A bluffer's guide to auctions

Simon Parsons[†,‡], Juan A. Rodríguez-Aguilar[§] and Mark Klein[‡]

[†]Department of Computer and Information Science
Brooklyn College
City University of New York
2900 Bedford Avenue, Brooklyn,
NY 11210, USA
parsons@sci.brooklyn.cuny.edu

[§]IIIA, Institut d'Investigació en Intel·ligència Artificial,
CSIC, Spanish Scientific Research Council
Campus de la UAB,
08193 Bellaterra, SPAIN
jar@iiia.csic.es

[‡]Center for Coordination Science,
Massachusetts Institute of Technology,
3 Cambridge Center, NE20-336,
Cambridge,
MA 02142, USA
m_klein@mit.edu

October 11, 2007

Abstract

There is a veritable menagerie of auctions — single dimensional, multi-dimensional, single sided, double sided, first price, second price, English, Dutch, Japanese, sealed bid — and these have been extensively discussed and analysed in the economics literature. The main purpose of this paper is to survey this literature from a computer science perspective, primarily from the viewpoint of computer scientists who are interested in learning about auction theory, and to provide pointers into the economics literature for those who want a deeper technical understanding. In addition, since auctions are an increasingly important topic in computer science, we also look at work on auctions from the computer science literature. Overall, our aim is to identifying what both these bodies of work these tell us about creating electronic auctions.

1 Introduction

For years most people's view of an auction was that formed by films and television news. Auctions were the way that paintings by Picasso or Van Gogh were sold. They took place in stately halls where a man with an imperious manner stood on a stage at a lecture calling out fairy-tale prices, "Am I bid 20 million?, Thank you ma'am, 21 million?", to a hushed room. Ranged in front of the stage were rows of chairs occupied by nervous bidders. As the prices were called, the bidders nodded, or winked, or coughed and the auctioneer would acknowledge them with a nod of his own. If this was a news report, the painting would be sold to a Japanese bank and the sale would be proclaimed a record. If it was a film comedy, somebody would cough in the wrong place, or wave to a friend, and end up accidentally buying a painting they couldn't afford[1].

If one had more direct experience of auctions, it would most likely have been with an auction that was recognisably a variation of the film/TV version. However, in recent years, this has changed. Now many more people than ever before are taking part in auctions over the internet, both buying and selling goods. Even those of us who don't use the auction facilities provided by eBay or its comeptitore[2] may know of the auctions run by governments worldwide to allocate the radio-frequency spectrum to mobile phone operators.

Even before the number of auctions started to multiply, there was a considerably variety of different types of auction — as dicussed below the traditional auction used to sell fish in Spain and agricultural goods in the Netherlands differs from our parody example, and there are many other kinds of auction described in [14]. In recent years, this variety has increased, driven by problems with traditional kinds of auction, both in terms of the way that they end up deciding who gets what is being auctioned and how they might be implemented in settings such as the internet. Indeed, a veritable menagerie of auctions exist and crop up in various places — single dimensional, multi-dimensional, single sided, double sided, first price, second price, kth price, English, Dutch, Japanese, open-cry sealed bid, and combinatorial. This profusion can be overwhelming for a newcomer to the field, specially computer scientists drawn to the area by the

[1] Such events do not only occur in film comedies. As [14, page 151] reports:

> The former president of Parke-Benet reports that a dealer attending a sale of eighteenth-century French furniture had arraged to unbutton his overcoat whenever he wished to bid; buttoning the overcoat again would signal that he had ceased bidding. The dealer, coat unbuttoned, was in the midst of bidding for a Louis XVI sofa when he saw someone outside to whom he wished to speak and suddenly left the room. The auctioneer continued to bid for the dealer who, when he returned to the room, found he had become the owner of the sofa at an unexpectedly high price. An argument then followed as to whether an unbuttoned coat not in the auction room is the same as an unbuttoned coat in the auction room.

[2] [62] lists 142 such sites that were operational in the autumn of 1998, and despite the fact that some high profile sites have ceased trading — Yahoo auctions in the uS and Canada retired on June 3rd 2007 — a quick search suggests that there are even more online auctions running today.

increasing volume of computer science research in the area of auctions. The aim of this paper is to help such newcomers cope.

We start, in Section 2, by discussing the more prominent members of the auction family, and then, in Section 3 describe some of the theoretical work that has established the properties of, and relationships between, members of the family. Next, in Section 4, we consider two abstract views of the auction process (both drawn from work in implementing different types of auction), before identifying related work, in Section 5, and concluding.

2 The zoology of auctions

Auctions have existed, in some form or other for many years. If, following [30], we take an auction to involve the exchange of money and goods, then it is clear that auctions could exist until the invention of money (which both [13] and [60] place around 700 BC). However, auctions seem to have been in place in Greece by the 5th century BC [94], and later became widely used in both Greece and Rome [114] and Greece [94], while [38, pages 192–194] indicates that auctions were a staple of medieval trade around the Mediterranean in the middle ages. Given this long history, it is not surprising that many different forms of auction have been developed, and in this section we will present a classification, based on that of [106], from which we also borrowed the title of this section, and extended with some distinctions of our own.

2.1 A classification

We can identify a number of different possible features of auctions. The following is a set of features that are broadly recognised in the literature, although the names we have used are not universal — it is the types into which auctions can be divided rather than the terminology that is most important.

- Auctions can be single dimensional, or multi-dimensional.
- Auctions can be one sided or two sided.
- Auctions can be open-cry or sealed bid.
- Auctions can be first price or kth price.
- Auctions can be single unit or multi-unit.
- Auctions can be single item or multi-item (combinatorial).

Since these properties are basically independent — so that an auction that is single dimensional can be open-cry or sealed bid and single or multi-unit — they may be used to draw up what Shoham [103, 106] calls a zoological classification as in Figure 1. In many ways this is not a very useful means of classifying auctions — we will consider the much more useful one from [124] below — but it does allow us to identify some of the main distinctions that can be drawn.

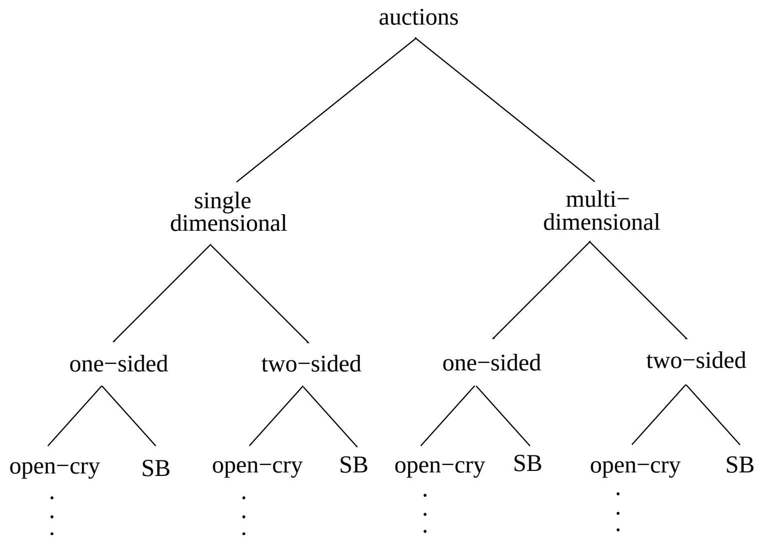

Figure 1: A zoological categorization of auctions

In *single dimensional* auctions, the only aspect of a bid that is important is the price offered for a good. In a *multi-dimensional* auction other aspects are distinguished. In a *one sided* auction, those bidding are either buyers or sellers — typically buyers — and the auctioneer has the job of deciding which is the winning bid. In a *two-sided* auction both buyers and sellers submit bids and the job of the auctioneer is to match buyers to sellers. In an *open-cry* auction, every bidder has access (can "hear") every other bid, whereas in a *sealed bid* auction, only the auctioneer has access to the bids.

In a *first-price* auction, the winning bidder pays the price of the winning bid, in a k*th price* auction the winning bidder pays the price bid by the bid that is ranked kth. The most familiar kind of kth price auction is the second-price auction where the bidder who bids highest (assuming the auction is ascending of course) pays the price bid by the second-highest bidder. In a single unit auction it is only possible to bid for a single good (though this may be a single item or a bundle of items grouped together — a box of fish, for instance or a computer along with software). In a multi-unit auction it is possible to bid for several units together, for instance in an multi-unit auction in which 20 boxes of cod are being sold simultaneously, it is possible to bid for, say, 10 boxes at $20 per box.

In a *combinatorial* auction, multiple, heterogeneous goods are auctioned simultaneously and bidders may bid for arbitrary combinations of goods. Thus, to extend the example we just used, a bidder may offer $100 for 3 boxes of cod and one box of halibut, while another may offer $120 for 2 boxes of cod and 3 of halibut. This kind of auction allows a bidder to express her preferences not just for particular goods but for sets or *bundles* of goods. A bidder's preference over a bundle of goods signals that a bidder's valuation for the bundle need not equal the sum of her valuations of the individual goods in the bundle — thus, to take an extreme example, our first combinatorial bidder might, if offered only cod or only halibut seperately, want neither.

This categorization allows us, for instance, to classify the "typical" auction described in the introduction as a single-dimensional, one sided, single good, open-cry auction. This, of course, does not classify it completely, and we shall discuss a more detailed description just below.

2.2 Some auctions, classified

Given the simple classification introduced above, we will consider the most common kinds of auction along with their position in this categorisation. The most familiar of these is the *English* auction — the standard auction-house auction described in the introduction. This is single dimensional since the only aspect of the good that is subject to bidding is the price. It is one-sided and typically sell-side — the seller may set a *reserve* price below which she will not sell the good and buyers make bids. It is open-cry since everyone is aware of every bid made, and it is first price since the winner pays the winning price. In addition, there are some distinguishing aspects that do not fit into the categorisation given above. For example, the bids are made in ascending order.

This is not, strictly speaking, a rule, but something that emerges from the other features of the auction. Since it is first price and everyone knows what the current winning price is, there is no point in calling a lower bid. As a result the protocol for implementing the auction (the way that the auction house works) rules out this possibility. In a traditional auction house this is achieved by having the auctioneer call out new winning bids and having bidders indicate that they are still interested in paying that price. This has another function[3] in providing a mechanism by which the end of the auction can be fixed. When the auctioneer is calling the bids, the end of the auction is when no bidder accepts the suggested price. It is also possible to fix a time limit for the receipt of final bids, either as an absolute time (2pm GMT) or relative to the time at which the last bid was received (5 minutes after the last bid is posted).

Possibly the next most familiar kind of auction is the *Dutch* auction. This again is a single-dimensional, one sided, open-cry, first price auction. The big difference between Dutch and English auctions is that the bids are descending rather than ascending. In a typical Dutch auction, the auctioneer starts at a price above that anyone is likely to pay, and then rapidly drops the price. The first bidder to accept the price gets the good. If the price falls below the reserve price, there is no sale. The Dutch auction is also known as the "descending clock" auction since, in practice, the auctioneer often indicates the price using a

[3]Indeed there is a third important function — that of encouraging bidders to keep bidding, and to ensure that bid increments are suitably large, both of which are to the advantage of the seller and the auction house, since the latter receives a fee based on the selling price. [14, pp 100–108] describes a number of mathods that auctioneers can use to manage this.

large clock-like dial (see [14] and, in particular, the photographs between pages 204 and 205).

Splitting the difference between English and Dutch auctions, we get *Japanese* auctions. [4]. Here the auctioneer calls out ascending prices, and bidders indicate that they are dropping out when the price gets too high for them. The last bidder to remain obtains the good at the price at which they become the last bidder. This makes the auction, once again, a single-dimensional, one sided, open-cry, first price auction. [74] describes an implementation of the Japanese auction (attributed to Cassady [14]) in which the rising price is displayed electronically and bidders hold down a button until they wish to withdraw from the auction.

A somewhat similar kind of auction is the *silent* auction which is commonly used in charity sales [73]. In a silent auction, bids are written down in open view along with the name of the bidder (typically next to the good being bid for — the goods for auction being laid out on tables around the auction room) and bidding closes at a predetermined time. Items are sold for the bid price to the highest bidder. Thus the silent auction is much like an English auction in form, but with many different goods being sold simultaneously. In fact silent auctions are almost identical to internet auctions, many of which operate under English auction rules.

All of the auctions described so far (including, rather confusingly, silent auctions) are open-cry, so every bidder knows the current best price for every good. While this makes it easy to implement the auction process, at least in a traditional auction house setting, open-cry auctions can be problematic. In particular, because of the openness, they can be disrupted in the sense of the outcome being distorted by some party who is trying to distort the outcome rather than participate in the auction. For example, in an English auction, a *shill*[5], working for the seller or the auction house, can take part in the auction to push the selling price up by bidding against a genuine bidder. Clearly there is no set of shilling tactics that are guaranteed to work[6], but the potential is obvious. Even if shills are not present, the way that bids are placed can be used to send signals as seems to have happened in the recent radio-frequency spectrum auctions in the US [23, 24]. This kind of activity, aimed at defrauding

[4] As Ausubel [2] points out, this is a misnomer, which Ausubel attributes to a misreading of Cassady [14]. [14, page63–64] uses the term "Japanese auction" to denote an auction whose

> distinctive aspect is that all bids are made by prospective buyers at the same time, or approximately the same time, using individual hand signs for each monetary unit. . . . The bidding starts as soon as the auctioneer gives the signal, and the highest bidder, as determined by the auctioneer, is awarded te lot.

In other words, the true Japanese auction, as used to sell fish in Tokyo, involves"simultaneous-bidding" [14, page 63], and so in theory will not be much different to a first-price sealed bid auction. As a result Ausubel suggests calling the kind of auction described here as "ascending-clock" rather than "Japanese". However, given the wide use of the name "Japanese" to describe this kind of auction, we will also use this terminology.

[5] Also called a "puffer" [14, page 212].

[6] The Auctionwatch site `http://www.auctionwatch.com/awdaily/tipsandtactics/buyshilling.html` gives tips for spotting shills which would seem to work pretty well as tips for shilling.

the seller and/or the auction-house is known as *collusion*.

A defence against this is to have private, sealed, bids. Such *sealed bid* auctions will set a deadline by which bids have to be received, and they are then "opened" by the auctioneer to determine the winner. Typically the highest bid wins, and the winner pays either the first or the kth price. Both first price and second price sealed bid auctions have been widely used and studied. Second-price sealed bid auctions are called *Vickrey* auctions after John Vickrey who first showed the advantage of the second price sealed bid auction [117]. Broadly speaking, the advantage is that making the winner pay the second highest price bid makes the auction proof against shilling and other forms of over-bidding. For more on this see Section 3.6.

Note that Vickrey auctions are not a panacea, as can be seen from the story of the New Zealand spectrum auctions. The New Zealand government adopted second price sealed bid auctions for their spectrum auctions with a number of embarassing results (taken from [70]). In one case a firm that bid NZ$100,000 paid NZ$6 since that was the second highest bid, in another the highest bid was NZ$7 million but the licence was sold for just NZ$5,000, and in a third case a bid of just NZ$1 was the winner and because there were no other bids the licence was acquired for nothing. The flaw here was that for a second-price auction to generate a decent revenue two or more bidders must value the licence relatively highly. Otherwise there is no guarantee that the outcome will extract anything like the value of the licence to the winner of the auction [72]. This illustrates the difference, which is not always clear, between mechanisms that aim to ensure that bids reflect the bidders' true value for the good in question, and mechanisms that aim to maximise revenue for the seller. Sometimes, as we will see below, these properties coincide, but this is not always the case.

2.3 Multi-unit, buy-side, and double auctions

As described all of these are single-unit auctions, and all can be made into multi-unit auctions by some adjustment of some part of the mechanism. Perhaps the simplest to adjust is the Dutch auction. For a multi-unit Dutch auction, the auctioneer will declare how many units are being auctioned and then proceed as before. Bidders simply call out a quantity rather than just indicating that they want the good. They then acquire the requested number of goods (up to the number being auctioned) and, if there are any unallocated units, the auction restarts. The restart can be from any price the auctioneer chooses, though typically it is from a price just above that at which it previously stopped, and the auction ends when either the reserve price is reached, or all the goods are allocated.

A multi-unit Japanese auction is created from a single-unit auction in much the same way as a multi-unit Dutch auction is, by allowing bidders to name the quantity of goods they wish to purchase. However, in order for it to work, the process has to change slightly. In a single unit Japanese auction, each bidder was implicitly in the bidding until they indicated their withdrawal. In a multi-unit auction each indicates, at every price increment, how many units they wish to buy at that price. Initially supply will be less than demand, but as the price rises, bidders are constrained not to increase the amount they wish to buy. When supply meets or exceeds demand, the goods are allocated. This is exactly the process described by Menezes [71] under the name "multiple-unit English auction".

Another form of the multi-unit English auctions is a little more complex. In this form buyers are allowed to name both the number of goods and the price that they are willing to pay for them. This means that when supply is no longer exceeded by demand, typically there will be a number of bids "on the table" each at a different price. The problem, then, is to decide how to decide on the price that bidders should pay given that they have bid different amounts. One solution is to get bidders to pay what they have bid, *pay-your-bid*, a form of *discriminatory* pricing — a name which covers all schemes where prices paid by the winning bidders are determined by the bids they made. Another option is to get bidders to all pay the same price, for example the lowest accepted price — a *non-discriminatory* or *uniform* pricing scheme[7]. A third option is for all successful bidders to pay the price bid by the highest losing bidder, a scheme which is a generalisation of the second-price auction. In general, multi-unit auctions are rather less well studied than their single-unit counterparts [3].

It is, of course, possible to have multi-unit sealed bid auctions, possibly the best known of these (though as [2] points out it is not widely used because of its perceived complexity) is the multi-unit Vickrey auction analysed in [63][8]. In this auction, each bidder submits a sealed bid which describes their demand curve — they specify how many units they are prepared to buy at every price. The auctioneer determines *clearing price*, the price at which demand equals the number of units for sale, and works out how may units each bidder will be prepared to buy at this price. (So far this process is basically the same as that for a multi-unit English auction.) Then the auctioneer determines the price each bidder pays, and this is set to be the price that would have been paid if the bid in question not have been made. In other words the bidder pays the price that is analogous to the second price in an auction for just one good.

All the auction mechanisms discussed so far have been sell-side auctions. That is the auction is set up from the perspective of the seller, and the buyers make bids. This is a natural set-up for auctions in which goods are distributed from single sellers to many buyers (or indeed from a small number of sellers to many buyers). Although not so common, it is perfectly possible to set up buy-side auctions in which a buyer receives bids from many sellers and picks a winner from whom she will buy. At the time of writing, the consumer credit markets in the United Kingdom and United States are, effectively, sealed-bid buy-side auctions. Any consumer of reasonable credit rating is bombarded by offers of

[7]An English multi-unit auction in which bidders pay the lowest accepted price is also, rather confusingly, sometimes called a "Dutch" auction, especially in the context of online auctions [62]

[8]As Varian [116] points out, this is a kind of "folk mechanism" — one that is widely known by economists, but doesn't seem to be written down anywhere and so isn't attributable to any one author.

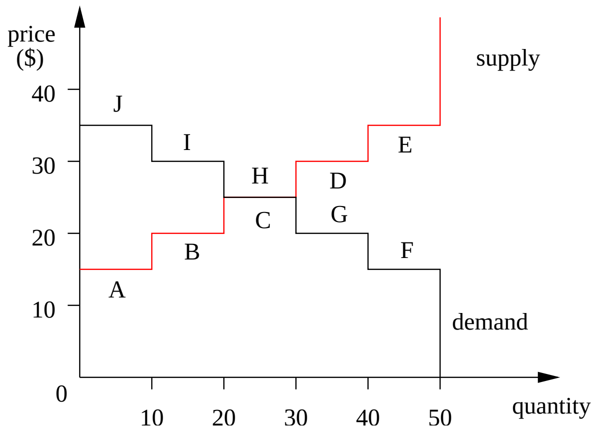

Figure 2: Illustrative supply and demand curves for a double auction.

credit — credit cards, loans, mortgages — from which particularly attractive options can be selected. Procurement auctions, which an increasing number of firms are using, are another form of buy-side auction. In this direction, a significant number of software tools to support buy-side procurement auctions are currently commercialised by several companies, including Ariba [9], Perfect [10], Emptoris [11], SAP [12], and iSOCO [13].

The important thing about buy-side versions of the auctions we have already considered, is that the price direction is reversed. An English buy-side auction, while still a first-price, one-sided, open-cry auction, now has descending prices, and the winner is the seller who is prepared to go lower than any other. Similarly, a buy-side Dutch auction will have ascending prices, with the sale going to the first bidder to accept the deal, while a buy-side Japanese auction will end when the price falls to a level that is acceptable to only one seller. To distinguish this feature of the offers to trade in buy-side auctions, we refer to them as *asks*, and use *offer* as a general term for something that is either a bid or an ask. An ask is a request for someone to buy some good from the asker at a given price. Again, we have first discussed the working of single-unit auctions — multi-unit auctions work much as they do on the sell-side, but just change the direction in which prices move.

Buy-side and sell-side auctions work as a means of matching, respectively, many sellers to one buyer and many buyers to one seller. What if we have many sellers and many buyers? Well, we put together a buy-side and a sell-side auction into one process to get a *two-sided* auction commonly known as a *double auction*. There is a range of different types of double auction [30].

[9] `http://www.ariba.com`
[10] `http://www.perfect.com`
[11] `http://www.emptoris.com`
[12] `http://www.sap.com`
[13] `http://www.isoco.com`

In the simplest kind of double auction, the *periodic* double auction — also known as a *call market* or *clearing house* — buyers and sellers are given a length of time to post their asks and bids. When the period is over, the auctioneer closes the market and constructs demand and supply curves from the pool of bids and asks, identifying how many units would be bought and sold at every price. The market is then cleared by matching buyers and sellers, and a clearing price is set on either a uniform-price or discriminatory price basis. The market may run just once in this way, or may re-open and repeat for some fixed number of iterations. Typical supply and demand curves will look like those in Figure 2[14]. Here the supply and demand curves intersect at a single price. It is equally possible for the intersection to be along a range of prices, and in such a case, the auctioneer must pick a price from this range. Whether the pricing is uniform or discriminatory, there will be a difference between the bid and ask prices, and the auctioneer has to choose how to set the price. A common rule for setting a uniform price is the so-called k-double auction rule [102], which sets the trade price as:

$$trade\_price = k \cdot bid\_price + (1 - k) \cdot ask\_price$$

where k is a number between 0 and 1, *ask_price* is the lowest ask that overlaps with a bid, and *bid_price* is the highest bid that overlaps with an ask. This rule is discussed more formally in Section 4.

A slightly more complex kind of double auction is the *continuous double auction* (CDA). This form of auction does not have to close to clear. Instead, the auctioneer immediately matches compatible bids and asks. Either perfect matches are found (where the number of units in the ask and bid are the same) or ones where, for example, the number in the ask is greater than the number in the bid. In the latter case after the bid is matched with some of the ask, the unsold units become another ask. The trade price is set in the same way as for a periodic double auction. The division into continuous and periodic is only one, coarse-grained, way to categorize double auctions, and some other variations are discussed in [90].

Common occurrences of continuous double auctions are stock markets — sellers look to sell blocks of shares at a particular price and buyers look to buy different sized blocks at another price — and these are implicitly multiple unit auctions. Although it is easy to see how this particular kind of multiple unit auction works, in general multi-unit auctions are hard to analyse, and theoretical results for single-unit auctions (which have been extensively analysed) do not seem to carry over to the multi-unit case too well. For example [71], ascending price multi-unit auctions do not seem to result in goods being sold at optimum prices from the point of view of the seller. (In contrast, in a single-unit English

[14] These supply and demand curves are taken from [19]. They are based on having traders A, B, C, D and E all looking to supply 10 goods which they will not sell for less than \$15, \$20, \$25, \$30, and \$35 per good respectively, and F, G, H , I, and J looking to buy at no more than \$15, \$20, \$25, \$30, and \$35 respectively. If the price is below \$15, no seller will be prepared to trade, while if the price is between \$15 and \$20 then only A will trade, and so on. Similarly, at a price above \$35 no buyer will trade, while at prices between \$30 and \$35, only J will trade, and so on.

auction the selling price is typically slightly higher than the winning bidder wishes to pay — see the discussion in Section 3 — and is therefore a good way to maximise revenue from the point of view of the seller.) This situation is analysed in [3], where it is proved that in uniform-price auctions, bidders who want to buy more than one item have an incentive to shade (reduce) their bid.

Despite these reservations, there are a number of multiple-unit auctions that have been proposed. One example, which is widely cited, is the ascending price multi-unit auction suggested by Ausubel [2]. The mechanism of this auction is easy to describe informally[15]:

> The auctioneer operates a continuously-ascending clock. For each price, p, each bidder i simultaneously indicates the quantity, $q_i(p)$ she desires, where demands are required to be non-increasing in number. When a price, p^*, is reached such that aggregate demand no longer exceeds supply, the auction is deemed to have concluded. Each bidder i is then assigned the quantity $q_i(p^*)$, and is charged the standing prices at which she "clinched" the respective objects.

Ausubel borrows the notion of clinching (which is defined formally in [2]) from discussions of sport. In the same way that a team playing in a league clinches a title when no other team can take the title instead of it, a good is clinched in an auction when no other bidder is prepared to pay enough to secure it. This makes the auction rather easy to understand, as well as provably efficient.

Another ascending price multi-unit auction that has received a lot of attention is the auction used by the FCC to sell spectrum right in the US. In order to avoid some of the problems of spectrum auctions in other countries (see Section 2.6 for a discussion of some of these), a new form of auction was invented, the *simultaneous ascending* auction. This is described in a number of places [21, 68, 73], and several of the authors of these papers had a role in the invention of the auction mechanism. The description given below is drawn from all of these.

The simultaneous ascending auction is an interesting mix of a number of features of auctions we have already considered. It is an auction for multiple items, and bidders are allowed to place sealed bids for any items which they are interested in obtaining. This is important given the nature of the items being auctioned — spectrum rights auctions sell the right to use a particular frequency band in a particular area, typically for mobile phone services. A single item does not typically provide enough geographical coverage for a mobile phone operator to establish a profitable service, so operators have to assemble a portfolio of items and so need to be able to bid on multiple items. Furthermore, it is important that items be auctioned simultaneously rather than sequentially, otherwise a failure in a later auction might make items won in earlier auctions worthless (if the key patch in a patchwork of coverage is not won then the other patches will not establish contiguous coverage and the service will not appeal

[15] This description assumes that the auction begins with more demand for goods than supply. This is easy to arrange if the price starts low enough.

to users). Furthermore, as Cramton [22] points out, when spectrum has been sold using sequential auctions, as it was in Switzerland, the results have been surprising and somewhat counter-intuitive. In the Swiss case, later auctions in the sequence raised much less revenue than earlier ones despite being for more resources (presumably because by the time of the later auctions, some potential operators had already been squeezed out of the market thus making the auctions less competitive).

Now, bidding in the simultaneous ascending auction takes place in rounds. At the end of each round the sealed bids are revealed, and the current winning bid and bidder are noted (the current winning bid is either the one from the previous round or a new, higher, bid). The reason for running the auction in rounds is again to enable operators to build sensible portfolios. With an auction that evolves over several rounds, operators can rebalance their bids if it turns out that the set of items they started out desiring are too highly valued by another bidder. In some variations of the auction, bidders are even allowed to withdraw bids as part of this rebalancing (though at a cost — if the item is sold for less than the withdrawn bid, then the withdrawer pays the difference between the withdrawn bid and the selling price). To reduce the number of rounds, there is a minimum bid for each round, which is made public along with the current winning bid and bidder identities.

The revelation of bidder identities is an interesting feature of the auction, especially in juxtaposition with the sealed bids during rounds. The reason for revealing bidders is that trials of the auction mechanism revealed that larger bidders seemed to be able to infer the identity of other bidders anyway, and so had an advantage. Revealing the information to everyone was a way of eliminating this advantage.

As pointed out in [97], each bidder participating in a simultaneous auction would prefer to wait and see until the end of the auction. In this manner, a bidder could observe the going prices and eventually optimise his bids to maximise his payoff. If every bidder waits and sees, there is a chance that no bidding starts. In [68] we find a patch to this problem, the so-called *activity rules* — rules on whether bids are accepted that are independent of the values of the bids. Before the start of the auction, every bidder has to indicate how much of the available spectrum they are interested in purchasing, and to pay a deposit for the right to bid on this. This establishes their "eligibility". Then, if a bidder makes a bid that exceeds its eligibility, the bid is rejected. Thus there is a limit to how much spectrum a single bidder can hold bids for, stopping anyone cornering the market for spectrum they do not intend to use (and, for instance, driving up the price for their competitors). In addition, a bidder has to be bidding (in the sense of making a new bid or holding the current winning bid from the previous round) on a given fraction of its eligibility at every round or else its eligibility will be reduced (though every bidder also has five "waivers" which it can use to avoid this reduction in a chosen round). This rule is an attempt to ensure that bidders are active and thus that the auction proceeds at a reasonable rate.

Once again, it is important to note that the use of a particular kind of auction is not a panacea — just adopting the simultaneous ascending auction will not guarantee success. As Cramton [22] points out, while the FCC auction and the simultaneous ascending spectrum auction in the UK are widely thought to have been successful, in terms of the revenue raised, the spectrum auction in the Netherlands was less so. This was despite being carried out under similar rules to those used in the UK. The problem seems to have been that the auction in the Netherlands was held under less competitive conditions — the number of licenses being sold was the same as the number of existing license holders. This discouraged anyone else from bidding to the extent that only one additional bidder joined the auction [58] and then quit early in the process.

2.4 Other kinds of auction

Ascending auctions can be considered to be strongly biased towards bidders who might only appear to have a weak advantage — for example [55] local knowledge of the area of a spectrum auction (see Section 3.4). To overcome this bias Klemperer [58] has suggested an alternative form of auction. This is the Anglo-Dutch auction which, as its name implies, combines aspects of both the English and the Dutch auction and was first proposed, but not named, in [55]. An Anglo-Dutch auction for a single good starts just like an English auction. The auctioneer runs an ascending auction in the usual way until all but two bidders have dropped out. The auction then changes. The two remaining bidders now make a sealed bid, a bid that is constrained to be not less than the price at which the last bidder(s) dropped out in the first stage of the auction. The winner is the bidder with the highest sealed bid and that is the price paid. The equivalence of the first-price sealed-bid auction with the Dutch auction (see Section 3) gives this part of the auction its name. An Anglo-Dutch auction is subtly different from the "Combination English-Dutch" auction described by [14, page 76–77]. This latter starts like an English auction, but the good is not awarded to the last bidder. Instead, the winning price of the English auction is used to set the starting price of a Dutch auction, and "the bidding starts downward from the highest price attained (in the English auction), with the prices in the descending phase including the maximum English-auction bid as well as the amount of the Dutch auction bid" [14, page 77]. In other words, the Dutch-auction stage starts from twice the price achieved in the English-auction phase.

The Anglo-Dutch auction works well when there is one or more bidder who is known to be so strong that other bidders would be unlikely to bid openly in a regular ascending auction[16]. Klemperer [58] argues that the sealed-bid stage introduces sufficient uncertainty that weaker bidders feel that they may be in with a chance of winning this final stage, and are therefore more likely to enter the auction to begin with. More bidders means that the price at the end of this

[16] Indeed this form of auction was originally suggested for the UK spectrum auction of 2000 which was originally going to offer four licenses in a market that had four incumbents. When the number of licenses was increased, the format of the auction was changed to a simple ascending auction [57].

first stage (and thus at the end of the auction as well) will be higher than if the auction had been a pure ascending auction. This leads Klemperer to suggest that the Netherlands spectrum auction would have raised more revenue had it been an Anglo-Dutch auction (a comment Klemperer made before the auction took place). Other properties of the auction are discussed in [58], and Goeree and Offerman [37] analyse a version of the Anglo-Dutch auction, which they call the "Amsterdam auction", in which the two final bidders receive a payment[17].

The aspect that distinguishes the Anglo-Dutch and simultaneous ascending auctions from the more traditional kinds of auction covered above is that it takes the form of a number of rounds, and some of the participants are eliminated between rounds. Another kind of auction with this general form is the *elimination* auction [32], also known as the *survival* auction. Indeed, the idea behind the elimination auction is almost exactly the same as that behind the Anglo-Dutch auction but has a different justification. The Anglo-Dutch auction combines ascending-bid and sealed-bid auctions to reduce the power wielded by strong bidders. The elimination auction combines ascending-bid and sealed bid to combine the information-revelation of the former with the guaranteed duration of the latter.

The way the two forms of auction are combined in the elimination auction is a little different from that in the Anglo-Dutch auction. In an elimination auction each round is sealed bid. After each round, the auctioneer eliminates some set of the lowest bidders and announces those losing bids along with a minimum bid for the next round. This, of course, defines a family of auctions, all of which can be shown to have outcomes that closely relate to a Japanese auction [32][18]

Amongst single dimensional models, we should briefly mention k-price auctions, where $k \geq 3$. Until recently these were thought to be of no practical use, but they seem to be useful in Internet auctions. As discussed in [75] and [76] (which also analyzes the equilibria attained in such auctions), it seems that, in general, buyers in internet auctions are risk seeking, and that for risk seeking buyers it is in the interests of the organizer of auctions to prefer k-price auctions with $k \geq 3$. In such situations their revenues are greater than for first or second price auctions. An example of this kind of auction is the 16th price auction used in the TAC Classic Trading Agent Competitions [43, 120, 121] involving software agents endowed with automated bidding strategies.

[17] As Goeree and Offerman explain, in a sense the Anglo-Dutch auction is a version of the Amsterdam auction since the latter has been used in Dutch and Belgian towns to sell real estate since medieval times.

[18] To be precise all members of the family are shown to be Nash outcome equivalent to Japanese auctions, meaning that if all bidders play their Nash equilibrium strategy — the strategy that will produce a Nash equilibrium [82, page 26] if all other bidders also bid according to their Nash equilibrium strategy, which in turns mean the first bidder can do no better if all others bid according to their Nash equilibrium strategy — the same bidder wins the good at the same price no matter which auction is used. Furthermore, the $n-1$ elimination auction — in which one bidder is eliminated at each stage so that an auction for n bidders will run for $n-1$ rounds — can be shown to be strategically equivalent to the Japanese auction, the strongest formal relationship that can be established between two auction mechanisms.

Finally, we should note that there are many other kinds of auction that we simply do not have space to consider. For example, Klemperer [56] describes the "all-pay" auction — which reflects lobbying competitions, in which every bidder pays their bid, but only the highest bidder wins — and Cassady [14, pages 71–74] describes the "handshake auction" and "whispered-bid" variants of the first-price sealed-bid auction. [78][19] comments that in Germany, antiques are often sold using an auction variant in which the top *two* bidders pay their bid, but only the top bidder gets the any return. This encourages very aggressive bidding, to the benefit of the seller, and apparently this kind of auction is known in Germany as an American auction. Another, odd, kind of auction is that held by Japanese major league baseball teams for the rights to players who want to leave for US teams (as was the case for Daisuke Matsuzaka [66] in late 2006). In this kind of auction, teams who want to sign the player make sealed bids for the rights to talk to the player, the highest bid wins, and the money goes to the team that the player will be leaving. However, if the bidding team fails to reach terms with the player in question, the money that was bid is returned.

2.5 Multi-dimensional auctions

For most of the above forms of auction, along with all the traditional forms of auction discussed in previous sections, the usual assumption is that the value which buyers and sellers place on goods is determined quite simply. Buyers have a value for goods and, broadly speaking, the more goods they can get the happier they are, and the cheaper the goods the happier they are. Sellers also have a value and the more goods they sell and the higher the price they sell them for, the happier they are. However, assuming that this is the case is a bit over-simplistic and more complex models of value, essentially multi-dimensional models, might be required. For example, in multi-good auctions, buyers and sellers may have definite limits on the number of goods they want to trade. As a buyer gets more goods of the same type, the value she places on each good may fall.

Even more complex situations may arise. Consider an agent bidding on plane tickets, for two tickets between New York and Boston. If the tickets are for a couple who want to spend a weekend in Boston together, there is no point in having just one ticket — the agent has to buy two tickets or none at all. Thus a single ticket has no value for the agent. Similarly, while getting one ticket on the Delta shuttle and one on the US Airways shuttle (or two tickets on two different flight on the same airline) is better than no ticket, it is not as good as two tickets on the same flight. Thus the way the agent decides the value of goods is a little complex. The same kind of *complementarities* between items occur in spectrum auctions — to put together a network an operator will need sets of licenses which fit together in a coherent way.

Such situations can be handled, as in the spectrum auctions, by having auctions over many rounds. This allows participants to adjust their bids on various

[19] Quoted at `http://www.magnolia.net/~leonf/sd/bargame9.html`.

items until they end up with a combination they are happy with. An alternative approach is to run *combinatorial* auctions [25], also known as auctions with *package bidding* [4]. In combinatorial auctions, participants are allowed to specify combinations of goods along with the price that they are willing to pay for them. For example:

$$\{2 \; tickets \; Delta \; shuttle\} : 300 \oplus \{2 \; tickets \; US \; shuttle\} : 200$$

indicating that the agent wishes to purchase two tickets on the Delta shuttle, and is willing to pay \$300 for them, or (in an exclusive-or sense, denoted by $\oplus$) two tickets on the US Airways shuttle but is only willing to pay \$200 for those. Other more complex kinds of bid may be possible (for instance giving combinations of values of different attributes[20]).

One obvious question to ask, given the advantages of combinatorial auctions, is why the FCC didn't use a combinatorial auction to sell spectrum. The answer, according to Milgrom [73] is that the auction designers were concerned about "free-riders". Consider a situation where three bidders are competing for two items. Bidder 1 wants licence A, bidder 2 wants licence B and bidder 3 wants both and bids only for the combination. A bid by bidder 1 on A that is relatively high compared with bidder 3's bid for both A and B helps 2 to obtain B relatively cheaply. While this can happen in any combinatorial auction, the problem in the FCC auction is that the multiple rounds mean that bidders can actively exploit this feature, identifying situations in which they can become bidder 2 and so choosing to free ride at the expense of bidder 1.

The problem of free-riding is not the only difficulty with combinatorial auctions. Another major problem with using combinatorial auctions is that determining the optimal allocation of goods (which bids to accept to generate the most revenue) is computationally intractable — Devries and Vohra [25] show that it is a variation on the set packing problem, and this is known to be NP-hard. As a result there has been considerable interest in trying to find special-purpose optimal algorithms, for example[98], as well as approximate solutions such as those provided by [39, 77, 127]. Another important issue that is the subject of ongoing research [10, 80, 81] is identifying a language in which to express bids. Such a language is required to enable agents to make compact statements of what they require rather than having to enumerate all bundles of goods that are acceptable.

However, the most important problem preventing the application of combinatorial auctions has to do with the complexity of bidding. Indeed, in combinatorial auctions the number of possible bundles is exponential in the number of goods. Hence, evaluating and submitting all possible bundles would be prohibitively time consuming both for the bidders and the auctioneer, who needs to solve the winner determination problem. In addition to the complexity of knowing their valuations over possible bundles, the bidders face the problem of deciding which bundles to submit and how much to bid [84]. Both the bid

[20] [105, 107] gives the example of bidding for a sofa and a chair where the requirement is that the colours match.

valuation and construction problem pose serious computational challenges to bidders, just as determining the winner does for auctioneers. We distinguish two strands of work focusing on bidding in combinatorial auctions. On the one hand, on a more theoretical side, some contributions [83, 122] assume that the bidders know their values for all possible bundles to subsequently consider a myopic best response bidding strategy where in each round bidders select new bundles to submit to maximise their utility given the current ask prices for bundles or goods. On the other hand, on a more practical side, further contributions [1, 110, 119] propose bidding strategies based on optimisation or heuristic techniques.

The use of combinatorial auctions is not limited to the simple sell-side auction that we used in the examples above. As described in [100], there is a wide range of combinatorial market designs — sell-side auctions, reverse (buy-side) auctions, and exchanges, with one or multiple units of each good, with and without free disposal[21]. [88] uses the term *combinatorial exchange* to denote a combinatorial double auction that brings together multiple buyers and sellers to trade multiple heterogeneous goods. Combinatorial exchanges combine and generalise two different mechanisms, double auctions and combinatorial auctions. As we have already discussed, in a double auction, multiple buyers and sellers trade units of an identical good, whereas buyers and sellers in a combinatorial exchange are allowed to both buy and sell bundles of goods, or to just buy bundles of goods or to just sell bundlles of goods. The example in [86] perfectly illustrates the differences between a double auction and a combinatorial exchange[22]:

> ...in an exchange for wireless spectrum, a bidder may declare that she is willing to pay $1 million for a trade where she obtains licenses for New York City, Boston, and Philadelphia, and loses her license for Washington DC.

As noticed by the authors of [86], a combinatorial exchange allows buyers and sellers to express more complex valuations as well as to buy and sell.

So far we have considered that bidders in combinatorial auctions submit bids for bundles of goods. Although this clearly allows them to better express their preferences, it many not suffice for some actual-world settings. Picture the following situation. An auctioneer is interested in obtaining cars, but is also in hearing from services to transform various components into a working car (at a certain cost), then it may solicit bids offering both ready-made cars, their individual components, and the services transforming components into cars. For this purpose, mixed multi-unit combinatorial auctions (MMUCA) [15] allow bidding agents to also bid for transformation services, that is an agent

[21]If the notion of free disposal is assumed, there is no requirement for all items to be sold (in sell-side auctions), or the auctioneer is permitted to balance the books by taking any extra items (in buy-side auctions).

[22]Notice though that sometimes we find in the literature, for example [125], that the notion of combinatorial double auction is used instead of combinatorial exchange, though it refers to the very same mechanism

may submit a bid offering to transform one set of goods into another set of goods. We should stress that there are important differences between mixed combinatorial auctions and double auctions or combinatorial exchanges. The most important difference is that these latter models do not have the concept of a *sequence of transformations*, which is required if the intention is to model some sort of production process. For instance, if an auctioneer is intended to obtain a cake and receives a bid to transform dough into a cake, the bid is not of much use unless some other bidders provide the service to transform flour into dough, and yet others provide some flour to start the production process with. Furthermore, notice that MMUCAs must be regarded as a generalisation of the model of combinatorial auction for supply chain formation presented in [5]. As one might imagine, despite the generalisation introduced by MMUCAs, winner determination remains NP-hard. Indeed, the winner determination problem poses particular computational challenge since the number of variables of an Integer Program — which will provide an exact solution to the optimum allocation — grows quadratically with the number of transformations mentioned in the bids. Therefore, the practical application of MMUCAs necessarily requires either introducing reduction techniques [35] or special-purpose algorithms.

As noticed above, the main issue when using combinatorial auctions is the winner determination problem, namely assessing the optimal allocation of goods. Nonetheless, in most real world settings there are further considerations besides maximising economic value. These considerations are usually business rules that are specified as a set of side constraints (such as budget, limit on the number of winners, or exclusivity among bids) that need to be satisfied while picking a set of winning bids. As noticed in [54, 99], the complexity of winner-determination can also be impacted by side constraints that represent business rules. Moreover, the addition of side constraints may create another interesting situation, an auction with no winner(s). This would be the case if the winner-determination cannot find any set of bids that comply with side constraints.

Multi-attribute auctions [8, 87] are another type of multi-dimensional auctions allowing entrants to bid over qualitative, non-price, attributes such as delivery time, weight, payment terms, as well as price. Borrowing the example from [27], in a multi-attribute auction for computers, the good may be defined by attributes such as processor speed, memory, and hard disk capacity. Bidders may have varying preferences (or costs) associated with the possible configurations. For instance, a buyer may be willing to purchase a computer with a 2 GHz processor, 500 MB of memory, and a 50 GB hard disk for a price no greater than \$500, or the same computer with 1GB of memory for a price no greater than \$600. As is the case for combinatorial auctions, multi-attribute auctions are intended to lead to higher market efficiency by providing more information on buyers' preferences and suppliers' offerings. As noticed by Engel et al. in [27], research in multi-attribute auctions has mainly focused on one-sided mechanisms, which automate the process whereby a single agent negotiates with multiple potential trading partners. Typically there is a buyer with some value function, v, ranging over the possible configurations K, and a set of sellers associated with cost functions $c_1, \ldots, c_n$ over K. The role of the auction is to elicit these functions (possibly approximate or partial versions) to identify the surplus-maximising deal as a result of solving

$$\arg\max_{i,x} v(x) - c_i(x)$$

In [17], bidders bid on both price and quality, and bids are evaluated by a scoring rule designed by a buyer. Analogs of the classic first- and second-price auctions correspond to first and second-score auctions described in [11, 17].

One-sided multi-attribute auctions present the auctioneer with a winner determination problem that is just as hard as in combinatorial auctions, and solving this problem presents a considerable research challenge [9]. Note that the techniques reported in [9] do also consider side constraints on the combination of trades comprising an overall deal along the lines of the side constraints considered in [54, 99]. In addition to optimising winner determination, the main issues arising when considering multi-attribute auctions have to do with their optimality and efficiency[23] compared to single-attribute auctions. Hence, it is obvious to pose the following questions:

1. Do multi-attributes auctions lead to higher utilities than single-attribute auctions?; and

2. Are multi-attribute auctions more economically efficient than single-attribute auctions?

The answer to these issues will indicate whether multi-attribute auctions are worthwhile. In [7], Bichler provides empirical answers to both issues. In his experiments, he shows that the utility scores achieved in multi-attribute auctions were significantly higher than those of single-attribute auctions. Nonetheless, the efficiency was similar in single-attribute and multi-attribute auctions, not finding any evidence for revenue equivalence between the multi-attribute auction formats. As noted in [54], in general, following the Myerson-Satterthwaite impossibility result there can be no economically efficient multi-attribute auction that does not run at a deficit because there is private information on two-sides of the market (for the buyer and for the sellers)

Multi-attribute auctions become more complex if we consider double-sided protocols where multiple buyers and sellers submit bids, and the objective is to construct a set of deals maximising overall surplus. We can distinguish two types of double-sided multi-attribute auctions: continuous double auctions [29, 40] and call markets [27]. Both [29] and [40] consider a matching problem for CDAs, where deals occur whenever a pair of compatible bids is found. In a call market, as in [27], the auctioneer collects bids until some deadline, be it periodic or fixed, at which he clears the auction by assessing a match over the entire set of bids. The complexity of clearing the market for the auctioneer is rather different for

[23] The reader should notice that concerns about *efficiency* and *optimality* are at the heart of auction design. Efficient auction design is concerned with how the surplus in an auction is divided among the bidders and the auctioneer, while optimal auction design concentrates on designs which maximise the expected revenue of the bid taker.

both auction types. On the one hand, clearing a multi-attribute CDA is similar to clearing a one-sided multi-attribute auction since the problem is to match a new incoming bid with the existing bids on the other side. On the other hand, multi-attribute call markets are argued to be much more complex because constructing an optimal overall matching may require to consider multiple trade combinations from all buyer-seller pairs. Despite the higher complexity inherent to their clearing algorithms, multi-attribute call markets are said to count on liquidity and efficiency advantages over multi-attribute CDAs [26]. Motivated by the higher complexity of call markets, Engel *et al.* [27] provide an expressive bidding language to subsequently develop network flow models that allow to study the clearing problem in call markets, and thus provide guidance for implementing optimisation algorithms.

Finally, combinatorial multi-attribute auctions have been investigated in the context of procurement auctions [36, 113]. In such procurement scenarios, each item is defined by several attributes, the buyer is the auctioneer, for example the government, and the sellers are the bidders. Furthermore, the auctioneer requests multiple items and both buyer and sellers can have arbitrary, that is substitutable or complementary, preferences on a bundle of items. The following example from [113] perfectly pictures the kind of scenarios wherein combinatorial multi-attribute auctions may be employed:

> For example, a task of constructing a large building can be divided into many subtasks. One constructor might be able to handle multiple subtasks, while another company is specialized to a particular subtask. In addition, since each constructor may contract processes under different conditions, i.e., their quality, appointed date, price and so on, the utility of the government may depend on these conditions in a complex fashion.

On the one hand, the work in [113] focuses on efficient protocols for combinatorial multi-attribute auctions. Therefore, this theoretical work must be regarded as a step beyond the work by [17], which did not treat multiple tasks (items). On the other hand, on the application side, Giovannucci *et al.* [36] develop an agent-aware decision support service for negotiation scenarios that operates as a winner determination solver for combinatorial multi-attribute auctions including side constraints.

2.6 Examples of real auctions

All the four basic kinds of auction — English, Dutch, first-price sealed-bid and second-price sealed-bid — have real life examples (many of which are given in [56]). The English auction is a common means of selling antiques and artwork. Dutch auctions are commonly used to sell perishable goods — flowers in the Netherlands, fish in Israel and Spain, and tobacco in Canada. First-price sealed-bid auctions are often used by governments to sell mineral rights, and are used to sell houses in Scotland. The UK Treasury sells securities using a first-price sealed-bid auction, albeit a multi-unit version, and the US Treasury used to use the same mechanism, having recently changed to the second-price sealed-bid auction. The latter is also used for most auctions of stamps by mail.

Apart from the kind of auction caricatured in the introduction, probably the most common kind of auction until very recently was the kind that one finds in stock markets. As mentioned above, stock markets are double auctions, bringing buyers and sellers of stock together to agree a price at which to trade[24] Typically, of course, stock markets are multi-unit auctions since few such deals are over single share. The classic view of the daily operation of the London and New York Stock Exchanges is that individual types of stock are traded in a continuous double auctions, so that the whole exchange is a set of parallel continuous double auctions in which individuals known as *specialists* organise trade in particular goods [30]. In addition, prices at the start of trading periods are set by a series of call markets, one for each stock.

Looking a little more closely [48], the situation in the New York Stock Exchange (NYSE) is a little more complex. At the time that [48] was written, in the early 1990s, some trade was carried out directly between brokers looking to buy and sell stock on the part of their clients. This was done exactly as [96] describe The Chicago Board of Trade operating (though the latter was a market for commodities rather than stocks), with traders literally calling out bids and asks until they identified a trading partner. In additon to this, the specialists operated as "market makers", guaranteeing to provide a market to sellers and a supply of stock to buyers. The proportion of trade carried out directly between brokers as opposed to that carried out through specialists depended on the stock — specialists tended to play a greater role in less heavily traded stocks. More recently, as a result of the ever-increasing volume of trade carried out electronically, the volume of trade moving through specialists is diminishing and the market as a whole is moving closer to the experimental markets studied in the economics literature, for example in [108] [25].

A slightly different mechanism, that of the Wunsch auction, is used by the Arizona stock market [69] for all trading, and by the NYSE for trading securities after the exchange has closed [30]. The Wunsch auction is a uniform-price variation on the call market which aims to gather many bids and asks and then compute a trade price which creates the maximum volume of transactions. The advantage of the Wunsch auction over a call market, the kind of market used by the NYSE to set starting prices after a period in which the market has not traded, is that prices are less volatile [69][26]. Of course [31], the NYSE operated as a call market before the 1860s, and only switched to the continuous form of

[24] Cassady [14, pages 13–14] suggests that double auctions might better be classified as multiple negotiations rather than auctions, but more modern writers do not seem to agree with this observation.

[25] Note that some markets, like NASDAQ, do not seem to have buyers and sellers dealing directly with one another, instead requiring that all transactions go through a market-maker.

[26] An auction that is similar to both the Wunsch auction and the call market is that used to determine the offer price of Google shares [47] in the first offering of Google stock (though this was confusingly described in the media as a "Dutch" auction). In the Google auction bids were collected over time, used to construct a demand curve, and this was then used to identify an offer price.

the auction in response to rising trade volumes.

Another kind of auction that has become very common is the internet auction. Of course is really a broad class of auctions, distinguished by the fact that they are carried out entirely electronically. Indeed the most obvious kind, those like eBay in which individuals take part, are a very small fraction in revenue terms of the whole, though they are responsible, as suggested above, for introducing many of us to the idea of taking part in auctions ourselves.

In fact the idea of personal auctions pre-dates sites like eBay by several years, and the earliest internet auctions really pre-date the World Wide Web — as Lucking-Reiley [62] points out, the earliest internet auctions were carried out on newsgroups and through email discussion lists. However, the real growth was sparked by Web-based auction sites like onSale (which launched in May 1995) and eBay (which launched in September of the same year), and which has run alongside the general growth in the use of the internet worldwide.

Between them onSale and eBay provide examples of the two main business models for internet auctions. When it launched onSale was an example of a *merchant site*, a retailer that explicitly allows buyers to set the price of transactions[27]. At the time of writing onSale has become a simple retailer of consumer electronics, but Priceline is an example of a site that stll holds to the original onSale model. eBay, on the other hand, is an example of a *listing agent* site which makes it possible for anyone to auction anything they choose to try and sell. While onSale and Priceline make money in the same way that retailers traditionally do, eBay makes money as auction houses traditionally have, by charging a commission on the sale price. Anyone who has been to sales of used household items (yard-sales in the US, car-boot sales in the UK) will not be surprised that eBay proved to be a popular place to unload unwanted possessions, but the magnitude of the market is surprising. [62] gives eBay's monthly revenue (sales) for August 1998 to be \$70 million and gives an estimate for Summer 1999 of \$190 million, and also estimates the volume of transactions across all internet auctions in 1998 to be \$1 billion. Meanwhile, Shoham [104] gives the number of new items registered daily in October 2000 on eBay as 600,000. While these figures are impressive, they are dwarfed by more recent data. eBay reports that in the first quarter of 2007, there were 588 million listings (over six million a day), while the gross value of all merchanise sold through the site over the same period was \$14.28 billion.

From the viewpoint of auction theory, one interesting aspect of internet auctions is their susceptability to sniping[28], which, as discussed in Section 3.6, subverts the intended format of an ascending bid auction (albeit one that is basically the same format as a silent auction) and turns it into a first-price sealed bid auction, as well as possibly annoying bidders who have played by the

[27] One might argue that all online retailers, especially in a world of smart consumers who make use of shopbots [44], are engaged in a sell-side auction in which they compete for buyers with other retailers.

[28] The practice by which bidders silently observe the auction and then place a bid just above the previous highest bid so close to the end of the auction that the now-usurped winner has no time to respond.

rules. As a result auction sites have developed a couple of anti-sniping strategies. The first is to have an automatic extension to the auction, for a period of several minutes, if there is any last-minute bidding. While this prevents sniping, it also means that bidders have to keep checking for such extensions and maybe bid again (thus increasing the similarity to the silent auction). Another alternative is to provide a proxy bidder — an agent that will ensure one has the highest bid for a good until a given limit is reached. A detailed analysis of last-minute bidding and the effect of the various responses to it is given by Roth and Ockenfels [95].

Another kind of auction that now has a significant turnover is the business-to-business (B2B) auction. Some of these auctions are conducted over the internet — for example Freemarkets Online is cited by [62] as allowing firms to run procurement auctions. Such auctions can be substantial. The General Electric Corporation (GE) of the USA, for example, purchased over \$6 billion worth of goods and services via on-line auctions in 2000 [33], and one can imagine that this figure has grown just as those relating to eBay have. There are also many such auctions. According to a PricewaterhouseCoopers report, there were over 1000 public eMarkets and around 30,000 private electronic exchanges at the start of 2001 [93], many of which use auction mechanisms to match buyers with sellers. These figures are backed up by Lucking-Reiley's survey [62], which looked at just 142 auction sites which between them generated nearly \$100 million a month, and again it seems likely that these figures have grown in the intervening period.

One concrete example of a class of a large buy-side auction market is that for spectrum licenses — we have already discussed this in various respects, but it is worth stressing the size of the the markets. For example, the total revenue raised in the nine spectrum auctions carried out in Europe in 2000 and 2001 was around \$100 billion, of which the largest single part was raised by the auction in the UK. That raised a remarkable 39 billion Euros, and a similar amount was raised in the German spectrum auctions [58]. Even the relatively unsuccessful auctions in the Netherlands raised 100 euros per head of population (as opposed to around 600 euros per head in the UK).

Another example, this time of large B2B double-auction markets, are those for electricity distribution. Of these the California power markets make for an interesting comparison with the spectrum auctions discussed above. While the latter have broadly been a success, with only a few auctions failing to raise reasonable revenues given the licenses on offer [57], the initial California power markets are almost universally considered to have operated in an unsatisfactory manner. This view is based upon aspects such as the frequent price spikes seen across the summer of 2000 and the 2000-2001 winter period [53], spikes that ultimately caused blackouts.

3 Analysing auctions

Given this wide variety of different types of auction it can be hard, as someone considering holding an auction, to decide what type of auction to conduct. There are a number of considerations, and we will discuss some of them in this section.

One approach to compare different auctions is to perform a theoretical analysis and use this to answer questions such as "which kind of auction will maximise the profit for the seller", or "which kind of auction will maximise the number of items sold" — both of which are relevant in different circumstances[29]. Since the outcome of any auction depends upon the way in which participants bid, any analysis will have to use some model of this behaviour. Here we discuss three such models, the *independent private values* model (which leads to revenue equivalence), Milgrom and Weber's *correlated values* model [74], and Klemperer's *almost common* values model [55]. We also cover the idea of *revenue equivalence*, which links different kinds of auction, and touch on the *revelation principle* and other aspects of basic *mechanism design*.

We start with the independent private value model, and our description of it and the results that follow from it are drawn from [74] which gives a deeper and more technical discussion of many of the points made below.

3.1 Independent private values model

In the independent private values model, many buyers bid for a single, indivisible, object. Each of the buyers is risk neutral — that is they bid to exactly their value for the object — and they have an accurate value for the good which is unknown to other bidders. This is the sense in which the value is private. Considering the whole set of bidders, each can be thought of as having been assigned a private value which is independently drawn from some distribution of values. Thus the value that one bidder has for the object does not affect, in a statistical sense, the value that another has. This is the sense that the value is independent. This model certainly seems reasonable (at least until reading Milgrom and Weber's critique) and leads to some interesting conclusions.

The first conclusion is that there is a form of equivalence between Dutch and first-price sealed-bid auctions. As Vickrey [117] points out, a bidder with a particular value for the object has a very simple strategy in a Dutch auction — claim the object if and when the price drops to the value she has for the object. With a set of bidders doing this, the winner is the one who values the good most, and the winner pays this value[30]. Now, exactly the same strategy will work for every bidder in a first-price sealed bid auction, and for a given set of bidders, the same bidder will win the object at the same price in a first-price sealed bid auction as in a Dutch auction.

[29] The first question is clearly of interest in attempting to raise money by selling the family silver, the second is of interest in situations like the sale of treasury bonds.

[30] In passing we should note that this outcome, that the object goes to the bidder who values it the most, is a good property of any auction — it ensures that the mechanism is "efficient" in economic terms, since it prevents goods ending up with buyers who value them less than others who do not have them. In an inefficient outcome, there is revenue that could have been extracted were the mechanism to have ensured that the goods passed to buyers who valued them more.

The second conclusion is that there is another kind of equivalence[31] between English auctions and second-price sealed bid auctions. In a second-price sealed bid auction, if a bidder knows her value for the good, which we assume is the case in the independent private values model, then she submits a bid for this amount. If this is the highest bid, then she wins and pays the amount bid by the second highest bidder. Similarly, in an English auction a bidder who knows her value for the object can just keep bidding up until the price reaches the value of the object to her. If all bidders follow this strategy, then the bidding will stop when the price reaches the value that the bidder who values the object second-highest places on the object, though the bidder who values it the most may have to pay a tiny amount more than this (that is in practice there may be some minimum increment by which the winner has to pay over the second-highest bid). Thus the two auctions will have effectively the same outcome.

The third result is that the outcome of English and second-price sealed bid auctions is Pareto optimal because the winner is the bidder who has the highest value for the object. As described (and obviously from the description given) the outcome of Dutch and first-price sealed bid auctions is Pareto optimal as well, though it will not be if the model is not symmetric (see below). A fourth result is that all four auctions give almost exactly the same expected revenue for the seller, and so are equally good choices from the point of view of selling the family silver. This last result depends on the assumption that there are many bidders, in which case, with private values being drawn from the same distribution, there will not be much difference between the highest and second-highest valuations for the good.

The final result that we will discuss (Milgrom and Weber [74] give seven) holds where bidders are not risk neutral, but instead are risk averse — that is they tend to bid below their actual valuation for the object. In such cases Dutch or first-price sealed bid auctions will generate higher revenues for the seller.

Now, several of these results rely on the assumption that values are known to the bidders. This is reasonable in may situations, for example in the auction of many consumer goods, but is much more likely to be violated in auctions such as those for mineral or spectrum rights where the companies concerned can only estimate, possibly inaccurately, the value. It seems reasonable, in such situations, to assume that the value of the object being auctioned has more or less the same value to every bidder, so that this is called the *common values* or *mineral rights* [74] model, but that their estimates of this unknown value vary. Now the winner in the auction will be the bidder who values the good the most and since this will be a bidder who has an extreme estimation (extreme in the sense that it is a point from the high end of the distribution of estimated values), this will, in general, be an overestimate of the true value[32]. This phenomenon is known as the *winner's curse*, reflecting the fact that the winner is typically paying too much.

[31] The kind of equivalence that Milgrom and Weber [74] state holds here is weaker than that between Dutch and first-price sealed bid auctions.

[32] For the winner to have underestimated the value, all other bidders must also have underestimated it and by more.

As Milgrom [72] points out, recognising that the winner's curse exists leads a rational bidder to want to shade their bid — knowing that if their bid is the highest it will, on average, be because they have over-estimated the value of the good. If all bidders know this, then auctions which set the sale price at the highest price bid will typically attract lower highest bids than those which use the second-price as the sale price. We will discuss the winner's curse more below.

3.2 Revenue equivalence

The result discussed in the previous section, that the four main types of auction — English, Dutch, first-price sealed bid and second-price sealed bid — all generated the same expected revenue, is sufficiently important that it deserves a little more discussion. For a start, this result has a name, and it is known as the *revenue equivalence theorem.* Furthermore, this theorem requires a model that has four assumptions (as classified by [67]):

1. The bidders are risk neutral
2. The independent private values assumption holds[33]
3. Bidders are symmetric; and
4. Payment is a function of bids alone.

The first two of these were, implicitly at least, discussed above. The last just says that the auction itself determines the exchange of goods and money (there are no additional payments between bidders, as for example when bidders collude to obtain a good which they then allocate among themselves). The assumption about symmetry is that all bidders are basically the same (technically, their private values are assumed to be drawn from the same distribution), so that the revenue equivalence result would not hold for bidders drawn from two distinct populations[34].

Looking at the situation in which these assumptions hold — a situation McAfee and McMillan [67] call "the benchmark model" — we find that a bidder in English and second-price sealed bid auctions has a dominant strategy (one that holds whatever the other bidders do). In the second-price auction this is to bid her private value. In the English auction, this is to stay in the bidding until the value reaches her private value. For the first-price sealed bid auction there is no dominant strategy, instead it is possible to determine the Nash equilibrium strategy for a given bidder, which is for that bidder to bid:

$$\left(\frac{n-1}{n}\right)V$$

[33]The fact that the common values model (see below) violates this assumption explains why first and second price auctions don't generate the same revenue in the common values model.

[34]McAfee and McMillan [67] give the example of foreign and domestic firms bidding for a government contract where there are systemic differences in the pricing structures (because of tariffs for example).

where there are n bidders and the bidder in question has private value V. Since the situation facing a bidder in a Dutch auction is exactly that facing a bidder in a first-price sealed bid auction, this result holds for the Dutch auction too (and it is clear that the Dutch auction and the first-price sealed bid auction will generate the same revenue [117] even when the first two assumptions do not hold).

The reason for identifying the four assumptions given above is in order to see what follows when they hold and fail to hold. Revenue equivalence holds when all four assumptions are true, but this does not mean that the auctions are the same. Revenue equivalence is a statement about the expected revenue, or, equivalently, the average revenue. The dominant strategy is easy to see (and state) for English and second-price auctions, whereas it is more complex for first-price and Dutch auctions. As a result, the variance of the revenue is greater for first-price and Dutch auctions (which, as [67] points out, might be a reason for a non-risk neutral seller choosing one kind of auction over another).

If the third of our assumptions fails, for example because there are two classes of bidder, revenue equivalence no longer holds. While the English auction is much as before, the first-price sealed-bid auction is not, because bidders from different classes can see that they are facing different degrees of competition (as a foreign firm facing a tariff would). The result is that the first-price sealed bid auction can become inefficient in the sense that it may end up allocating whatever is being auctioned to a bidder other than the one that values it the most (which is where revenue equivalence fails), and neither the first-price sealed-bid, nor the English auction are optimally efficient with asymmetric bidders.

The last of the four assumptions is violated whenever bidders receive incentive payments or have to pay royalties[35], and it is possible to show that the expected revenue increases as does the royalty rate. However, as the royalty rate increases, so does the temptation for the winning bidder to misreport the amount she must pay royalties on, reducing the optimum royalty rate below 100%. For further discussion of these issues, see Section 3.5.

Finally, consider what happens if the assumption about risk neutrality is violated because bidders are risk averse. In this case it turns out that the English auction produces less revenue than the first-price sealed bid auction, and that the first-price auction is not optimal. Instead the optimal auction involves subsidising high bidders who lose, and penalising low bidders. Since "the optimal auction is very complicated" [67] we will say no more about it.

3.3 The correlated values model

In the case of mineral and spectrum auctions one of the assumptions of the independent private values model is violated — the assumption that the buyers all have an accurate value for the good being auctioned. As Milgrom and Weber [74] argue, many auctions also violate the assumption that values are independent. To take the mineral rights auction as an example, independence means

[35] In auctions of oil rights on government land, for example, the seller can observe the amount of oil extracted and demand royalties based on the amount extracted.

that none of the bidders establish their value as a result of related information (geological surveys for example), or make any inference of value based on the interest other bidders have in the particular area for auction.

To allow for an analysis which fits the real world more precisely, Milgrom and Weber developed a model in which bidders have values that are positively correlated[36] so that as one bidder's estimate of the value of the good rises, so does that of all others (meaning, for example, that as bids rise in an English auction, so do the valuations of the bidders). The resulting *correlated values* model is a generalisation of the independent private value model, and also includes the situation of the common values model as a special case. In fact, the independent private value model, where buyer valuations have no effect on one another, and the common values model, where buyers have the same valuation (give or take errors in estimation) are the two extremes of the correlated values model — the model generalises both independent and common values to allow arbitrary relations between valuations to be captured.

Analysing English, Dutch and first and second-price sealed bid auctions using their model, Milgrom and Weber found a number of interesting results. First, just as in the independent private value model, Dutch and first-price auctions are found to be equivalent, so that result is robust with respect to the accuracy and independence of buyer valuations. The second result is that when buyers are uncertain about their valuations, then English and second price sealed bid auctions are not equivalent since the English auction will generally lead to higher prices.

The reason for this second result is quite interesting. What Milgrom and Weber actually show is first that in the general case revealing information in their model means that prices will increase, and that, in the specific case of the second price sealed bid auction, it is always better to reveal information than to not reveal it (that is even if the information is bad it is better to report it). Then Milgrom and Weber argue that, in the two bidder case, English auctions and second price sealed bid auctions are equivalent (using much the same argument as in the independent private value model). Finally, they analyse the general case of an English auction as follows. At some point there will be two bidders left, at which point the auction gives the same result whether it is an English auction or a second price sealed-bid auction. If the auction is sealed-bid, we can "add in" the other bidders without changing the auction — the "final two" are just the two with the highest values and the presence of bidders with lower values makes no difference to the outcome of the auction. However, if the auction is English, the phase of bidding which eliminates all but the final two bidders will, because of the information reported, tend to raise final price.

Another result that follows from the model is that when bidders' estimates of value are independent, then the second-price sealed bid auction generates higher prices than the first-price sealed bid auction. This result is obtained [72] by an extension of the argument given above for higher prices in English auctions. In the English auction, prices rise because of the declared value that other bidders

[36] Actually the necessary notion to get this effect is somewhat stronger than correlation [72].

put on the good. It is the positive influence of these declared values on the final price paid that causes the latter to rise compared with a sealed-bid auction. Similarly, since in a second-price sealed bid auction the price paid is determined by another bidder's valuation, the positive influence will again work (albeit less strongly) and a second-price sealed bid auction will lead to higher prices than a first-price sealed bid auction.

This result in turn allows the four auctions studied to be ranked by the prices (and thus revenues for the sellers) — prices are higher in second-price sealed bid auctions than in first-price sealed bid or Dutch auctions (which have the same values[37]), and higher in English auctions than in second-price sealed bid auctions.

3.4 Almost common values and the winner's curse

The correlated values model provides one way to relax the assumption that bidders in an auction have independent private values for a good. Another approach is Klemperer's [55] *almost common values* model. This model is like the common values model in assuming that bidders come close to having common values, because the good is one that actually does have the same value to each of them (since they will be selling it on, or exploiting it in exactly the same way) but that they have different estimates as to what this value is. The different estimates introduce small asymmetries between bidders.

The idea is that a bidder with a higher value is inclined to bid a little more aggressively. This may be a small effect, but it means that competing bidders with a lower value face an increased winner's curse if they win, and this tends to make them bid more conservatively than they would otherwise do. This, in turn, means that the more aggressive bidder faces less competition, and so will win the auction for a lower price and suffer a lower winner's curse than if they hadn't valued the good so highly. This effect magnifies a small advantage into a much larger one, and, as Klemperer [55] discusses, offers a convincing explanation of the results of a number of auctions that fit the almost common values mould. One is the case of the Los Angeles license in one spectrum auction, in which the license can be assumed to have a similar value to all bidders other than Pacific Telephone (PacTel). PacTel had a small advantage — a database of existing customers in the license area, brand-name recognition, and executives who were familiar with the area (in which PacTel already operated). PacTel duly won the auction, and for a price that was lower per head ($26 as opposed to $31) than the license for Chicago, a license which might be assumed to be less valuable.

Given that it can be an advantage to be the most aggressive bidder, it can be in a bidder's interests to *appear* to be the most aggressive bidder. Then other bidders may be encouraged to be more cautious and reduce the winner's curse for the aggressive bidder. Klemperer [58] points out that Glaxo's assertion that it "would almost certainly top a rival bid" when acquiring Welcome in 1995, could

[37] Milgrom [73] reports that experimental studies have shown that Dutch auctions lead to lower prices than first-price sealed bid auctions, and suggests that this may be because in practice people do not analyse the auction and determine strategy.

be read as an attempt to impress its aggressive stance on its competitors. If it was, it certainly paid off. Glaxo won the bidding for Welcome with its initial bid of £9 billion, apparently frightening Roche and Zeneca into not bidding £11 billion and £10 billion respectively (bids that the companies were thought to have contemplated making). Bulow and Klemperer [12] give a theoretical analysis of such situations.

3.5 The revelation principle

All the results we have presented so far are specific to particular auction mechanisms. It is natural to ask if it is possible to carry out more general analyses, and indeed it is. Here we are in the realm of *mechanism design*, and it is a large area of research in economics. In this section we will give a brief, and largely non-technical, overview of some of the basic results in the area. However, to really understand mechanism design one has to get technical, and the interested reader is referred to the excellent introductions in [52] (from which this section is largely drawn) and [85, Chapter 2].

We will start by roughly sketching the basic framework in which one typically thinks about mechanism design and introducing some of the terminology[38]. The first element is the set of agents taking part in the mechanism — the set of traders in an auction. The mechanism itself has a set of possible outcomes — in the case of an auction these are the various ways the good(s) can be allocated to the traders. Agents have preferences over the set of outcomes — which can be used to define the utility to each agent of a given outcome — and some private information that is usually referred to as the *type* of the agent. Commonly in an auction setting, the preferences of an agent depend only on its type, a situation that is known as a private values setting. Finally, the mechanism has a *decision rule* which determines the outcome as a function of the types of the agents in the mechanism.

This background is enough to identify one measure of a good decision rule. (A decision rule in an auction setting is, of course, the way in which the auctioneer determines who the winner is.) A decision rule is *efficient*, or *allocatively-efficient*, if it maximises the total utility of all the agents taking part in the mechanism. In contrast, a decision rule is *dictatorial* if the outcome is completely determined by the type of one of the agents (a dictatorial auction would always be decided by the bid of one particular agent). Another measure is *Pareto efficiency.* A decision rule is Pareto efficient if the outcome it chooses is such that no other outcome gives one agent more utility without giving at least one other agent less utility. Allocative efficiency and Pareto efficiency coincide if utility can be transferred between agents (for example by passing money from one to another).

Such transfers turn out to be very important in mechanism design. As laid out so far, the theory requires that agents truthfully report their type (in an

[38] The early emphasis of much of this work was from the perspective of identifying rules by which societies should be organised. This influences some of the terminology.

auction setting this requires that they bid their private value). If agents don't do this, then they can distort the result of the decision rule in their favour, just as an agent can distort the outcome of an auction it its favour by, for example, bidding below its private value. To encourage truthful reporting of type, it is possible to tax or subsidize agents based on the type they report, thus creating *transfers*, and the key question in mechanism design is how to do this in order to make mechanisms efficient and to not penalize agents who report truthfully.

To incorporate this into the theoretical framework, we need to add a *social choice function* which extends the decision rule with an associated transfer. This means that once all the agents have reported (possibly untruthfully) their types, the social choice rule identifies the outcome and also a tax or subsidy for each agent. Thus the benefit an agent gets from the social choice function is the actual utility it obtains from the outcome (which may differ from the utility identified by the type it reported) plus the utility of the transfer. At least this is the case for an agent that has *quasi-linear* preferences, which is the usual case considered in the mechanism design literature. A social choice function is said to be *budget balanced* if the sum of all the transfers is zero, and this is clearly a desirable property. A weaker property is *feasibility*, where the sum of all the transfers is negative[39]. A social choice function that is feasible is also said to be *weakly budget balanced*.

Now we have the concept of social choice function, we can identify what mechanism design theory calls a *mechanism*. This moves away from the idea that we have been working with, of outcomes being defined by agent types, and considers them (more realistically) being based upon messages that agents send. (In other words we think of the mechanism as operating on the bids made, not directly on agents' private values.) A mechanism, in these terms, is a function that defines what we have been calling the outcome and transfers for all combinations of possible messages. A mechanism, therefore, *implements* a social choice function. The social choice function specifies what should be done given knowledge of agent types. A mechanism specifies what should be done given what the agents actually report. Social choice functions can indeed be viewed as a special class of mechanisms, *direct* mechanisms, in which the decision is taken as if the agents are reporting their type directly.

The key thing in designing a mechanism is to have it implement some desirable social choice function. Mechanism design typically achieves this end by making it desirable for agents to be truthful about their types. In the language of mechanism design, such mechanisms are *incentive compatible*.

Now, we can think of the message that an agent sends in a mechanism as a strategy in the sense of game theory. In exactly the usual way, a strategy is dominant for an agent if it is optimal no matter what any other agents do [82, page 181]. If all agents have a dominant strategy to truthfully report their type, then the mechanism is said to be *dominant strategy incentive compatible* or simply *strategy-proof*. In other words, in such mechanisms, there is no reason

[39] A non-feasible social choice function would require money to flow into the mechanism. A feasible but non-budget balanced social choice function would generate money that must be spent by agents other than those in the mechanism.

for an agent to do any strategising — its best move is to just report its type (or in an auction, just bid its private value). The big question is, how can we build such mechanisms.

Part of the answer is provided by the *revelation principle*. This result[40], which is relatively quick to prove, states that any strategy-proof mechanism can be implemented as a direct mechanism. In other words, to find the kind of mechanism we want, we only have to look at mechanisms in which the messages the agents send are ones about their type.

The revelation principle is one of the positive results of mechanism design. A more negative result is the Gibbard-Satterthwaite theorem [34, 101]. This states that if agents have general preferences (in other words they are not held to be quasi-linear, or otherwise constrained), the mechanism includes at least two agents, and there are at least three different optimal outcomes, then a social choice function is strategy-proof if and only if it is dictatorial. Thus, in the most general case, it is not realistic to implement mechanisms using dominant strategies, but (as [85] points out) often the mechanism being designed provides additional structure on the problem. This can take us into areas in which the Gibbard-Satterthwaite theorem does not apply — where agents do have quasi-linear preferences, for example.

Further negative results that are directly applicable to auctions are the impossibility results of Hurwicz and others [42, 49, 50, 51], which boil down to the fact that it is not possible to create a mechanism which is efficient, budget-balanced and strategy-proof in a simple exchange environment, even when the agents have quasi-linear preferences (a simple exchange is an auction-like market in which buyers seek single units of the same good from a seller).

Despite these negative results, there is a family of direct mechanisms for agents with quasi-linear preferences, which are allocatively efficient, and strategy proof. These are the Vickrey-Clarke-Groves (VCG) mechanisms [18, 46, 117]. The form of these mechanisms [46] is the following. The social choice rule picks the outcome that maximises the total value according to the messages. In other words, the social choice rule treats the messages as if they were truthful reports of agent types, and leaves it to the transfer function to ensure that truthful reporting occurs. The transfer function identifies a tax or subsidy for agent i by considering how the total value maximised by the choice rule would differ had i not taken part in the mechanism.

If the agents in a VCG mechanism are competing in the purchase of a single good, then we get the Vickrey [117] auction. Agents report their valuation for the good, and the social choice rule allocates it to the agent who reports the highest value for it. The transfer function makes every agent "pay" the difference between the total value given their bid and the total value if they had not bid. For all agents except the winner, this transfer is zero. For the winner it is exactly the (negative) difference between the second-highest bid and the winning bid. This reduces the payment required so that it is equal to the second-highest

[40] There is another version of the revelation principle which holds for mechanisms that are implemented in Bayesian-Nash equilibrium [82, page 26], but this is beyond our scope.

bid. As already discussed, the Vickrey auction can be generalised [63, 116], and this gives a mechanism that is applicable to combinatorial auctions. This generalised Vickrey auction is itself a special case of the *pivotal* mechanism originally identified by Clarke. Transfers in the various specialisations of pivotal mechanisms have a nice interpretation. For an agent i we compute the outcome without that agent's message. If it has no effect on the outcome, then the transfer is zero. If this is not the case, then the agent is pivotal in determining the outcome, and the transfer represents the loss imposed on other agents as a result of i's presence. This provides a suitable incentive for i to be truthful in reporting his type.

This completes our brief look at some of the more general results that can be obtained for auction mechanisms — much more information can be found in [52], [85] and [65]. These results help to set the limits on what can be achieved — the Hurwicz result, for example, rules out finding a mechanism that is simultaneously efficient, budget-balanced and strategy-proof — and thus focus the search for useful mechanisms that have some compromise set of desirable properties. However, from a computer science perspective, even the positive results from mechanism design theory leave something to be desired. As both [85] and [20] point out with respect to the revelation principle, knowing that any strategy-proof mechanism can be implemented as a direct mechanism doesn't actually tell you *how* to do the implementation, and certainly doesn't tell you how to do the implementation *efficiently* (in a computational, rather than an economic, sense).

3.6 Collusion, lying, and other sharp practice

The analyses considered in the previous sections assume that all bidders are behaving fairly. That is bidders are assumed to be acting independently, and to not be attempting to alter the outcome of the auction other than by attempting to win the auction by bidding at their valuation of the good in question (give or take a little bid shading). In practice, bidders are not so benign, and in this section we consider a number of ways that they can attempt to gain some kind of advantage.

Different varieties of auction are susceptible to different kinds of manipulation. Milgrom [72] cites [41]'s analysis of *rings* of bidders (also called "pies" [14, page 187] and "kippers" [14, page 189]) in English auctions. In a ring, a group of bidders get together and agree that they are all interested in a set of items and will attempt to purchase these together, re-auction them as a way of deciding which of the group actually get the items, and then divide the proceeds of the sale amongst themselves. So long as the other bidders do not know the true value of the good, and so will not be prepared to bid above the ring, this can be profitable for the ring members. When a ring is formed, one member of the ring alone bids up to the highest price that any individual member of the ring is willing to pay. In an English auction, no member of a ring can exploit the ring agreement (introducing any additional bidding outside the ring is equivalent to not being in the ring in the first place), so such coalitions are stable. However, in first-price sealed bid auctions, any defector from the ring can steal a march on their erstwhile collaborators by bidding above the price the ring has agreed on, and the same kind of result will hold in a Dutch auction where the defector, or her accomplice, can jump in just before the point that the ring would do so making the ring unstable. Furthermore [14, page 180], to be effective, any ring needs to incorporate all the highest bidding buyers — otherwise the ring will simply be outbid.

Related to work on bidding rings is work on *bidding clubs* [61] in which buyers aggregate to increase their market power[41]. Bidding clubs invite agents to join, and these then hold a "knockout" auction amongst themselves, the winners of which proceed to the main auction. Agents that join such clubs benefit from doing so (because of the reduced competition in the auction itself) as do other agents in the main auction who are not part of the club. The only loser, then, is the seller, who realises a lower price for the goods being sold.

In any kind of auction with a closing time, *sniping* can be a problem, and is observed in both internet auctions and their low-tech cousins silent auctions. When there is a long deadline, there is no advantage in bidding early. All that does is to signal that one is interested in the good and potentially push the price up. Indeed [62], the strategy of bidding up to one's true valuation (these auctions are intended to be English auctions with an implicit open-cry) is dominated by the strategy of submitting the same bid just before the auction ends. If all bidders do this, it effectively turns the auction into a first-price sealed bid auction (which generates less revenue for the seller), and if all bidders don't do this they run the risk of having a "sniper" narrowly top their bid without giving them time to reply (though arguably if they are willing to reply they haven't bid their true valuation).

Another kind of manipulation that buyers can practice is *bid shielding*, though this is restricted to auctions, such as internet auctions, in which retraction of bids is allowed. If retraction is permitted, a bidder can make a low bid and then get an accomplice to make a very high bid which discourages other bids. Just before the auction closes (or even afterwards if it is allowed) the high bid is retracted, and the low bid wins. But retraction may also occur after the auction closes. Swiss auctions [115], run on a first-price sealed-bid basis, try to cope with such type of retractions. Once the auction is over, the winner is not allowed to change his bid, but instead he can choose whether to withdraw it or not. The practical reason for this auction is to award construction contracts and eventually the winning companies may find that they cannot meet the project specifications because they change or they find themselves overcommitted. The overall effect is that often bidders bid more aggressively because the post-auction bid retraction allows for flexibility. However, when the second highest bid is significantly different from the highest bid, the auctioneer can

[41] though we have classified bidding clubs under "collusion", Leyton-Brown *et al.* are careful to point out that they take a neutral stance, and one can imagine situations where such behaviour is justified. Indeed, [14, page 184] suggests that rings and auctioneers can develop mutually profitable relationships, with rings stepping in to make sure that surplus goods are purchased in return for later concessions.

force the company making the highest offer to meet its bid.

There has been considerable interest in the suggestion that there was some kind of collusion in the FCC spectrum auctions. As described in [24], bidders in the spectrum auctions used their bids to signal to one another — typically by using the last few digits of a bid (bids were typically at least six figure dollar numbers) for a "block" of spectrum. Some of the time this was just used to identify the bidder, as in the case of GTE ending bids with "483", which spells "GTE" on a telephone keypad[42]. However, in some auctions, the signalling was more collusive in nature. For example [24] describes signalling between two bidders, Mercury PCS and High Plains Wireless, about two markets, Lubbock and Amarillo in Texas:

> After trading bids on block F of Lubbock for several rounds, with the price rising by 10% in each round, Mercury bumped High Plains in round 121 from Amarillo, a market on which High Plains had been the standing high bidder since round 68. This was Mercury's first bid on Amarillo during the auction. The bid served as punishment to High Plains for bidding against Mercury on Lubbock, a punishment made clear since it contained as its last three digits "246" the market number for Lubbock.

This example shows bidders both indulging in *code bidding*, that is using the market numbers as part of a bid, and *retaliation*, that is punishing a rival by placing a bid on a market the rival is known to be interested in. Cramton and Schwartz [23] exhaustively analysed the DEF-block auctions held August 1996–January 1997, and found that although only a small fraction of bidders employed these tactics, the bidders that did win more than 40% of the available spectrum in terms of population covered (476 of the 1,479 licences for sale), and paid significantly less for it than bidders who did not signal or retaliate ($2.50 per person in the covered areas rather than $4.34).

This kind of activity was not limited to the FCC auctions. Klemperer [58] reports a similar exchange between Mannesman and T-Mobil[43] in the spectrum auction in Germany in 1999. In that auction it was stipulated that any new bid on a block had to be at least 10% higher than the previous bid. Mannesman's first bids on one set of blocks was 20 million DM per MHz and 18.18 million DM per MHz on another. Since 18.18 plus a 10% increase is almost exactly 20, T-Mobil interpreted this as an offer to split the blocks, accepted it, and the auction ended at that point.

One interesting thought on collusion is suggested by Menezes [71] in his analysis of multi-unit Japanese auctions. His analysis suggests that such auctions will end with all buyers bidding at the reservation price[44], and that this might

[42] In the United States, each key on the keypad of a phone has letters associated with it, just as mobile phone keypads do.

[43] Then the name for the company that now calls itself T-Mobile.

[44] Which raises the question as to why a seller would use such a mechanism. Menezes gives an answer — such auctions may be appropriate if the seller wants to guarantee the sale of all units rather than maximise revenue.

then appear as if the sellers had been colluding in order to fix the price. Extending this idea, in any auction in which it is easy for the buyers to determine the equilibrium price, it will be hard to tell if collusion has taken place.

In general, the literature indicates that bid collusion is a pervasive problem with criminal cases in highway construction [92], the distribution of school milk [91], utility procurement and other auction markets. The problem with colluding behaviour is that it is extremely hard to detect as widely admitted in the economics literature. The general trend in collusion detection is summarised in the work of Porter *et al.* [92] and Bajari *et al.* [6] — they create models for competitive and collusive behaviour to subsequently assess, over recorded auction data, whether there was competition or collusion.

Of course, it is not only buyers who can attempt to manipulate the outcome of an auction. Auctions are particularly vulnerable to fraudulent behaviour of auctioneers (who typically receive a percentage of the sale price[45] and so have a vested interest in raising it), since they often act as a trusted third party. A case in point is the second-price sealed bid auction. The auctioneer can't change the winning bid, because that would be easily detected, but the auctioneer could easily increase the second-price bid (or insert a bid of higher value than was actually submitted) increasing the price the winner must pay to virtually that of the winning bid, and [72], effectively converting the auction into a first-price sealed bid auction. But probably the best known case of auction collusion is represented by Sotheby's and Christie's scandal [79]. Sotheby's and Christie's, the world's two major auction houses, conspired to engage in price-fixing between 1993 and 1999. Traditionally both auction houses made their profits by taking a commission on the amount received by the seller. However, strong competition between the two had lead to an arms race to reduce commissions. Therefore, they had come to rely for profits on commissions from smaller consignments and those charged to buyers. Former *CEO*s of the companies decided to avoid risking profitability by stopping competition and fixing the price of commission rates charged to sellers. This led to indictments by a federal jury on the United States.

Shills, mentioned in Section 2.2, are an easy way for sellers or auctioneers to manipulate the price in their favour. In any kind of open-cry auction some accomplice of the seller or auctioneer can enter the bidding and attempt to move the price upwards by bidding appropriately (indeed, according to the correlated values model, even having the shill in the auction at all will help to increase the final price). Such practice can be very hard to detect, and it is complicated by some common, legitimate, practices by auctioneers. First, [14, page 105], an auctioneer in an English auction may bid on behalf of the seller — this is apparently legitimate as long as the seller has set a reserve price, and the conditions of sale include a notice to the effect that the auctioneer may bid for the seller. As [14, page 105] points out, it is quiet possible for the auctioneer to bid the price up even when these conditions are not met. Second, bidders

[45] According to [62], Sotheby's charges the buyer 15% of the final bid (on top of the sale price) and the seller 20% of the sale price whereas eBay charges around 5%.

in an auction may place "book bids" [14, page 152], where a potential buyer who cannot attend the auction instructs the auctioneer the maximum that she is prepared to pay, and the auctioneer then bids on her behalf[46]. It is unclear how to tell an auctioneer who is executing boiok bids from one who is spuriously bidding up the price on behalf of their client.

Shills are a particular problem in Internet auctions. Since anyone can easily obtain a number of Internet identities, there is nothing to stop a devious seller placing an item for auction and then bidding for it himself under a false name, playing the role of their own shill. Consider the case of Kenneth Walton and his shill bidding practices on eBay [118] — Walton and his cohorts were indicted after placing shill bids on hundreds of eBay auctions for a year. Many Internet auctions even alleviate the main disadvantage of shilling, which is the risk of overbidding so that the seller "wins" the item, by allowing bids to be retracted even after the auction is over. The problem of shilling, though, can be addressed by suitable design of the auction protocol, and exactly this has been done by Yokoo *et al.* [126].

Of course, in internet auctions, devious sellers can do far more than use shills in order to deceive, so mechanisms for preventing *false name bids* being advantageous have only limited effectiveness. As [59] explains, it is simple to set up an Internet identity, use false names to build up a good reputation for that identity, and then hold a real auction for non existent goods after which the identity just fades away.

Finally, we should note that there are ways of affecting the outcome of auctions that are perfectly legal with respect to the rules of the auction, but which only exist because of loopholes in the auction design (in fact sniping, which we have already mentioned, falls into this class of activities). A good example of this, given by [70], is that of the 1993 satellite television licenses in Australia. The auction itself was a first-price sealed-bid auction, which was duly won by the two highest bidders (there were two licenses on offer) with extremely high bids. However the two bidders, both of which were new to the satellite television business, defaulted on their bids. This did not cost them any money because the Australian government did not require a deposit from bidders.

Now, under the rules of the auction, if a winning bidder defaulted in this way, then the next highest bidder became the winner. However, in both cases the next highest bidder was the same as the highest bidder — each of the winners had placed a series of progressively higher bids for the same license. Naturally these second-highest bids were defaulted on as well, and the price fell progressively through a number of defaults until the licenses were finally sold for around A$100 million less than originally bid. Both ended up in the hands of the same company, one of the original two high bidders, and were immediately sold to other companies for a large profit. As Klemperer [58] points out, not setting a high cost for defaulting means that the bidding is for options on goods rather

[46] This is an idea that has been adopted by eBay, which offers bidders the services of an automate bidding agent that will maintain the highest bid for an item up to some specificed limit.

than the goods themselves, which changes the game somewhat. Furthermore, companies like the high bidders in the Australian auction, which were small and had few resources [70] are actually favoured in this kind of market. If their option is overvalued they can avoid commitments through bankruptcy, something that is not a possible course of action open to other bidders in the auction.

Another kind of loophole was exploited when Turkey auctioned two telecom licenses in sequence with the condition that the reserve price for the second auction would be the selling price of the first (a measure that might be thought to prevent the kind of problem seen in the Swiss spectrum auction). However, one company used this rule to its advantage, bidding far more for the first license than would be sensible were there to be a second operator. However, the result of this bid was to raise the reserve for the second license so high that nobody could afford it, giving the first company a monopoly.

4 Abstract Models

Clearly there are commonalities between different varieties of auction, and this has led to attempts to classify them in abstract terms. We have already seen the "zoological" approach, and there are a number of more sophisticated models along the same lines including Friedman's characterisation of double auctions [30] (and our own variation on that [90]), and Smith's classification of microeconomic models in general [109]. Here we will discuss two models from the computer science literature which we think are particularly helpful from the perspective of anyone implementing an auction (at least that is our experience).

4.1 The parametric model

The first model we will consider is that from [124], a paper that stems from the authors' work on the Michigan Internet AuctionBot [123]. AuctionBot implements a wide variety of auction mechanisms, and providing this functionality led the authors to identify a range of parameters that describe the variety. Their description of the parameters is broken down into three parts:

1. Common auction characteristics;
2. The auction parameter space; and
3. Matching functions.

We will discuss these in turn.

4.1.1 Common auction characteristics

We begin by identifying three main activities that take place in auctions — receiving bids and asks[47], clearing (that is determining what gets traded at

[47] Wurman *et al.* use a slightly different terminology from that we have gleaned from the literature. They use "bid" in the generic sense in which we use "offer", that is as an indication

what price), and revealing information to traders (in the form of issuing *quotes*). Before these can be elaborated, though, some additional features have to be identified.

First, it is necessary to identify how *net allocations*, the way in which goods and money are distributed, are specified. Basically they can be specified using discrete or continuous values, and the payment associated with each allocation can be specified linearly (the sum of quantity of each good multiplied by its price) or non-linearly (in which case a value has to be given for every possible allocation). Given these choices, the next thing to consider is how traders make offers, and these can be made either in terms of prices per good, or payments per allocation (the latter generalises the former). Offers are *divisible* whenever an allocation can always be sub-divided, and whether described as prices or payments, offers can be monotone or not monotone.

4.1.2 The auction parameter space

The space of parameters can be explored by considering the three activities mentioned in the previous section. We will start by considering offers, and note that one related parameter is whether traders may bid or ask, or do both.

To handle offers it is necessary to specify a language [81] in which they are expressed. This determines if prices or payments are to be used, and exactly what may be expressed with them. For example, traders may be allowed to make combinatorial offers, may be restricted to make offers related to a fixed number of units, or may be able to make offers that are continuous functions of prices or payments, or even provide expressions combining their own bids (e.g. OR-bids, XOR-bids, OR-of-XOR, and so on [80]).

There may be rules about what offers may be made, namely *acceptance rules*. The *ascending rule* requires bids from a given trader to increase, and asks to decrease, over time. The *decreasing rule* requires the opposite. Such restrictions may be applied across all agents as well as to a given agent. As an example, the New York Stock Exchange requires that all bids for a particular instance of a good increase over time, so what one trader bids affects the next bid from *all* traders. One may also impose rules requiring bids to "beat the quote" — quote prices are typically posted in more complex auctions, like double auctions. In a double auction the quote price is that which would win the current good were it made at the time the quote is issued (but since many buyers and sellers are simultaneously bidding, the quote may be overbid by many buyers, and sellers may even issue asks at a lower price).

There may also be rules about when and if offers may be withdrawn, and *activity rules* that state whether traders have to bid periodically to be allowed to stay in the auction (such rules are an attempt to stop the kind of last-minute "sniping" discussed above, and are used by the FCC in its spectrum auctions).

---

that a trader wishes to participate in an exchange, either as a buyer or a seller. They also distinguish the content of the bid from the bid itself and call the content the "offer". To avoid confusing readers, here we stick with the terminology we have used up until now.

Turning to the mechamism the auction uses for clearing, Wurman *et al.* [124] base their discussion around the notion of a *clearing function* that maps a set of offers to a set of net allocations and payments. Of course part of the specification of the function is how it determines what each trader pays given the offer(s) that it made. The range of possibilities will be discussed below along with the matching function. Another parameter of the clearing function determines when it is called — at a predetermined time, at a random time, based on bidder activity (as in the continuous double auction this might be when an offer is made), or based on bidder inactivity (as in the English auction). A further, similar, parameter determines when a given clearing marks the close of the auction. Other properties of the clearing function are the mechanism it uses to break ties (if appropriate) and what fees the auctioneer collects from which traders.

The final activity that we need to consider the parameters of is the mechanism for information revelation. This is the mechanism for managing quotes. Here there are many possibilities, and we will just touch on a few of the possible variations. In specifying an auction, one can decide to send the same quote to every trader, an *anonymous quote*, or tailor the quote to the trader, a *discriminatory quote*. One also needs to decide when quotes should be sent (with options rather like those for clearing), and also what to quote. Typically, as mentioned above, the quote price is an indication of what must be offered to secure a trade, and so represents information about the highest bid and/or lowest ask. However some auctions publicise the *order book*[48], that is the list of unsatisfied bids and asks, to some or all of the traders. Auctions may also reveal information about past trades (as, for example, is revealed by listings of the previous days' trading prices in newspapers). Finally, it is common practice in buy-side multi-round procurement auctions to signal each bidder to let them know their position with respect to the winning bid in order to spur competitiveness.

4.1.3 Matching functions

The final part of the description in [124] is the description of matching functions. These are the functions that determine which agents trade, where the decision is based upon the offers made. Clearly the choice of matching function depends, to some extent, on the clearing policy. For example, in a clearing house auction, one can use a matching function that pairs offers in such a way that surplus is maximised across all offers that have been received by the clearing deadline. This will involve, in general, choosing between different sets of matched offers (every bid and ask may match several offers). In contrast, in a continuous double auction, only the current highest ask and lowest bid are candidates for matching.

Wurman *et al.* roll the mechanism that sets trade prices into the matching function[49], and different classes of matching function can be distinguished on the

[48] To use the terminology of, for example, the New York Stock Exchange.

[49] To be precise, they consider the matching function has two parts — decide which agents will trade, then determine the exact terms of the trade.

basis of how they set these prices. As already mentioned, a matching function is *uniform price* if all trade prices for a given market clearing are carried out at the same price (this is the kind of thing that happens in a clearing house). A matching function is *discriminatory* if prices are tailored to pairs of traders (as in a continuous double auction). Irrespective of whether the price is uniform or discriminatory, the price in a double auction must be set at a value between that of the buyer(s) and seller(s). In a discriminatory price auction, we can have a *pay-seller's-price* policy, in which trades take place at the price asked by the seller, and its dual *pay-buyer's price* policy. These options set the price at the two extremes of the buyer-seller interval, and we can also set the price between, so that:

$$price = \kappa \ offer\_price_{buyer} + (1 - \kappa) offer\_price_{seller}$$

where $\kappa \in [0, 1]$.

These ideas can be extended to uniform pricing, by applying the same ideas to the set of all buyer and seller offers that will result in trades. As Wurman *et al.* explain, the matching function first determines the n bids that are above a price $\underline{p}$ that is acceptable to at least one seller, and the m asks that are below a price $\overline{p}$ that is acceptable to at least one buyer. There are then $l = \min(m, n)$ winning bids (the l highest bids) and l winning asks (the l lowest asks), and all the traders making these offers are happy with a transaction price in the range $[\overline{p}, \underline{p}]$. The transaction price is then set, for example using the k-double auction rule [102]:

$$price = k\overline{p} + (1 - k)\underline{p}$$

$k \in [0, 1]$. The two extreme values of k produce the $(m + 1)$st and mth price rules for uniform pricing.

Considering multi-unit offers as sets of single-unit offers, this approach can be extended to work for multi-unit auctions. The *surplus* of any exchange is the difference between the bid and ask prices (and so is independent of the trade price). An allocation determined by a matching function is said to be *locally efficient* if it maximises the total surplus across the offers it is considering.

All these mechanisms determine trade-price only based on information about bid and ask prices. It is also possible to use other information, for example the time of the offers — in such a time-based mechanism either the earliest or latest offer of a matching pair could be the one to determine the trade price — and there are also matching functions that deal with matching wholly indivisible offers (whereas the examples so far have assumed matching divisible (monetary) offers), and functions that deal with matching offers for bundles of goods. The interested reader is directed to [124], where parameter choices for a number of common auction types are also detailed.

4.2 The abstract process model

As a somewhat alternative stance from that of the parametric model, we can think of all auctions as being instantiations of a general auction process, which is broken down into a set of sub-processes[50]. These are the processes from the point of view of the auctioneer — thus an auctioneer equipped with this model would be able to handle any kind of auction by instantiating each process appropriately (which might mean not carrying out that process).

In this section we discuss this view of auctions, giving first a description of the full set of abstract processes, and then describing how some of the more common auctions fit into this framework.

4.2.1 The model

We have the following set of processes which should be familiar from the previous sections.

Bid call The auctioneer asks for bids. In any kind of sell-side auction (including sealed-bid) the bid call may be accompanied by the reserve price on the item(s). In open-cry sell-side auctions this call may also be accompanied by another price, either in addition to or instead of the reserve price — the price at which bids are requested.

Ask call The buy-side equivalent of the bid-call. In theory this could be accompanied by the analogue of the reserve price, which would be a price ceiling, but we are not aware of any auctions in which this actually happens. In a Dutch or Japanese buy-side auction, the ask call explicitly gives the price at which asks are solicited.

Bid collection Bids are collected. This activity not only covers the receipt of bids (which typically does not require the auctioneer to do much) but also the validation of bids. For instance in a typical multi-unit ascending auction [71] bidders cannot bid for more units than are for sale, and may not increase their bids over time. In addition, activity rules, like those discussed above for the FCC spectrum auctions, would be implemented as part of bid collection.

Ask collection The corresponding activity for the buy-side. Again restrictions are enforced — these will be analogous to the sell-side restrictions, for instance not to decrease asks over time in a multi-unit auction.

Bid retraction In some forms of auction, bidders are allowed to retract bids either before or after the auction closes. The auctioneer has to be able to respond to this by, for example in the case of an open-cry auction, reinstating the next highest bid. Thus being able to carry out this process requires that the auctioneer keep track of all bids.

Ask retraction In a buy-side auction, the auctioneer will have to be able to accomodate the retraction of asks.

[50]Some of these are considered by Wurman *et al.* [124], so probably the best way to think of the abstract process model is as a refinement of some components of the parametric model.

Winner determination In single-sided open-cry auctions the winner is obvious to all. In sealed-bid auctions, or call market double auctions, it is less clear who the "winner" is. The auctioneer needs to open the sealed bids and/or calculate supply and demand curves in order to establish the winner(s).

Clearing However the winner is determined, the auctioneer needs to clear the market — ensuring that the relevant traders are notified and that payments are made at the appropriate level. This will include price determination in auctions like the multi-unit Vickrey auction where such calculations are required.

Information revelation In different auctions the auctioneer may need to reveal different types of information to the participants. Reserve prices are one form of information, changes to the closing time another, and quote prices yet a third. In addition in the FCC ascending auction (as discussed in Section 2.6) the bids are sealed but the auctioneer reveals information about the bids between rounds of bidding. This is part of the information revelation process.

Tie breaking In auctions that admit ties, the auctioneer has to provide a means of breaking them.

Stage switch Auctions are a particular type of negotiation. Therefore, and following [112] they may involve several *stages*[51]. Thus, in a single-stage auction, the rules are the same from the beginning to the end of the process (this is also the case if the auction is multi-round); whereas in a multi-stage auction the rules are allowed to change (e.g. changing the auction protocol, adding or removing goods, etc.). The stage switch process takes care of changing the rules from stage to stage.

Closing The auctioneer must be able to close the auction, and, indeed, postpone closing if required as a defence against last-minute bidding [95].

These processes can be composed into an abstract process model as shown in Figure 3. Notice that the process diagram specified in Figure 3 is a business process diagram using the Business Process Modeling Notation (BPMN)[52] released by the Object Management Group [45].

We argue that our abstract model is sufficiently expressive to capture a wide range of auction types. As evidence of this: (i) we describe the instantiation of the abstract model for a number of common auctions in section 4.2.2; and (ii) we describe a more general instantiation of the abstract model for a number of auction modes in section 4.2.3.

[51] Indeed it is a common procurement practice to carry out negotiations through multi-stage auctions, and current commercial tools support such practices [16]

[52] BPMN provides a graphical language for the specification of business process models. It is intended to provide businesses with the capability of understanding their internal business procedures in a graphical notation as well as with the ability to communicate these procedures in a standard manner.

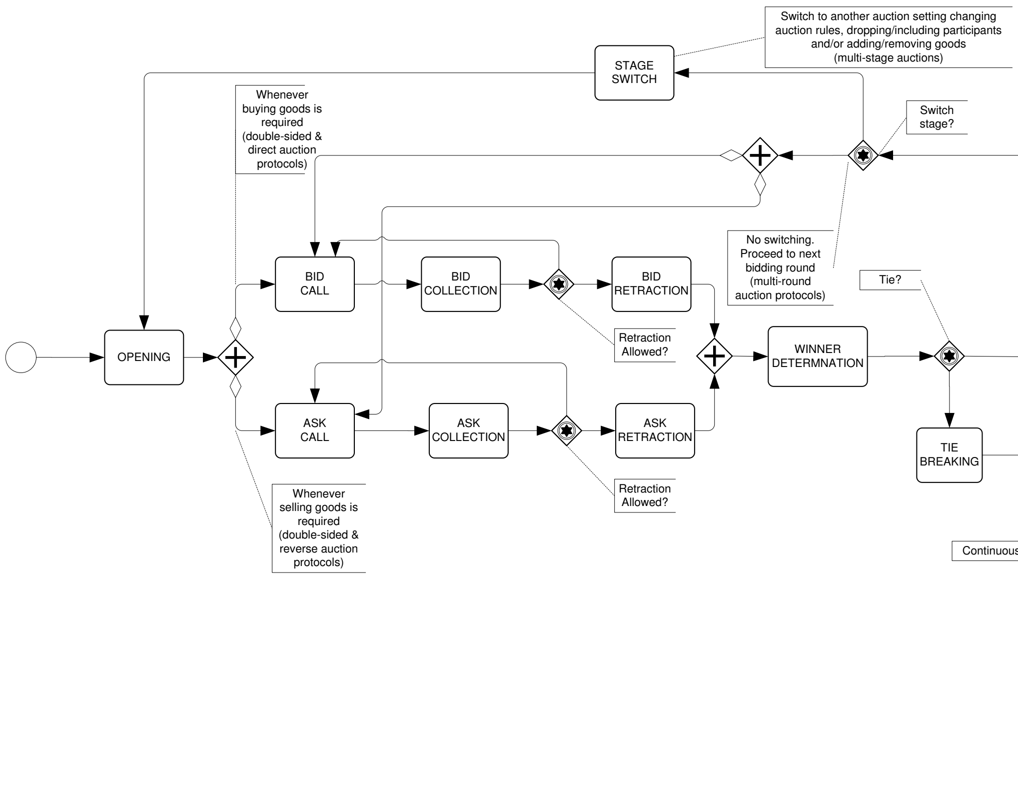

Figure 3: Abstract auction process model

Before doing this, however, there are a number of issues related to the abstract process which is worth discussing. These centre around some aspects of auctions which do not, at first blush, appear to be captured in the model as presented. These aspects are:

1. Some auction mechanisms seem to require the need to keep a record of all bids and asks made. When this is not due to the need to accomodate bid retraction, this logging of offers might be considered to be an additional process.

2. Some auctions require a number of rounds, and this might be modelled in the abstract process model.

3. Some multi-round auctions, such as the Elimination/Survival auction [32] and the Anglo-Dutch [58], involve the explicit removal of some participants between rounds, and this elimination could be considered a separate process.

While we could expand the abstract model to include these new aspects, for now we prefer to keep the abstract model simple and capture these different aspects as parts of appropriate processes which are already in the model.

For example, the process of recording all bids and asks can be considered to be something that happens across the bid/ask collection and retraction processes, with the auctioneer recording every message relating to bids and asks, and archiving these as well as recording the current bid/ask of every buyer/seller. Something similar can be done to model an auction with several rounds. Such an auction can be captured by modifying the bid/ask call processes to both signal the total number of rounds at the outset of the auction and then indicate the beginning of a new round at the appropriate time[53]. In an auction in which buyers or sellers are eliminated at each round, the elimination could either be a form of information revelation, or carried out through the bid/ask call (with, for instance, eliminated buyers and sellers not being informed of the next round).

4.2.2 Instantiations of the model for common auctions

In order to show the applicability of the process model, as well as to further explain what the various sub-processes are, we present the following instantiations of the abstract model. Although we make some remarks about the forms of auction that take place in the real world, our discussion is biased towards electronic auctions between intelligent autonomous agents.

We start in Figure 4 with a standard single-unit sell-side English auction. While hopefully this is non-contentious, there are a couple of aspects of this description which are arguable. For example, one might consider that the Winner Determination process to be explicitly that of finding the highest unretracted bid rather than folding it into the Bid Collection process. This might be considered to run after every bid is received (since bids will typically be higher than

[53] An alternative would be to signal multiple rounds implicitly — there is another round if the end of one round is followed by another bid/ask call for the same auction.

Bid call The auctioneer declares the auction open and names the good which is being auctioned, inviting bids. It may also name a starting price, and may (if we model the typical human English auction rather than a typical electronic auction) periodically invite bids of a particular value (bids that would be the current winning bid).

Ask call There is no Ask Call in a sell-side auction.

Bid collection The auctioneer receives bids. These may either explicitly denote a price, or implicitly denote a price by responding to a bid call which named a price. Either way the auctioneer updates its record of the current winning bid.

Ask collection There is no Ask Collection in a sell-side auction.

Bid retraction The auctioneer may allow bid retraction. Allowing this provides more flexibility, but also may encourage shills by making it easy for the shill to avoid buying in error [62].

Ask retraction There is no Ask Retraction in a sell-side auction.

Winner determination At any point in the auction the winner is clear, so all the process has to do is to inform the winner (necessary if the winner is not on-line in an electronic auction).

Clearing There is no clearing process. The sale price follows directly from the last bid to be collected. The auctioneer informs the winner of the price they must pay, and puts the winner in touch with the seller.

Information revelation The auctioneer reveals to each bidding agent, by some broadcast mechanism, what the last bid it received was (for instance by posting it on a web-site). In some implementations this is not necessary since bids are explicitly broadcast to all participating agents.

Tie breaking There is no tie-break process since there can be no ties in an English auction.

Closing The auctioneer decides when the auction should close according to some predetermined rule, and then transmits the fact of the closure to all participating agents. Closure may be at some pre-determined time, or may be determined by the time of the last bid.

Figure 4: The abstract process model instantiated for a single-unit sell-side English auction.

all previous bids) or once after the auction closes. Something similar might be said about the Clearing process.

Variations on the English auction theme are easy to obtain by varying the relevant portion of the abstract model. A single-unit English buy-side auction will clearly have no processes relating to bids but will have analogous processes for asks instead. Making the auction multi-unit is a little more complex. Bid Collection will no longer be able to determine the sale price. Instead we can imagine that it keeps track of supply and demand, and triggers Winner Determination and Clearing. The altered portions of the model are:

Bid collection The auctioneer receives bids and records them. The bids must explicitly denote a price and a quantity. The auctioneer also determines when the auction should close by establishing when supply exceeds demand.

Winner determination The auctioneer establishes which of the bids it currently has will obtain a sale, and which bidders these relate to. It then informs the winners.

Clearing The clearing process determines the price paid by the "winners" based on whatever pricing rule is in force (for example pay-your-bid or uniform pricing). It also informs the winners of the number of goods they will be purchasing.

Closing The auctioneer closes the auction implicitly at clearing. It may broadcast the fact of clearing making this explicit (for example so that bidders don't have to determine this for themselves).

This model assumes that all goods are allocated at a single clearing, which is a reasonable assumption if the auctioneer can determine, before clearing, that not only does demand exceed supply, but that it does so at a price above the reservation price (if any). Again this model can be made buy-side by swapping Bid Call for Ask Call, Bid Collection for Ask Collection, and Bid Retraction for Ask Retraction.

Single unit sell-side Dutch auctions may be captured by the instantiation in Figure 5 (here we have not mentioned uninstantiated processes). Obtaining multi-unit Dutch auctions is, as discussed above, relatively simple. Bidders don't just call out "mine" but instead give the number of goods they wish to obtain at the standing price, they are allocated those goods at that price (or up to as many as are for sale), and the auction restarts from that price. The processes which change are:

Bid call The auctioneer issues bid calls by naming the (descending) price of the good being auctioned and the number of goods for sale.

Bid collection As soon as the auctioneer receives a bid, the auction stops. The Winner Determination and Clearing processes are immediately invoked even if several bids are collected simultaneously or near-simultaneously.

Winner determination The winner(s) is/are obvious from the Bid Collection. The auctioneer informs the winner(s) of their winning and the number of units purchased (which may not be the same as the number requested if demand exceeds the supply). If two simultaneous or near-simultaneous bids exceed the supply, Tie Breaking is invoked. If all the units are allocated, Closing is invoked. If unallocated units remain, a new bid call is issued for the remaining units.

Clearing The clearing price is obvious from the last bid made. The auctioneer just informs the winner(s) of the price.

Bid call The auctioneer issues bid calls by naming the (descending) price of the good being auctioned.

Bid collection As soon as the auctioneer receives a bid, the auction stops. Unless two or more bids are received within some time interval, the winner and clearing price are immediately determined. If two or more bids are collected within the relevant time interval, the tie-breaking process is invoked.

Bid retraction There is typically no bid retraction allowed in Dutch auctions, but if one is (within some small interval of the bid being put in, for instance), the same procedure as for tie-breaking can be invoked.

Winner determination The winner is obvious from the Bid Collection. The auctioneer informs the winner of their winning.

Clearing The clearing price is obvious from the last bid made. The auctioneer just informs the winner of the price.

Information revelation All bids[54] are transmitted to all registered bidders or everyone monitoring the auction.

Tie breaking Ties are broken by issuing a new bid call slightly above the price of the bids that are tied.

Closing The auctioneer closes the auction implicitly at clearing. It may broadcast the fact of clearing making this explicit (for example so that bidders don't have to determine this for themselves).

Figure 5: The abstract process model instantiated for a single-unit sell-side Dutch auction.

Closing The auctioneer closes the auction explicitly once all units are allocated.

As ever buy-side Dutch auctions are obtained by swapping the Bid and Ask processes, and prices rise rather than fall over time.

Finally, we recall that a sell-side Japanese auction is like a reverse sell-side Dutch auction, with the auctioneer calling out a rising price and the buyers being implicitly in the bidding until the price exceeds what they are willing to pay and they indicate that they are dropping out. We therefore have a set of processes that are a lot like those of a Dutch auction, and these are shown in Figure 6. The multi-unit version of this auction is discussed in detail by Menezes [71], and it is rather different to the single unit auction since each buyer has to indicate, at each price increment, how many units they wish to buy. The auction closes when supply meets, or exceeds, demand. The processes that have to change to capture a multi-unit sell-side Japanese auction as opposed to single-unit sell-side Japanese auctions are as follows:

Bid call The auctioneer issues bid calls by naming the (ascending) price of the good being auctioned and the number of goods being auctioned (though the latter need only be announced at the first bid call).

Bid collection After each bid call, every buyer has to name the number of units they will buy at that price. If the total number is less than or equal

Bid call The auctioneer issues bid calls by naming the (ascending) price of the good being auctioned.

Bid collection As soon as the auctioneer has received a bid (indication of withdrawal) from all but one buyer (or all buyers), the auction stops. Unless bids from the final two or more buyers are received within some time interval, the winner and clearing price are immediately determined. If bids from the final two or more buyers are collected within the relevant time interval, the tie-breaking process is invoked.

Bid retraction There is typically no bid retraction allowed in Japanese auctions, but if one is (within some small interval of the bid being put in, for instance), the same procedure as for tie-breaking can be invoked.

Winner determination The winner is obvious from the Bid Collection. The auctioneer informs the winner of their winning.

Clearing The clearing price is obvious from the last bid made. The auctioneer just informs the winner of the price.

Information revelation All bids are transmitted to all registered bidders or everyone monitoring the auction.

Tie breaking Ties are broken by issuing a new bid call slightly below the price of the bids received[55].

Closing The auctioneer closes the auction implicitly at clearing. It may broadcast the fact of clearing making this explicit (for example so that bidders don't have to determine this for themselves).

Figure 6: The abstract process model instantiated for a single-unit sell-side Japanese auction.

to the number being auctioned, the winner determination and clearing processes are invoked.

Bid retraction The same comments apply as for bid retraction in the single-unit auction with the additional remark that there is very little need for bid-retraction in multi-unit Japanese auctions unless the bid being retracted was one which ended the auction — otherwise the bid can be corrected at the next price increment.

Winner determination The winner is obvious from the Bid Collection. All bidders still in the auction at the end of the auction win the number of items they had bid for.

Tie breaking The only "ties" will be when one buyer retracts a bid which would have closed the auction. The same process as for the single-unit auction can be used to resolve the tie.

Other auctions which can be captured in this abstract framework include continuous and periodic double auctions, the relevant instantiations for which are given in [89].

4.2.3 Instantiations of the model for auction modes

In section 4.2.2 we have provided particular instances of several auction protocols in order to show the applicability of the abstract process model depicted in figure 3. Nonetheless, we observe that we can take that exercise one step further in order to specify a general process model for each auction mode departing from the process model in figure 3. Our purpose is twofold. On the one hand, we aim at demonstrating that our general process model is general enough to express a wide range of auction types. On the other hand, we aim a providing auction process templates from which process models for particular auction types can be readily derived via specialisation. For instance, the common auctions analysed in section 4.2.2 are all examples of *multi-round direct (sell-side)* auctions. Therefore, we should expect their process models to be a specialisation of a general process model for multi-round direct auctions.

$x.y$	x followed by y	$x\|y$	x or y occurs
x^*	x occurs 0 or more times	x^+	x occurs 1 or more times
$x\|\|y$	x and y interleaved	$[x]$	x is optional

Table 1: Operators for regular expressions (x and y stand for process names)

In table 2 we provide the process model specifications for a selection of auction modes as regular expressions that we build by combining the processes in figure 3 using the operators in table 1. A single-round direct (sell-side) auction calls and collects bids in a single round, eventually allows bid retraction, runs a winner determination process to assess the winner(s) eventually followed by a tie-breaking process, and ends up by revealing information concerning the auction's result to subsequently the auction. A single-round reverse (buy-side) auction is composed of the same structure but deals with asks instead of bids. A multi-round direct (reverse) auction extends the former scheme by allowing to iterate over the call and collection of bids (asks), the assessment of (the) winner(s), and the revelation of information to participants at the end of each round. Finally, the double-sided of single-round and multi-round auctions incorporate the interleaving of bid and ask call, bid and ask collection, and (eventually) bid and ask retraction. Notice that the continuous format of a double-sided auction is slightly different since clearing may occurr at the end of a bidding round. Hence, such process model allows to readily specify a CDA.

Table 3 shows the process specification for the common auctions provided in section 4.2.2. Notice that the results indicate that the process specifications for the three auction protocols can be readily obtained as specialisations of the process model for a multi-round direct (sell-side) auction as shown in table 2. Recall that they also differ in the way they implement each sub-process in the process model specification as noticed at the beginning of section 4.2.2.

Auction mode	Process specification
single-round direct	$Opening.BidCall.BidCollection.[BidRetraction].WinnerDetermination.$ $[TieBreaking].InfoRevelation.Closing$
single-round reverse	$Opening.AskCall.AskCollection.[AskRetraction].WinnerDetermination.$ $[TieBreaking].InfoRevelation.Closing$
multi-round direct	$Opening.((BidCall.BidCollection)^{+}.[BidRetraction].WinnerDetermination.$ $[TieBreaking].[InfoRevelation])^{+}.[Clearing].Closing$
multi-round reverse	$Opening.((AskCall.AskCollection)^{+}.[AskRetraction].WinnerDetermination.$ $[TieBreaking].[InfoRevelation])^{+}.[Clearing].Closing$
doubled-sided single-round	$Opening.(BidCall\|\|AskCall).(BidCollection\|\|AskCollection).$ $[BidRetraction\|\|AskRetraction].Clearing.Closing$
doubled-sided multi-round	$Opening.((BidCall\|\|AskCall).(BidCollection\|\|AskCollection).$ $[BidRetraction\|\|AskRetraction])^{+}.Clearing.Closing$
continous double-sided	$Opening.((BidCall\|\|AskCall).(BidCollection\|\|AskCollection).$ $[BidRetraction\|\|AskRetraction]).[Clearing].[InfoRevelation])^{+}.Closing$

Table 2: Instantiation of auction models

Auction protocol	Process specification
Dutch	$Opening.((BidCall.BidCollection)^{+}.WinnerDetermination.$ $TieBreaking.[InfoRevelation])^{+}.Clearing.Closing$
English	$Opening.(BidCall.BidCollection.[BidRetraction].$ $WinnerDetermination.InfoRevelation)^{+}.Closing$
Japanese	$Opening.(BidCall.BidCollection.[BidRetraction].$ $WinnerDetermination.TieBreaking.InfoRevelation)^{+}.Closing$

Table 3: Instantiation of common auctions

5 Related work

This survey is far from being either comprehensive or the first attempt to survey work on auctions. Given the large number of papers on auction theory in the economics literature, it is unsurprising that this should be the case, but the number and range of auction surveys is quite impressive. The classic text that provided the first real attempt to survey auctions is [14], written in the late sixties, from which we have quoted liberally. More recently, most of the big names in auction theory have made a contribution to the meta-literature. Engelbrecht-Wiggans [28], writing at the start of the 1980s surveys most of the results in auction theory that were known at the time, and Stark and Rothkopf [111] provide a comprehensive bibliography from the same period. McAfee and McMillan [67] provide a very useful guide to the literature in the late 1980s, which centers around the revenue equivalence theorem, the assumptions behind it, and the implications of relaxing those assumptions. Section 3.2 is a very brief precis of this paper, though we hope we have done more than just summarise McAfee and McMillan. Milgrom [72], writing at much the same time, also centres his "primer" around the revenue equivalence theorem, though he broadens the analysis to include some aspects of his correlated-values model, and a discussion of non-auction bargaining institutions. Klemperer, too, has provided a "guide to the literature" (which also recommends [64]) in [56], covering revenue equivalence and various kinds of models of bidder values up to, and including, his own contribution on almost-common values.

All of the surveys of modern auctions tend to concentrate on results pertaining to the four main kinds of auction — English, Dutch and first and second price sealed bid, the four that are related by the revenue equivalence theorem. Klemperer [56] does, however, deal with the double auction as well (along with a number of more exotic types of auction), and this is the main subject of [30]. This latter predates the huge recent interest in double auctions in computer science, some of which is reflected in [90] (a paper that is intended as a companion to this one), while [19] gives some background on double auctions and the basic economic theory behind them.

All but the last two of these are written by economists for economists (though not necessarily auction theorists) and as a result reflect the views and mindset of that community. In contrast, we are computer scientists writing for computer scientists (especially those interested in implementing and experimenting on electronic auctions), and doubtless we reflect the biases and preoccupations of our own community.

6 Summary

This paper aims to provide a guide to auctions and bidding for computer scientists. This is an area with a large, and sometimes forbiddingly technical literature, written for and by economists. Much of this work is in the auction theory sub-field of game theory, and is largely theoretical (though, following Smith [109] there is a large body of work in experimental auction theory as well), and its main thrust is somewhat different to that of most computer scientists. Economists are largely interested in designing auctions in which humans will participate, while computer scientists are increasingly interested in auctions in which software systems will participate, more or less autonomously but on behalf of humans. Economists also tend to be interested more in the theory of auctions than computer scientists — at least in theory that explains human bidding behaviour — while computer scientists are increasingly interested in how one implements auctions. As a result, trawling through the literature on auctions can be a little off-putting for those of us working in computer science. Our aim in writing the paper was to make such a trawl less forbidding.

The paper is structured into three main sections. First, it takes an introductory look at the breadth of different auction variants — attempting to scope out the coverage of the field — and relating these variants back to the physical and virtual auctions one can encounter. Second, the paper looks at some of the theoretical results that have been obtained in auction theory. This section revolves around the various models of bidders adopted by auction theoreticians (essentially the assumptions under which they obtain results), and the results that follow from them. This section looks in particular detail at the ways that loopholes in auctions can be exploited by bidders. These are cautionary tails for anyone contemplating auction design. Finally, the paper looks at two attempts from the computer science literature to identify frameworks for comparing auctions. The second of these is here published in detail for the first time.

Acknowledgements

This work was funded in part by HP under the "Always on" grant (SP and MK), by NSF IIS-0329037 "Tools and Techniques for Automated Mechanism Design" (SP), and by IEA (TIN2006-15662-C02-01), OK (IST-4-027253-STP), eREP(EC-FP6-CIT5-28575) and 2006 5 OI 099 (JAR).

References

[1] N. An, W. Elmaghraby, and P. Keskinocak. Bidding strategies and their impact on revenues in combinatorial auctions. *Journal of Revenue and Pricing Management*, 3(4):337–357, 2005.

[2] L. M. Ausubel. An efficient ascending-bid auction for multiple objects. Working Paper 97-06, Department of Economics, University of Maryland, 1997.

[3] L. M. Ausubel and P. Cramton. Demand reduction and inefficiency in multi-unit auctions. Working Paper 96-07, Department of Economics, University of Maryland, 1998. First draft published November 1995, most recent draft published March 1998.

[4] L. M. Ausubel and P. Milgrom. Ascending auctions with package bidding. Working paper, Department of Economics, University of Maryland, 2001.

[5] M. Babaioff and W.E. Walsh. Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. *Decision Support Systems*, 39:123–149, 2005.

[6] P. Bajari and L. Ye. Deciding between competition and collusion. *The Review of Economics and Statistics*, 85(4):971–989, 2003.

[7] M. Bichler. An experimental analysis of multi-attribute auctions. *Decision Support Systems*, 29(3):249–268, 2000.

[8] M. Bichler. *The Future of e-Markets: Multi-Dimensional Market Mechanisms*. Cambridge University Press, New York, NY, USA, 2001.

[9] M. Bichler and J. Kalagnanam. Configurable offers and winner determination in multi-attribute auctions. *European Journal of Operational Research*, 160:380–394, 2005.

[10] C. Boutilier and H. H. Hoos. Bidding languages for combinatorial auctions. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, pages 1211–1217, San Francisco, CA, 2001. Morgan Kaufmann.

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

[12] J. Bulow and P. Klemperer. Prices and the winner's curse. *RAND Journal of Economics*, 33(1):1–21, 2002.

[13] T. F. Carney. *The Economies of Antiquity*. Coronado Press, Lawrence, KS, 1971.

[14] R. Cassady, Jr. *Auctions and auctioneering*. University of California Press, Berkeley, CA, 1967.

[15] J. Cerquides, U. Endriss, A. Giovannucci, and J. A. Rodriguez-Aguilar. Bidding languages and winner determination for mixed multi-unit combinatorial auctions. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 1221–1227, Hyderabad, India, January 6-12 2007.

[16] J. Cerquides, A. Reyes-Moro, J. A. Rodriguez-Aguilar, and M. López-Sánchez. Enabling assisted strategic negotiations in actual-world procurement scenarios. *Electronic Commerce Research*, 2007. Accepted for publication.

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

[18] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[19] D. Cliff. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical Report HP-97-91, Hewlett-Packard Research Laboratories, Bristol, England, 1997.

[20] V. Conitzer and T. Sandholm. Computational criticisms of the revelation principle. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC-04)*, New York, NY, USA, 2004.

[21] P. Cramton. The FCC spectrum auctions: An early assessment. *Journal of Economics and Management Strategy*, 6(3):431–495, 1997.

[22] P. Cramton. Spectrum auctions. In M. Cave, S. Majumdar, and I. Vogelsang, editors, *Handbook of Telecommunications Economics*, chapter 14, pages 605–639. Elsevier, Amsterdam, 2002.

[23] P. Cramton and J. Schwartz. Collusive bidding in the FCC spectrum auctions. Working paper, University of Maryland, 1998.

[24] P. Cramton and J. Schwartz. Collusive bidding: Lessons from the FCC spectrum auctions. *Journal of Regulatory Economics*, 17(3):229–52, May 2000.

[25] S. de Vries and R. Vohra. Combinatorial auctions: A survey. *INFORMS Journal of Computing*, 15(3):284–309, 2003.

[26] N. Economides and R. A. Schwartz. Electronic call market trading. *Journal of Portfolio Management*, 21(3):10–18, 1995.

[27] Y. Engel, M. P. Wellman, and K. M. Lochner. Bid expressiveness and clearing algorithms in multiattribute double auctions. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 110–119, New York, NY, June 2006. ACM Press.

[28] R. Engelbrecht-Wiggans. Auctions and bidding models: A survey. *Management Science*, 26:119–142, 1980.

[29] E. Fink, J. Johnson, and J. Hu. Exchange market for complex goods: Theory and experiments. *Netnomics*, 6(1):21–42, 2004.

[30] D. Friedman. The double auction institution: A survey. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 1, pages 3–25. Perseus Publishing, Cambridge, MA, 1993.

[31] D. Friedman and J. Rust. Preface. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, pages 199–219. Perseus Publishing, Cambridge, MA, 1993.

[32] Y. Fujishima, D. McAdams, and Y. Shoham. Speeding up ascending bid auctions. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, San Francisco, CA, 1999. Morgan Kaufmann.

[33] GE. Letter to Share Owners. Annual report, General Electric Corporation, USA, 2000.

[34] A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–602, 1973.

[35] A. Giovannucci, J. A. Rodríguez-Aguilar, J. Cerquides, and U. Endriss. On the winner determination problem in mixed multi-unit combinatorial auctions. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-agent Systems*, New York, NY, 2007. ACM Press.

[36] A. Giovannucci, J. A. Rodríguez-Aguilar, A. Reyes, F. X. Noria, and J. Cerquides. Towards automated procurement via agent-aware negotiation support. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 244–251, New York, NY, 2004. IEEE Press.

[37] J. K. Goeree and T. Offerman. The amsterdam auction. *Ecomometrica*, 72(1):281–294, 2004.

[38] S. D. Goiten. *A Mediterranean Society: The Jewish Communities of the Arab World as Portrayed in the Documents of the Cairo Geniza*, volume 1, Economic Foundations. University of California Press, Berkeley and Los Angeles, 1967.

[39] R. Gonen and D. Lehmann. Optimal solutions for multi-unit combinatorial auctions: Branch and bound heuristics. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 13–20, New York, NY, 2000. ACM Press.

[40] J. Gong. Exchanges for complex commodities: Search for optimal matches. Master's thesis, University of South Florida, 2002.

[41] D. A. Graham and R. C. Marshall. Collusive bidder behaviour at single-object second-price and English auctions. *Journal of Political Economy*, 95:579–599, 1987.

[42] J. R. Green and J-J Laffont. On coalition incentive compatibility. *Review of Economic Studies*, 46:243–254, 1979.

[43] A. Greenwald and P. Stone. Autonomous bidding agents in the Trading Agent Competition. *IEEE Internet Computing*, 5(2):52–60, 2001.

[44] A. R. Greenwald and J. O. Kephart. Shopbots and pricebots. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 506–511, Stockholm, Sweden, 1999.

[45] Object Management Group. BPM 1.0: BPM business process modeling notation specification. Technical report, Object Management Group, February 2006. Final adopted specification.

[46] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[47] S. Hansell. Google's slow search for a good share price. New York Times: Business Day, August 14th 2004.

[48] J. Hasbrouck, G. Sofianos, and D. Sosebess. New York Stock Exchange systems and trading procedures. Working Paper 93-01, New York Stock Exchange, April 1993.

[49] L. Hurwicz. On informationally decentralized systems. In C. McGuire and R. Radner, editors, *Decision and Organisation: A Volume in Honour of Jacob Marchak*. North-Holland, 1972.

[50] L. Hurwicz. On the existence of allocation systems whose manipulative Nash equilibria are Pareto optimal. Unpublished manuscript, 1975.

[51] L. Hurwicz and M. Walker. On the generic nonoptimality of dominant-strategy allocation mechanisms: A general theorem that includes pure exchange economies. *Econometrica*, 58:683–704, 1990.

[52] M. O. Jackson. Mechanism theory. In U. Devigs, editor, *Optimization and Operations Research*, The Encyclopedia of Life Support Science. EOLSS Publishers, Oxford, UK, 2003. The working paper version of this article includes a more comprehensive bibliography and some bibliographic notes.

[53] A. E. Kahn, P. C. Cramton, R. H. Porter, and R. D. Tabors. Pricing in the California Power Exchange electricity market: Should California switch from uniform pricing to pay-as-bid pricing? Blue Ribbon Panel Report, January 2001. A study commissioned by the California Power Exchange.

[54] J. Kalagnanam and D. C. Parkes. Auctions, bidding and exchange design. In David Simchi-Levi, S. David Wu, and Max Shen, editors, *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*, Int. Series in Operations Research and Management Science, chapter 5. Kluwer, 2004.

[55] P. Klemperer. Auctions with almost common values: The "Wallet Game" and its applications. *European Economic Review*, 42(3–5):757–769, 1998.

[56] P. Klemperer. Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13(3):227–286, 2002.

[57] P. Klemperer. How (not) to run auctions: The European 3G telecom auctions. *European Economic Review*, 46(4-5):829–845, May 2002.

[58] P. Klemperer. What really matters in auction design. *Journal of Economic Perspectives*, 16:169–189, 2002.

[59] J. 8. Lee. Making losers out of auction winners. New York Times: Circuits, March 7th 2002.

[60] J. Lévy. *The Economic Life of the Ancient World*. University of Chicago Press, Chicago, IL, 1967.

[61] K. Leyton-Brown, Y. Shoham, and M. Tennenholtz. Bidding clubs in first-price auctions. Working paper, Department of Computer Science, Stanford University, 2002. An extended abstract of this paper appeared at AAAI 2002.

[62] D. Lucking-Reiley. Auctions on the internet: What's being auctioned, and how? *Journal of Industrial Economics*, 48:227–252, 2000.

[63] J. K. Mackie-Mason and H. R. Varian. Generalized Vickrey auctions. Technical report, Department of Economics, University of Michigan, 1994.

[64] E. S. Maskin and J. G. Riley. Auction theory with private values. *American Economic Review*, 75:150–155, 1985.

[65] E. S. Maskin and T. Sjöström. Implementation theory. In K. J. Arrow, A. K. Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice Theory and Welfare*. North-Holland, Amsterdam, 2002.

[66] `http://en.wikipedia.org/wiki/Daisuke_Matsuzaka`.

[67] R. P. McAfee and J. McMillan. Auctions and bidding. *Journal of Economic Literature*, 25(2):699–738, 1987.

[68] R. P. McAfee and J. McMillan. Analyzing the airwaves auction. *Journal of Economic Perspectives*, 10:159–176, 1996.

[69] K. A. McCabe, S. J. Rassenti, and V. L. Smith. Designing a uniform-price double auction: An experimental evaluation. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, pages 307–332. Perseus Publishing, Cambridge, MA, 1993.

[70] J. McMillan. Selling spectrum rights. *Journal of Economic Perspectives*, 8:191–199, 1994.

[71] F. Menezes. Multiple-unit English auctions. *European Journal of Political Economy*, 12(4):671–684, 1996.

[72] P. Milgrom. Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3(3):3–22, 1989.

[73] P. Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108:245–272, 2000.

[74] P. Milgrom and R. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122, 1982.

[75] D. Monderer and M. Tennenholtz. Distributed games. *Games and Economic Behavior*, 28(1):55–72, 1999.

[76] D. Monderer and M. Tennenholtz. k-price auctions. *Games and Economic Behavior*, 31(1–2):220–244, 2000.

[77] A. Mu'alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In *Proceedings of the 18th National Conference on Artificial Intelligence*, pages 379–384, Edmonton, Alberta, Canada, 2002. AAAI Press.

[78] J. K. Murnighan. *Bargaining Games*. William Murrow and Company, 1992.

[79] BBC News. 2001: Sotheby's and Christie's chiefs charged. Available on: `http://news.bbc.co.uk/onthisday/hi/dates/stories/may/2/newsid_2480000/2480711.stm`, May 2001.

[80] N. Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pages 1–12, Minneapolis, Minnesota, United States, 2000. ACM Press.

[81] N. Nisan. Bidding languages for combinatorial auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, pages 215–232. MIT Press, Cambridge, MA, 2006.

[82] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, MA, 1994.

[83] D. Parkes. ibundle: an efficient ascending price bundle auction. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 148–157, New York, NY, 1999. ACM Press.

[84] D. Parkes. Optimal auction design for agents with hard valuation problems. In F. Ygge A. Moukas, C. Sierra, editor, *Agent Mediated Electronic Commerce II: Towards Next-Generation Agent-Based Electronic Commerce Systems*, volume 1788 of *LNAI*, pages 206–219. Springer, Berlin, 2000.

[85] D. Parkes. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. PhD thesis, University of Pennsylvania, Department of Computer and Information Science, May 2001.

[86] D. Parkes, R. Cavallo, N. Elprin, A. Juda, S. Lahaie, B. Lubin, L. Michael, J. Shneidman, and H. Sultan. ICE: An iterative combinatorial exchange. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, pages 249–258, New York, NY, 2005. ACM Press.

[87] D. Parkes and J. Kalagnanam. Models for iterative multiattribute procurement auctions. *Management Science*, 51:435–451, 2005.

[88] D. Parkes, J. R. Kalagnanam, and M. Eso. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *Proceedings 17th International Joint Conference on Artificial Intelligence*, pages 1161–1168. Morgan Kaufman, 2001.

[89] S. Parsons. Exception analysis for double auctions. Research Note, Center for Coordination Science, Sloan School of Management, Massachusetts Institute of Technology, 2002.

[90] S. Parsons, M. Klein, and M. Marcinkiewicz. Everything you wanted to know about double auctions, but were afraid to (bid or) ask. Research Note, Center for Coordination Science, Sloan School of Management, Massachusetts Institute of Technology, 2007.

[91] M. Pesendorfer. A study of collusion in first-price auctions. *Review of Economic Studies*, 67:381–411, 2000.

[92] R. H. Porter and J. D. Zona. Detection of bid rigging in procurement auctions. *Journal of Political Economy*, 101(3):518–538, 1993.

[93] PricewaterhouseCoopers. E-markets: Realism, not Pessimism, 2001.

[94] F. Pringsheim. The Greek sale by auction. In *Scritti in Onore di Contardo Ferrini Pubblicati in Occasione della sua Beatificazione*, volume XXVIII of *Pubblicazioni dell' Universita Cattolica del Sacro Cuore, Nuova Serie*, pages 284–343. Societa Editrice "Vita e Pensiero", Milan, 1949.

[95] A. E. Roth and A. Ockenfels. Last minute bidding and the rules for ending second-price auctions: Theory and evidence from a natural experiment on the internet. Working paper, Department of Economics, Harvard University, 2000.

[96] J. Rust, J. H. Miller, and R. Palmer. Characterizing effective trading strategies. *Journal of Economic Dynamics and Control*, 18:61–96, 1994.

[97] T. Sandholm. An algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, 2002.

[98] T. Sandholm. Optimal winner determination algorithms. In P. Cramton, Y. Shoham, and Steinberg, editors, *Combinatorial Auctions*. MIT Press, Cambridge, MA, 2006.

[99] T. Sandholm and S. Suri. Side constraints and non-price attributes in markets. *Games and Economic Behaviour*, 55:321–330, 2006.

[100] T. Sandholm, S. Suri, A. Gilpin, and D. Levine. Winner determination in combinatorial auction generalizations. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 69–76, New York, NY, 2002. ACM Press.

[101] M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[102] M. A. Satterthwaite and S. R. Williams. The Bayesian theory of the k-double auction. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 4, pages 99–123. Perseus Publishing, Cambridge, MA, 1993.

[103] Y. Shoham. A survey of auction types. Lecture notes: Stanford University CS206—Technical Foundations of Electronic Commerce, Winter 2000.

[104] Y. Shoham. Auctions on the internet: what's actually happening. Lecture notes: Stanford University CS206—Technical Foundations of Electronic Commerce, Spring 2001.

[105] Y. Shoham. Combinatorial auctions. Lecture notes: Stanford University CS206—Technical Foundations of Electronic Commerce, Spring 2001.

[106] Y. Shoham. The zoology of auctions. Lecture notes: Stanford University CS206—Technical Foundations of Electronic Commerce, Spring 2001.

[107] Y. Shoham. Combinatorial auctions. Lecture notes: Stanford University CS206—Technical Foundations of Electronic Commerce, Spring 2002.

[108] V. L. Smith. An experimental study of competitive market behaviour. *The Journal of Political Economy*, 70(2):111–137, April 1962.

[109] V. L. Smith. Microeconomic systems as an experimental science. *American Economic Review*, 72(5):923–955, 1982.

[110] J. Song and A. C. Regan. Combinatorial auctions for transportation service procurement: the carrier perspective. *Transportation Research Record*, 1833:40–46, 2002.

[111] R. Stark and M. H. Rothkopf. Competitive bidding: A comprehensive bibliography. *Operations Research*, 27:364–391, 1979.

[112] M. Strobel and C. Weinhardt. The montreal taxonomy of electronic negotiations. *Group Decision and Negotiation*, 12:143–164, 2003.

[113] T. Suyama and M. Yokoo. Strategy/false-name proof protocols for combinatorial multi-attribute procurement auction. *Autonomous Agents and Multiagent Systems*, 11:7–21, 2005.

[114] J. A. C. Thomas. The auction sale in Roman law. *The Judicial Review*, pages 42–66, April 1957.

[115] T. Von Ungern-Sternberg. Swiss auctions. *Econometrica*, 58:341–357, August 1991.

[116] H. R. Varian. Economic mechanism design for computerized agents. In *Proceedings of the USENIX Workshop on Electronic Commerce*, May 1995. Minor typos fixed 3rd March 2000.

[117] W. Vickrey. Counterspeculation, auctions, and competitive sealed bids. *Journal of Finance*, 16(1):8–37, 1961.

[118] K. Walton. *FAKE: Forgery, Lies & eBay*. Simon Spotlight Entertainment, 2006.

[119] X. Wand and M. Xia. The combinatorial bid generation problem for transportation service provision. *Transportation Research Record*, 1923:189–198, 2005.

[120] M. P. Wellman, A. Greenwald, P. Stone, and P. R. Wurman. The 2001 trading agent competition. In *Proceedings of the 18th National Conference on Artificial Intelligence*, pages 935–941, Edmonton, Alberta, Canada, 2002. American Association for Artificial Intelligence.

[121] M. P. Wellman, P. R. Wurman, K. O'Malley, R. Bangera, S. Lin, D. Reeves, and W. E. Walsh. Designing the market game for a trading agent competition. *IEEE Internet Computing*, 5(2):43–51, 2001.

[122] P. Wurman and M. Wellman. Akba: a progressive, anonymous-price combinatorial auction. In *Proceedings of ACM Conference on Electronic Commerce (EC00)*, pages 21–29, Minneapolis, 2000.

[123] P. R. Wurman, W. E. Walsh, and M. P. Wellman. Flexible double auctions for electronic commerce: Theory and applications. *Decision Support Systems*, 24:17–27, 1998.

[124] P. R. Wurman, M. P. Wellman, and W. E. Walsh. A parameterization of the auction design space. *Games and Economic Behavior*, 35(1/2):304–338, 2001.

[125] M. Xia, J. Stallaert, and A. B. Whinston. Solving the combinatorial double auction problem. *European Journal of Operation Research*, 164(1):234–251, 2004.

[126] Makoto Yokoo, Yuko Sakurai, and Shigeo Matsubara. Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence*, 130(2):167–181, 2001.

[127] E. Zurel and N. Nisan. An efficient approximation algorithm for combinatorial auctions. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, pages 125–136, Tampa, Florida, USA, 2001. ACM Press.