisan, 2000).

Secondly, the simple encoding of bidders' types by listing *all* subsets for *every* bidder is not feasible for even small numbers of bidders and items. For 20 items and 100 bidders the matrix has already 20 rows and more than 20 million columns. It's not only the storage that causes a problem here, but also the fact that every bidder would have to communicate more than a million values. A sealed bid combinatorial auction has thus to restrict itself to a small number of bids from each bidder. Subsets for which a bidder does not make a bid are assumed to have a default value. For example, one might assume that such subsets have a value equal to zero. Many authors make also the free *disposal assumption* saying that getting additional items does not decrease the value. If the model does not exclude that the same bidder wins several subsets, the value of every subset that can be composed as disjoint union of subsets for which a bid has been made is assumed to be equal to the sum of the values of the subsets.

The restriction to a reasonable number of bids is less serious when a combinatorial auction is done in several rounds, versus in one round of sealed bids. In this case a bidder has a chance to submit in a later round bids which she did not consider in the beginning, and thus even learn her true valuation during the bidding process (Parkes, & Ungar, 2000). This can be seen as adding in every round columns to the integer linear programming problem, which relates nicely to a *column-generation algorithm* in combinatorial optimisation. In such an algorithm we start with a small number of columns and find the optimal solution (in our terms this means a small number of initial bids). We then "generate" new columns which, once added to the linear program, can improve the optimal solution. Note that every new column extends the number of possible solutions of the optimisation problem (all previous solutions plus solutions that use that column), thus with a column added the optimal solution is at least as good as before. In a multi-round combinatorial auction, column generation means that bidders can submit bids in every round which improve the revenue for the auctioneer. In a one-round sealed bid auction one can use a similar approach by letting bidders submit a *pricing procedure*. This is a software agent that is able to calculate for every subset the price that a bidder is willing to pay for that subset. An approach like that can be found in various online auctions for single items. It is called a *bidding agent*. A bidder parameterises this agent by her maximal price, and the agent will generate minimal incremented bids in an ascending price auction, until this price reached.

Thirdly, and most important the above integer linear programming problem is NP-hard. This means that there will be no algorithm that can solve this problem in a number of operations that is polynomial in the encoding length of the problem, as far as P≠NP, which, although never formally proven, is generally assumed to be true. This is already true if every bid contains not more than 3 items, and all bid prices are *1* or if every bidder makes exactly one bid with bid price *1* (Rothkopf, Pekec, & Hastard, 1998, van Hoesel, & Müller, 2000). On the other hand the problem is polynomial solvable if there exists a sorting of items such that every bid contains a set of neighboured items (Rothkopf, Pekec, & Hastard, 1998). This latter result holds only in the case of complements. If bids substitute each other, and therefore additional constraints are required telling that every bidder is allocated at most one bid, it becomes NP-complete. This follows from a result about the complexity of certain scheduling problems (Keil, 1992).

Complexity results for the winner determination problem have to be taken with some care. Note that we want to have an algorithm that is polynomial in the length of the encoding of a problem. Now, if every bidder would submit a bid for every subset, the encoding would already take a huge amount of space. Indeed, it can be shown that with this huge input the winner determination problem can be solved in polynomial time (Müller & Schulz, 2000). But with a large number of items, an exponential number of bids would not be realistic from a bidders point of view, therefore it is reasonable that in many applications we are not in this extreme case of an exponential number of bids.

The fact that the winner determination problem is NP-hard is the most serious obstacle towards the application of sealed bid combinatorial auctions, and will form the core of the discussion of the remaining part of this chapter. We will first explain the classical Vickrey-Clarke-Groves pricing mechanism for sealed bid combinatorial auctions. We give a very simple proof that it is truth revealing. If it is used in combination with a heuristic or an approximation algorithm it requires a warranty on the allocation algorithm that leads at the same time to some abnormal effects. These are recent results by Nisan and Ronen (Nisan, & Ronen, 2000), which we summarise without giving proofs. We will then illustrate that the complexity of computing a good allocation in a combinatorial auction has to do with the complexity of defining market clearing prices in a market with indivisible items. We finally give a framework to solve the allocation problem with a primal-dual algorithm. This framework computes such prices, however with the disadvantage of being not optimal. If it is combined with a carefully adjusted pricing mechanism, it gives a truth revealing combinatorial auction mechanism, however only for special cases.

The Vickrey-Clarke Groves Mechanism

A combinatorial auction is one of many mechanisms that could be thought of as a clearing mechanism for a market with indivisible items. A general analysis of such markets is far beyond the scope of this paper, so we will have to make some assumptions in this section.

Our first assumption is that the auctioneer has the goal to maximize own revenue. The best result that can be achieved, given certain types of bidders, and assuming that bidders do not bid more than their valuation, is the solution of the above integer linear optimisation problem with $w = v$. A bidder may however be better off in bidding less than her valuation. In order to achieve nevertheless an allocation that is close to this optimum the auctioneer has to set incentives to bidders to report their *true* valuation in their bids.

Take for example the case of a sealed-bid auction for a single item. Bidders submit sealed bids with a price w_j. If the auctioneer chooses the highest bid as winning bid, bid 1, say, and sets $p_1 = w_1$, then the corresponding bidder would have been better off if he had submitted a bid that is just above the second highest bid. Thus the auction mechanism sets incentives for strategic behaviour. However, a classical mechanism is available to avoid strategic behaviour. We have seen it already in the introduction. It is based on an appropriate modification of the payment scheme: winning bid 1 has only to pay the second highest bid, i.e., $p_1 = w_2$, say. Under this design, which is called the Vickrey auction, the best strategy for bidders is to reveal their true valuation (Vickrey, 1961). We observe that in a highly competitive market the auctioneer makes still almost optimal revenue, as the second highest valuation is likely to be almost equal to the highest valuation.

To put these ideas into formulas for the general case we need the following definitions. A set J_c of bids of a bidder *c* is said to constitute a *weakly dominant* strategy, if for every set of bids by other bidders any other set J_c' of bids from *c* would not improve her revenue. In other words, independent of what other bidders bid, the set J_c is always the best response. An auction mechanism is called *truth revealing* if for every bidder, bidding her true valuation is a weakly dominant set of bids.

Remember that we call an auction mechanism *efficient* if the total valuation of the allocation is maximal with respect to all feasible allocations, i.e. it maximises total welfare. Suppose we have a truth revealing mechanism and all bidders use their weakly dominant strategies. We will then have an efficient mechanism if the allocation *x* is an optimal solution for the above integer linear optimisation problem. Thus a truth revealing mechanism does provide the auctioneer with the right data to achieve efficiency.

For multi-item auctions the *Vickrey-Clarke-Groves mechanism* (Clarke, 1971, Groves, 1973), which is frequently also called the *Generalized Vickrey Auction*, is an auction design that satisfies truth revealing. It works as follows. Suppose *w* are the bid prices reported by bidders, and *A* is the matrix of bids, as described above. Let *x*' be the optimal allocation with respect to these bids, i.e. the optimal solution of the integer linear program given by A and *w*. The vector *x*' defines the winning bids. Next the auctioneer computes the prices for the winning bids. Let c be some winning bidder, and let x_c' be the part of *x*' related to bids from bidder *c*, and x_{-c}' the other components of *x*', and the same convention apply to other vectors and matrics involved. The auctioneer now deletes for a moment all bids from bidder c, and solves the integer optimization problem given by the objective $w_{-c}\, x_{-c}$ and the constraint $A_{-c}\, x_{-c} \leq 1$. Let z_{-c} be the optimal solution, in other words the best allocation of winning bids if bidder *c* would not have particiapted in the auction. The price *p(c)* to be paid by *c* is $p(c) = w_c\, x_c' - (w\, x' - w_{-c}\, z_{-c})$. The term in brackets that is subtracted from the bidder's bid price can be viewed as the marginal contribution of bidder *c* to the auction. This price is computed for every bidder who wins a bid under the allocation x_c'.

To give an illustration, consider our example from the previous section. We have seven bids, 3 from bidder *A*, and two from bidders *B* and *C*, respectively. The optimisation problem to solve is:

$$\begin{array}{ll}
\text{Max} & 100\, x_1 + 20\, x_2 + 20\, x_3 + 30\, x_4 + 30\, x_5 + 30\, x_6 + 30\, x_7 \\
\text{s.th.} & x_1 + x_2 + x_4 + x_6 \leq 1 \\
 & x_1 + x_3 + x_5 + x_7 \leq 1 \\
 & x_1 + x_2 + x_3 \leq 1 \\
 & x_4 + x_5 \leq 1 \\
 & x_6 + x_7 \leq 1 \\
 & x_1, x_2, x_3, x_4, x_5, x_6, x_7 \in \{0,1\}
\end{array}$$

The optimal solution is $x_1 = 1$, and all other $x_i = 0$. The price to be payed by bidder *A* is *100* minus her marginal contribution. The latter is equal to (100 – 60), as excluding every bid from A leaves us with the integer linear programming problem:

$$\begin{array}{ll}
\text{Max} & 30\, x_4 + 30\, x_5 + 30\, x_6 + 30\, x_7 \\
\text{s.th.} & x_4 + x_6 \leq 1 \\
 & x_5 + x_7 \leq 1 \\
 & x_4 + x_5 \leq 1 \\
 & x_6 + x_7 \leq 1 \\
 & x_4, x_5, x_6, x_7 \in \{0,1\}
\end{array}$$

which has an optimal value of *60*.Therefore *A* wins the set *{Q,R}* for a price of *60*.

We shall briefly give the proof for truth revelation. For that we have to compare a bidder's revenue for arbitrary bid prices *w* with the revenue when she changes to her valuation. With similar conventions as above, and an easy calculation, the first value is $v_c\, x_c' + w_{-c}\, x_{-c}' - w_{-c}\, z_{-c}$, and the second value is given, as $v_c\, x_c^t + w_{-c}\, x_{-c}^t - w_{-c} z_{-c}$, using x^t for the optimal allocation under the true valuation. The last term is identical in both formulas. The sum of the first two terms is in the first case equal to the objective value of the feasible allocation *x'* with respect to the objective composed from v_c and w_{-c}. In the second case the sum of the two terms is the objective value of the *optimal* feasible allocation under this objective. Therefore the value of the second formula is larger than the value of the first.

As VCG is truth revealing it is also efficient. And again it is likely to get very close to the optimum value for the auctioneer, when for example for every winning set there is a loosing bid on that set (or on a subset) that makes almost the same price offer. Truth revealing requires however to reveal the complete type, which might be information that is exponential in the number of items, as we discussed above. Another disadvantage of VCG is that it requires solving the winner determination problem once in phase 1 and again for each winner of a bid in phase 2. Although the later would be polynomial in the size of the input if all bids are reported, the overall mechanism of reporting bids and computing an allocation and prices needs an exponential number of operations.

One might argue that a good heuristic algorithm could replace the exact computation of the winner determination. Based on this idea, Nisan and Ronen defined *VCG-based mechanisms* (Nisan & Ronen, 2000). Such a mechanism uses the same second price formula as the VCGmechanism. However the vectors *x* and *z* are replaced by the solutions computed by the heuristic. Truth revealing VCG-based mechanisms can be characterized as follows (Nisan & Ronen, 2000). Let the *range* of an allocation mechanism be defined as the set of all different allocations that can be proposed by a mechanism. For example, a very simple mechanism might assign the whole set of items always to one bid and chose for that the bidder with highest bid for the set (or a subset, if we assume free disposal). In this case the range of the algorithm consists of as many different allocations as there are bidders, namely assigning everything to bidder 1, or to bidder 2, or to bidder 3, etc.. A sufficient condition for a truth revealing VCG-based mechanism is that it is *maximal in its range*. This means that the allocation algorithm chooses always an allocation that is at least as good with respect to the current bids as any other allocation in its range. If one looks at the proof above one observes that this property guarantees indeed truth revealing as it leaves no benefit for strategic behaviour. Furthermore, Nisan and Ronen showed that being maximal in the range is also *necessary* for a VCG based mechanism in order to be truth revealing.

From this result they make some astonishing observations. If the range does contain every allocation, then it is clear that the allocation algorithm has actually to find the optimal allocation. But this problem is NP-hard. If the range is however a subset of all allocations, then the algorithm must have some very counter intuitive behaviour. To see this, let *x* be an allocation that assigns some set S_1 to bidder 1, a set S_2 to bidder 2 and so forth, but that is *not* in the range of the algorithm. Whatever bidders will bid, the algorithm will not suggest this allocation. Then look at the following valuation: bidder 1 values all items in S_1 by *1*, bidder 2 values all items in S_2 by *1*, etc., while every bidder values all the other items by *0*. The allocation algorithm would in this case not do the obvious allocation *x*, namely giving S_1 to 1, S_1 to 2, etc..

Supported Allocations

It is important to mention that the direct equivalence of the winner allocation problem in sealed bid combinatorial auctions and the set packing problem from combinatorial optimisation gives us as well a huge family of polynomial solvable cases. For example, we may associate every bid with a node in a graph, give it the weight w_j, and connect two nodes by an edge if the bids compete on a common item. This leaves us with the problem to calculate a *maximum weighted stable set* in a graph. Combinatorial optimisation research has identified many classes of graphs for which the stable set problem is polynomial solvable (see, e.g., Skiena, 1998). In particular in cases where the bids have a geometric or geographic interpretation the graph is often a so-called *intersection graph*. In many such cases the geometric representation can be used to design polynomial algorithms (see, e.g., Felsner, Müller & Wernisch, 1997) .

Important, and this not only from the algorithmic perspective, are those cases for which winner determination is polynomial due to the fact that the linear program that we get when we drop the integer constraints from the above ILP has nevertheless an integer optimal solution. In this case we can use a polynomial algorithm to solve the winner determination problem, or to speak more pragmatically, every commercial linear programming solver can be used to compute the best allocation. Nisan observed that the fact that the LP has integer optimal solution is actually directly related to an economic property underlying the combinatorial auction (Nisan, 2000). This property is the topic of this section.

An allocation x in a combinatorial auction is *supported by single-item prices* y_i, if for all bids j, $x_j = 1$ implies $\sum_{i \in j} p_i \leq w_j$, and $x_j = 0$ implies $\sum_{i \in j} p_i \geq w_j$. It is *exactly* supported if $x_j = 1$ implies $\sum_{i \in j} p_i = w_j$. An allocation x is called *full* if every item is contained in some winning bid. From an application point of view a supported allocation has the advantage of publishable prices that explain to the bidders why they lost or won a bid. From a computational point of view it is of importance because of its relation to the solvability of the allocation ILP. It is an immediate consequence of linear programming duality that a combinatorial auction instance admits an optimal allocation x that is exactly supported by single item prices if the linear relaxation of the allocation ILP has an integral optimal solution. Indeed, the single item prices are equal to the dual variables, one for every item, and the dual feasibility together with the complementary slackness condition $x_i(\sum_{i \in j} y_i - w_j) = 0$ prove that the solution is exactly supported. The dual of the linear relaxation of the allocation ILP is

$$\begin{array}{ll} \min & y\,1 \\ \text{s.t.} & yA \geq w \\ & y \geq 0 \end{array}$$

Nisan showed almost the opposite direction as well. Namely, if a combinatorial auction instance admits a *full* allocation supported by single-item prices, then this allocation is optimal and the allocation LP has an integral optimal solution. Furthermore, it has a full allocation that is exactly supported (Nisan, 2000). Let us show the very simple proof of that result. Suppose x is a full allocation supported by single item prices y. First we may assume without loss of generality that the allocation is supported exactly, since for a winning bid j with $\sum_{i \in j} y_i < w_j$ we can increase y_i to make the left-hand-side equal to the right hand side (Note that every i is contained in at most one winning bid). After this modification, the prices y_i form a feasible dual solution. Let us now compare the primal and dual objectives of x and y, respectively. We see that $\sum_j w_j x_j = \sum_{i : xj = 1} \sum_{i \in j} y_i = \sum_i y_i$. The last equation holds because x is a full allocation, thus this condition is essential in the proof. By this equality we have a primal feasible x, and a dual feasible y, with identical objective value, proving the optimality of both.

There are many cases for which the allocation LP has an integral optimal solution, and is therefore supported (Müller & Schulz, 2000, Vohra & de Vries, 2000). In many cases the solution procedure does not have to rely on a LP solver, but could be done directly by a combinatorial algorithm, which is in most cases the faster approach. There are also classes for which one has a polynomial algorithm, but for which the LP does not necessarily have an integer solution. Here the allocation is in general not supported.

Primal-dual Algorithms

Primal-dual algorithms have particularly nice properties when applied to the winner determination problem in combinatorial auctions. Firstly, they compute solutions that are exactly supported. Secondly, for some special cases the payment phase can be adjusted in a way that makes the auction mechanism truth revealing. Thirdly, a primal-dual approach certificates the quality of the current allocation with respect to the criterion revenue maximization for the auctioneer. Indeed, the primal objective is a lower bound and the dual objective gives an upper bound on the revenue that can be achieved with respect to the current bids.

When adapted to our context a primal-dual algorithm works as follows. We are given the binary optimization problem to maximize the function wx subject to the intersection constraint $Ax \leq 1$ and the constraint that primal variables are binary. The dual program we look at is the dual of the linear relaxation of this binary program. It is given by the objective $y1$ and the constraint $yA \geq w$. We try to construct a feasible primal solution x, i.e. an allocation, and a dual feasible solution y, i.e., single-item prices such that x and y together fulfil the primal complementary slack constraint $x_i(\sum_{i \in j} y_i - w_j) = 0$. Together with dual feasibility, this slack condition translates to the condition that the allocation x is strictly supported by prices y .

Remember that a Dutch auction is an auction in which prices are descending. A *Dutch primal-dual algorithm* (Müller & Schulz, 2000) for winner determination mimics the principle of a Dutch auction. Single-item prices (dual variables) are iteratively reduced to levels at which the price of a bundle j, say, with respect to these single-item prices hits the bid price for the bundle. At such a level further reduction of prices of items in bundle j is not possible without becoming dual infeasible. The Dutch primal-dual algorithm decides for j whether it is assigned or not, fixes all prices of items in j and continues to reduce prices of other items. The algorithm terminates when no more dual variables can be reduced, or all bids are decided. The choice of variables, which are going to be reduced in the next iteration, and how much they are reduced, makes different versions of this scheme.

The following example illustrates the algorithm. There is a set of two items, *a* and *b,* and there are three bids: *12* for *{a}*, *10* for *{b}*, and *23* for *{a,b}*. Suppose we do a greedy allocation of bids, which means that we assign a bid whenever its dual constraint becomes tight, and all items in the bid are still available. We start with prices for *a*, and *b* to be equal to *23*. Prices will first go down to *12* for *a*, and *b,* the bid for *{a}* is assigned. Now, only the price for *{b}* can be further reduced. At price *11* the dual constraint on the set *{a,b}* becomes tight. This bid cannot be assigned. Now all prices are frozen and the algorithm stops. If we had not assigned *{a}*, then *{a,b}* could have been chosen. In both cases the dual objective equals the optimal solution, only in the second case the primal solution is optimal, too.

Lehmann, O'Callaghan and Shoham investigated the special case of combinatorial auctions in which every bidder is interested in at most one subset of items (Lehmann, O'Callaghan & Shoham, 1999). For this case they prove sufficient conditions for a combinatorial auction mechanism, i.e., a winner determination and payment algorithm, to be truth revealing. Key to their result is to adjust the payment scheme in a way that bidders cannot regret too high bids. The payment scheme is not equal to the VCG scheme. In order to explain their approach we need some definitions.

A combinatorial auction mechanism for single-minded bidders is called *exact*, if it assigns to a bidder either the empty set or exactly the set she is interested in. It is called *monotone* if winning bids would keep winning, if they would be replaced by a higher bid on the same set, or the same bid on a subset of items. In other words if a bid j with $x_j = 1$ is replaced by a bid k with $I_k \subseteq I_j$, and $w_k \geq w_j$, then $x_k = 1$. Lehmann et al. prove that in a mechanism that satisfies exactness and monotonicity, given a bid j together with a fixed set of other bids, there exists a *critical value* c_j such that for $w_j < c_j$ the mechanism will set $x_j = 0$ and for $w_j > c_j$ the mechanism will set $x_j = 1$. This motivates the following definition. An exact and monotone mechanism is called *critical* if $x_j = 1$ implies $p_j = c_j$. As a final condition in order to make a winner determination algorithm truth revealing they need *participation*. A mechanism fulfils *participation* if $x_j = 0$ implies $p_j = 0$. The result is then that a mechanism for single-minded bidders, which fulfils exactness, monotonicity, critical and participation is truth revealing.

The Dutch primal-dual allocation algorithm fulfils exactness. Furthermore it fulfils monotonicity. First, assume a bidder replaces a bid j by a bid k that bids a higher price on the same set. Then tightness of the dual constraint for *k* is not later achieved than it has been achieved for *j*. Second, suppose a replacement of a bid *j* by a bid *k* on a subset of its items. Then the dual constraint becomes tight at higher individual prices for each item in the bid, which again moves the decision about this bid to an earlier stage of the algorithm. But at earlier stages of the algorithm no bid *l* with $I_k \cap I_l \neq \emptyset$ can exist. From monotonicity it follows that there exists a critical price p_j for every bid. Using this in the payment phase of the mechanism guarantees participation and by that we get a truth-revealing mechanism.

The primal-dual view has recently proven to be rather helpful in understanding combinatorial auctions. In particular one can interpret ascending price auctions with combinatorial bids as primal-dual algorithms (Parkes & Ungar, 2000). Furthermore, the modeling of markets for indivisible items by integer linear programming models, their relaxations, and the dual programs of the later can lead to valuable insights in market design (Bikchandani & Ostroy, 2000).

Summary

This chapter has given an introduction on auction mechanisms as means of dynamic pricing on the Internet. We have explained how the popularity of auctions can be explained by general trends in the digital economy, like new forms of intermediation, customization, and customer involvement. We then outlined several design issues of an auction. Starting at simple single-unit, single-item auctions, we showed how in particular business-to-business applications require more complex auction formats. Such auction formats have to address the needs coming from complementarity or substitutability of items in a multi-item setting. Designing such auctions requires an integrated treatment of the computational aspects of auction design, i.e. how much information has to be communicated and how does the auctioneer process this information, and the economical aspects of the design, i.e., what are the strategies of bidders, what revenue can the auctioneer make, and is the allocation of items efficient with respect to the bidders preferences.

We could certainly not cover all aspects that should be covered if this chapter would claim to be a survey paper. However this has not been the intention. We rather want to encourage the reader to use the references of this chapter and dive into the theory. Touching this theory at least to some extent, rather than leaving it aside, had the intention to convince the reader that auction design is a far more complex engineering task than one might expect at first glance. The author maintains a Web portal listing resources on the subject at www.etrade.infonomics.nl.

Acknowledgements

I would like to thank Andreas Schulz for many fruitful discussions on the subject when I visited MIT in summer and autum 2000, and Vincent Feltkamp and Sander Onderstal for their valuable suggestions.

About the author

Dr. Rudolf Müller is Associate Professor for Operations Research at Maastricht University. He is program leader of the e-organisation research unit of the International Institute of Infonomics. He did completed his PhD in Mathematics in 1993 at Technical University Berlin, and has been Assistant Professor for IS at Humboldt-University in Berlin. His research interests include electronic markets, Internet information services, and combinatorial optimisation.

References

Agorics Inc. (1996). *Auctions – going, going, gone*. *A survey of auction types*. Los Altos, CA : Agorics Inc. Available at http://www.agorics.com/new.html (11.12.2000).

Barun Sarkar, M., Butler, B., & Steinfield, C. (1998). Intermediaries and Cybermediaries: A Continuing Role for Mediating Players in the Electronic Marketplace. *JCMC 1*(3). Available at http://www.ascusc.org/jcmc/vol1/issue3/sarkar.html (11.12.2000).

Bikchandani S., & Ostroy, J. M. (2000). *The package assignment model*. Working paper. An earlier version is available at http://www.cramton.umd.edu/conference/auction-conference.html (11.12.2000).

Brewer, P. J. (1999). Decentralized computation procurement and computational robustness in a smart market. *Economic Theory 13* (pp. 41-92).

Caplice, C. G. (1996). *An Optimisation Based Bidding Process: A New Framework for Shipper-Carrier Relationships.* PhD thesis. Department of Civil and Environmental Engineering, School of Engineering, MIT.

Clarke, E. H. (1971), Multipart pricing of public goods, *Public Choice*, *11*, (pp. 17-33).

Felsner, F., Müller, R., Wernisch, L. (1997). Trapezoid Graphs and Generalizations, Geomety and Algorithms, *Discrete Applied Mathematics 74* (pp.13-32).

Fujishima, Y., Leyton-Brown, K., Shoham, Y. (1999). Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of IJCAI'99*, Stockholm, Schweden. Morgan Kaufmann. Available at http://robotics.stanford.edu/~kevinlb/ (11.12.2000).
Groves, T. (1973). Incentives in teams, *Econometrica, 41*, (pp. 617-631).

Herschlag, M., & Zwick R. (2000). Internet Auctions – a popular and professional literature review, *Quarterly Journal of Electronic Commerce, 1*(2), 161- 186. Available at http://home.ust.hk/~mkzwick/download.html (11.12.2000).

Hoesel, S. van, & Müller, R (2000). Optimization in Electronic Markets: Examples in Combinatorial Auctions. *Netnomics,* forthcoming.

Keil, J. M. (1992). On the complexity of scheduling tasks with discrete starting times. *Operations Research Letters, 12*, (pp. 293-295).

Kennanvision.com (1998). *Exchange in the Internet Economy, The Keenan Report No. 1*. Available at http://www.keenanvision.com (11.12.2000).

Klemperer, P. (1999). Auction Theory – A Guide to the Literature. *Journal of Economic Surveys*, *13* (3), (pp. 227-286). Available at http://hicks.nuff.ox.ac.uk/economics/people/klemperer.htm (11.12.2000).

Lehmann, D., O'Callaghan, L. I., & Y. Shoham, Y. (1999). *Truth Revelation in Rapid, Approximately, Efficient Combinatorial Auctions*, Working Paper, 1999).

Leyton-Brown, K., Pearson, M., Shoham, Y. (2000). Towards a Universal Test Suite for Combinatorial Auctions. *Proceedings of the 2000 ACM Conference on Electronic Commerce (EC'00).* Available at http://robotics.stanford.edu/~kevinlb/ (11.12.2000).

Leyton-Brown, K., Shoham, Y., Tennenholtz, M. (2000). An algorithm for multi-unit combinatorial auctions. *Proceedings of National Conference on Artificial Intelligence (AAAI)*, Austin, TX, July 31-August 2. Available at http://robotics.stanford.edu/~kevinlb/ (11.12.2000).

McMillan, J. (1994). Selling Spectrum rights. *Journal Economic Perspectives 8* (3), (pp. 145-162).

Monderer, D., & Tennenholtz, M. (1998). *Optimal Auctions Revisited*. In Proceedings of AAAI-98.

Müller, R., & Schulz, A. (2000). *Combinatorial Auctions from a computational perspective*. Working paper. Available from the authors.

Myerson, R. B. (1979). Incentive Compatibility and the Bargaining Problem. *Econometrica 47*, (pp 61-73).

Myerson, R. B. (1981). Optimal Auction Design. *Mathematics of Operations Research*, *6* (1), (pp. 58-73).

Nisan, N. (2000). Bidding and Allocation in Combinatorial Auctions, *in Proceedings of the ACM Conference on Electronic Commerce (EC-00)*. Available at http://www.cs.huji.ac.il/~noam/ (11.12.2000).

Nisan, N., & Ronen, A. (2000). Computationally Feasible VCG Mechanisms, in *Proceedings of the ACM Conference on Electronic Commerce (EC-00)*. Available at http://www.cs.huji.ac.il/~noam/ (11.12.2000).

Ockenfels, A., & Roth, A. (2000). *Late Minute Bidding and the Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet.* Working Paper, Harvard University. Available at http://www.uni-magdeburg.de/vwl3/axel.html (11.12.2000).

Parkes, D. C., & Ungar, L. H. (2000). Iterative Combinatorial Auctions: Theory and Practice. In *Proceedings of AAAI-00*.

Parkes, D. C., Ungar, L. H., & Foster, D. P. (1999). Accounting for Cognitive Costs in On-line Auction Design. In Noriega, P. & Sierra, C. eds., *Agent Mediated Electronic Commerce (LNAI 1571)*, (pp. 25-40). Springer-Verlag.

Rassenti S. J., Smith V. L., Bulfin R. L. (1982). A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economis 13* (2) (pp. 402-417).

Rothkopf, M. H., Pekec, A., & Harstad, R. M. (1998). Computationally manageable combinatorial auctions. *Management Science*, *44* (8), (pp. 1131-1147).

Sandholm, T. (2000). Approaches to Winner Determination in Combinatorial Auctions. *Decision Support Systems*, 28(1-2), (pp. 165-176).

Sandholm, T. and Suri, S. (2000). Improved Algorithms for Optimal Winner Determination in Combinatorial Auctions and Generalizations. *Proceedings of National Conference on Artificial Intelligence (AAAI)*, Austin, TX, July 31-August 2 (pp. 90-97).

Scott, J. (2000). Emerging Patterns from the Dynamic Capabilities of Internet Intermediaries, *JCMC 5*(3). Available at http://www.ascusc.org/jcmc/vol5/issue3/scott.html (11.12.2000).

Skiena, S. S. (1998). The Algorithm Design Manual. Springer, The Electronic Library of Science Series.

Tennenholtz, M. (2000). Some tractable combinatorial auctions. *Proceedings of National Conference on Artificial Intelligence (AAAI)*, Austin, TX, July 31-August 2.

Vickrey, W. S. (1961). Counterspeculation, auctions and competitive sealed tenders, *Journal of Finance*, *16* (pp. 8-37).

Vohra R., & Vries, S. de, *Combinatorial auctions – a survey*, working paper, 2000. Available at http://www.kellogg.nwu.edu/faculty/vohra/htm/res.htm (11.12.2000).

Wellman, M., Walsh, W., Wurman, P., MacKie-Mason J. (1998). Auction protocols for decentralized scheduling. *Proceedings of the 18th International Conference on Distributed Computing Systems.*

Wolfstetter, E. (1996). Auctions: An introduction. *Journal of Economic Surveys*, *10*, (pp. 367-421).

MERIT-Infonomics Research Memorandum series
- 2001-

2001-016 **Auctions - the Big Winner Among Trading Mechanisms for the Internet Economy**
Rudolf Müller

Papers can be purchased at a cost of NLG 15,- or US$ 9,- per report at the following address:

MERIT – P.O. Box 616 – 6200 MD Maastricht – The Netherlands – Fax : +31-43-3884905
(* Surcharge of NLG 15,- or US$ 9,- for banking costs will be added for order from abroad)

Subscription: the yearly rate for MERIT-Infonomics Research Memoranda is NLG 300 or US$ 170, or papers can be downloaded from the internet:

http://meritbbs.unimaas.nl
http://www.infonomics.nl
email: secr-merit@merit.unimaas.nl