

Auctions and bidding: A guide for computer scientists

SIMON PARSONS

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`

and

JUAN A. RODRIGUEZ-AGUILAR

IIIA, Institut d'Investigació en Intel·ligència Artificial,
CSIC, Spanish Scientific Research Council,
Campus de la UAB,
08193 Bellaterra, SPAIN.

`jar@iiia.csic.es`

and

MARK KLEIN

Center for Coordination Science,
Massachusetts Institute of Technology,
3 Cambridge Center, NE20-336,
Cambridge,
MA 02142, USA.

`m.klein@mit.edu`

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.

Categories and Subject Descriptors: J.4 [**Computer Applications**]: Social and Behavioural Sciences—*Economics*; I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms: Economics, Design, Algorithms

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 lectern 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 competitors² 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 considerable variety of different types of auctions — as discussed 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 auctions described by [Cassady 1967]. 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 exists and crops up in various places — single-dimensional, multi-dimensional, single-sided, double-sided, first-price, second-price, k th 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 increasing volume of computer science research in the area of auctions. The aim of this paper is to help

¹Such events do not only occur in film comedies. As [Cassady 1967, page 151] reports:

The former president of Parke-Benet reports that a dealer attending a sale of eighteenth-century French furniture had arranged 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[Lucking-Reiley 2000] lists 142 auction 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 United States and Canada retired on June 3rd 2007 — a quick search suggests that there are even more on-line auctions running today.

such newcomers cope.

We start, in Section 2, by discussing the more prominent members of the auction family, and then, in Section 3, we describe some of the theoretical work that has established the properties of, and relationships between, members of the auction family. Next, Section 4, we consider two abstract views of the auction process, both drawn from work in implementing different types of auction before identifying ways in which work on auction theory feeds and is fed by computer science (Section 5). Section 6 covers related work, and Section 7 concludes.

2. THE ZOOLOGY OF AUCTIONS

Auctions have existed, in some form or other for many years. If, following [Friedman 1993], we take an auction to involve the exchange of money and goods, then it is clear that auctions could have existed since the invention of money (which both [Carney 1971] and [Lévy 1967] place around 700 BC). Indeed, auctions seem to have been in place in Greece by the 5th century BC [Pringsheim 1949], and later became widely used in the Roman Empire [Thomas 1957], while Goiten [1967] (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 auctions have been developed, and in this section we will present a classification, based on that of [Shoham 2001c], 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 k th-price.
- Auctions can be single-unit or multi-unit.
- Auctions can be single-item or multi-item (combinatorial).

Since these properties are largely independent — so that an auction that is single-dimensional can be open-cry or sealed-bid and single-unit or multi-unit — they may be used to draw up what Shoham [2000; 2001c] 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 [Wurman et al. 2001] below — but it does allow us to identify some of the main distinctions that can be drawn.

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, aspects such as quality and delivery date for example. 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




Fig. 1. A zoological categorisation of auctions

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 *kth-price* auction the winning bidder pays the price of 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 pays the price bid by the second-highest bidder.³ 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 separately, want neither.

This categorisation allows us, for instance, to classify the “typical” auction described in the introduction as a single-dimensional, one-sided, single-unit, open-cry

³Assuming the bidders are buyers — in an auction where the bidders are sellers, then the lowest price would win and the winner would pay the second-lowest price.

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 positions 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 function⁴ 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 large clock-like dial (see [Cassady 1967] and, in particular, the photographs between pages 204 and 205).

Splitting the difference between English and Dutch auctions, we get *Japanese* auctions.⁵ Here the auctioneer calls out ascending prices, and bidders indicate

⁴Indeed 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. Cassady [1967] (pages 100–108) describes a number of methods that auctioneers can use to manage this.

⁵As Ausubel [1997] points out, this name is a misnomer, which Ausubel attributes to a misreading of Cassady [1967]. [Cassady 1967, pages 63–64] 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

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. Milgrom and Weber [1982] describe an implementation of the Japanese auction (attributed to Cassady [1967]) 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 that is commonly used in charity sales [Milgrom 2000]. 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 and with the burden on bidders to check what the current highest bid is. 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. Shilling, for example, where an accomplice of the seller makes bids that are only intended to push the price up, can be particularly problematic in open-cry auctions. 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 [Cramton and Schwartz 1998; 2000]. This kind of activity, aimed at defrauding the seller and/or the auction-house is known as *collusion*. Collusion is discussed in more detail in Section 3.7.

A defense against some problems of collusion 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 k th 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 this kind of auction [Vickrey 1961]. Broadly speaking, the advantage is that making the winner pay the second highest price bid makes the bidders’ best strategy to bid their best estimate of the value of the good, taking away any strategic manoeuvring. (This does not prevent shills being effective in Vickrey auctions.)

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-

bidder, as determined by the auctioneer, is awarded the lot.

In other words, the true Japanese auction, as used to sell fish in Tokyo, involves “simultaneous-bidding” [Cassady 1967, 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.

price sealed-bid auctions for their spectrum auctions with a number of embarrassing results (taken from [McMillan 1994]). In one case a firm that bid NZ\$100,000 for a license paid NZ\$6 since that was the second highest bid, in another the highest bid was NZ\$7 million but the license 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 license 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 license relatively highly. Otherwise there is no guarantee that the outcome will extract anything like the value of the license to the winner of the auction [Milgrom 1989]. 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 the auctions described so far 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 he/she wishes 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 [1996] 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 the prices offered have risen so that demand no longer exceeds supply, there will typically 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.⁶ 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 [Ausubel and Cramton 1998].

It is, of course, possible to have multi-unit sealed-bid auctions, possibly the best known of these (though as [Ausubel 1997] points out it is not widely used because of its perceived complexity) is the multi-unit Vickrey auction analysed in [Mackie-Mason and Varian 1994].⁷ 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 many 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 had not 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 a single seller 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 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 vein, a significant number of software tools to support buy-side procurement auctions are currently commercialised by several companies, including Ariba.⁸, Perfect⁹, Emptoris¹⁰, SAP¹¹, and isOCO¹²

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

⁶An 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 on-line auctions [Lucking-Reiley 2000]

⁷As Varian [1995] 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.

⁸<http://www.ariba.com>

⁹<http://www.perfect.com>

¹⁰<http://www.emptoris.com>

¹¹<http://www.sap.com>

¹²<http://www.isoco.com>




Fig. 2. Illustrative supply and demand curves for a double auction.

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 [Friedman 1993].

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 number of iterations. Typical supply and demand curves will look like those in Figure 2.¹³ The supply and demand curves will sometimes intersect at a single price, and this defines the price at which supply and demand are in equilibrium. In cases, as in Figure 2 the intersection

¹³These supply and demand curves based on examples from [Cliff 1997]. They are based on having traders *A*, *B*, *C*, *D* and *E* all looking to supply goods that they will not sell for less than \$10, \$15, \$20, \$25, and \$30 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 \$10, no seller will be prepared to trade, while if the price is between \$10 and \$15 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.

between supply and demand curves is 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 [Satterthwaite and Williams 1993], which sets the trade price as:

$$\text{trade_price} = k \cdot \text{bid_price} + (1 - k) \cdot \text{ask_price}$$

where k is a number between 0 and 1, ask_price is the lowest bid that overlaps with a bid, and bid_price is the highest ask that overlaps with an ask. This rule is discussed more formally in Section 4 along with a related discriminatory price rule, but note that it means that no buyer pays more than it bids, and no seller receives less than it asks. In Figure 2 this rule will pick a price in the range \$20–\$25.

A slightly more complex kind of double auction is the *continuous double auction* (CDA) [Friedman and Rust 1993a]. 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 categorise double auctions, and some other variations are discussed in [Parsons et al. 2008].

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 [Menezes 1996], 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, this is not the case in a single-unit English auction — see the discussion in Section 3). This situation is analysed in [Ausubel and Cramton 1998], 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 [1997]. The mechanism of this auction is easy to describe informally:¹⁴

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

¹⁴This description assumes that the auction begins with more demand for goods than supply. This is easy to arrange if the price starts low enough.

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 [Ausubel 1997]) 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 being provably efficient in the economic sense.

Another ascending price multi-unit auction that has received a lot of attention is the auction used by the FCC to sell spectrum rights in the United States. 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 [Cramton 1997; McAfee and McMillan 1996; Milgrom 2000], 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 to users). Furthermore, as Cramton [2002] 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 the identity of the bidders 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 re-balance 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 re-balancing (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 [Sandholm 2002], 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 [McAfee and McMillan 1996] we find a patch for 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 [2002] points out, while the various FCC auctions 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 [Klemperer 2002c] 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 [Klemperer 1998] local knowledge of the area of a spectrum auction (see Section 3.4). To overcome this bias Klemperer [2002c] 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 [Klemperer 1998]. 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 [Cassady 1967, 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” [Cassady 1967, 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.¹⁵ Klemperer [2002c] 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 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 [Klemperer 2002c], and Goeree and Offerman [2004] analyse a version of the Anglo-Dutch auction, which they call the “Amsterdam auction”, in which the two final bidders receive a payment.¹⁶

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 [Fujishima et al. 1999], 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,

¹⁵Indeed 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 [Klemperer 2002b].

¹⁶As Goeree and Offerman explain, the Amsterdam auction has been used in Dutch and Belgian towns to sell property since medieval times.

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 [Fujishima et al. 1999].¹⁷

Amongst single dimensional models, we should briefly mention k th-price auctions, where $k \geq 3$. Some authors — [Wolfstetter 1996] is one example given by [Monderer and Tennenholtz 2000] — have suggested that such auctions are of no practical interest, but they seem to be useful in the context of auctions conducted on the Internet. As discussed in [Monderer and Tennenholtz 1999] and [Monderer and Tennenholtz 2000] (which also analyses 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 organiser of auctions to prefer k th-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 [Greenwald and Stone 2001; Wellman et al. 2002; Wellman et al. 2001] involving software agents endowed with automated bidding strategies.

Finally, we should note that there are many other kinds of auction that we simply do not have space to consider in detail. For example, Klemperer [2002a] describes the “all-pay” auction — which reflects lobbying competitions, in which every bidder pays their bid, but only the highest bidder wins — and Cassady [1967] (pages 71–74) describes the “handshake auction” and “whispered-bid” variants of the first-price sealed-bid auction. [Murnighan 1992]¹⁸ 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 any return. In such an auction, any bidder whose bid is exceeded stands to lose the bid as well as the good being auctioned, and in this situation [Murnighan 1992] reports that many bidders will respond by bidding very aggressively, apparently not realising that each additional bid that doesn’t win makes the potential loss even worse. In Germany, this kind of auction is known as an American auction. Another 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.¹⁹ 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, that

¹⁷To 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 set of strategies with the property that no bidder can benefit by changing her strategy while the other players keep their strategies unchanged, namely the strategies constituting a Nash equilibrium [Osborne and Rubinstein 1994, page 26]— 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.

¹⁸Quoted at <http://www.magnolia.net/~leonf/sd/bargame9.html>.

¹⁹See http://en.wikipedia.org/wiki/Daisuke_Matsuzaka.

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 flights 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 items until they end up with a combination they are happy with. An alternative approach is to run *combinatorial* auctions [de Vries and Vohra 2003], also known as auctions with *package bidding* [Ausubel and Milgrom 2001]. 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:

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

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 [2000] is that the auction designers were concerned about "free-riders"²¹. Consider a situation where three bidders are competing for two

²⁰[Shoham 2001b; 2002] gives the example of bidding for a sofa and a chair where the requirement is that the colours match.

²¹Free-riders, in this sense, are bidders who obtain good for less than they are prepared to pay

items. Bidder 1 wants license A , bidder 2 wants license 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, maybe for less than 2 would have paid if it had been in direct competition for the license (bidder 2 is the free-rider in this case). 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 (refer to [Lehmann et al. 2006] for a discussion on the NP-hardness of the winner determination problem). As a result there has been considerable interest in trying to find special-purpose optimal algorithms, for example [Sandholm 2006], as well as approximate solutions such as those provided by [Gonen and Lehmann 2000; Mu'alem and Nisan 2002; Zurel and Nisan 2001]. Another important issue that is the subject of ongoing research [Boutilier and Hoos 2001; Nisan 2000; 2006; Uckelman and Endriss 2008] 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 [Parkes 2000]. Both the bid 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 theoretical side, some contributions [Parkes 1999; Wurman and Wellman 2000] 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 a more practical side, further contributions [An et al. 2005; Song and Regan 2002b; Wang and Xia 2005] 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 [Sandholm et al. 2002], 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.²² [Parkes et al. 2001] use the term *combinatorial exchange*

for them. Thus free-riders reduce revenue for the seller, and they are hence a concern for the designers of auctions intended to maximise revenue for the seller.

²²If the notion of free disposal is assumed, there is no requirement for all items to be sold (in

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 bundles of goods. The example in [Parkes et al. 2005] perfectly illustrates the differences between a double auction and a combinatorial exchange:

... 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.

Thus a combinatorial exchange allows buyers and sellers to express more complex valuations than are possible in a double auction as well as making it possible to buy and sell. Note that the literature, for example [Xia et al. 2004], uses the term *combinatorial double auction* to denote a combinatorial exchange.

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 may not suffice for some actual-world settings. Picture the following situation. An auctioneer is interested in obtaining cars, but she is also interested in hearing offers for services that transform parts into a working car (at a certain cost). The auctioneer 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) [Cerquides et al. 2007] allow bidding agents to also bid for transformation services, that is an agent 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 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 intends to obtain a cake and receives a bid to transform batter into a cake, the bid is not of much use unless some other bidders provide the service to transform flour, eggs and sugar into cake batter, and yet others provide the flour, eggs and sugar to start the production process with. Furthermore, notice that MMUCAs must be regarded as a generalisation of the model of combinatorial auctions for supply chain formation presented in [Babaioff and Walsh 2005]. 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 devising algorithms that manage to reduce the search space (by taking advantage of the topological structure of the

sell-side auctions), or the auctioneer is permitted to balance the books by taking any extra items (in buy-side auctions).

problem) [Giovannucci et al. 2007; Giovannucci et al. 2008] or using local (heuristic) 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 [Kalagnanam and Parkes 2004; Sandholm and Suri 2006], 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 [Bichler 2001; Parkes and Kalagnanam 2005] are another type of multi-dimensional auction allowing entrants to bid over qualitative, non-price, attributes such as delivery time, weight, payment terms, as well as price. Borrowing the example from [Engel et al. 2006], 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. [2006], 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, \dots, 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

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

Thus, the auctioneer tries to assess the deal that maximises the difference between the deal's valuation and the purchasing cost. In [Che 1993], 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 [Branco 1997; Che 1993].

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 [Bichler and Kalagnanam 2005]. Note that the techniques reported in [Bichler and Kalagnanam 2005] do also consider side constraints on the combination of trades comprising an overall deal along the lines of the side constraints considered in [Kalagnanam and Parkes 2004; Sandholm and Suri 2006]. In addition to optimising winner determination,

the main issues arising when considering multi-attribute auctions have to do with their optimality and economic efficiency²³ compared to single-attribute auctions. Hence, it is obvious to pose the following questions:

- (1) Do multi-attribute 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 questions will indicate whether multi-attribute auctions are worthwhile and [Bichler 2000] provides some empirical answers. In his experiments, Bichler shows that the utility scores achieved in multi-attribute auctions were significantly higher for the buyer than those of single-attribute auctions. However, the efficiency (measured as the percentage of auctions where the bidder with the highest valuation won) was similar in single-attribute and multi-attribute auctions. Moreover, Bichler did not find evidence for revenue equivalence between the multi-attribute auction formats. As noted in [Kalagnanam and Parkes 2004], in general, following the Myerson-Satterthwaite impossibility result (see Section 3) 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 (the buyers' side and the sellers' side).

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 that are analogous to the single-attribute double-sided auctions described above, namely continuous double auctions [Fink et al. 2004; Gong 2002] and call markets [Engel et al. 2006]. Both [Fink et al. 2004] and [Gong 2002] consider a matching problem for CDAs, where deals occur whenever a pair of compatible bids is found. In a call market, as in [Engel et al. 2006], the auctioneer collects bids until some deadline, periodic or fixed, at which time she 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 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. Moreover, the problem can be further complicated by considering side constraints on overall assignments and indivisible demand²⁴ as studied in [Kalagnanam et al. 2000]. Kalagnanam et al. show that although auctions with side constraints on

²³The reader should notice that concerns about *efficiency* and *optimality*, as defined in Section 3, are at the heart of auction design. Efficient auction design is concerned with achieving allocative efficiency, namely with achieving an allocation that maximises the total value over all bidders. Hence, the goods are allocated to those who value them most. Optimal auction design concentrates on designs which maximise the expected revenue of the bid taker. Therefore, when referring to optimality, we shall be referring to revenue maximisation.

²⁴The demand requested in a single bid cannot be met using supply from more than one ask.

overall assignments can be cleared in polynomial time using network flow algorithms, when the demand is indivisible, assessing an optimal allocation becomes computationally intractable and requires solving NP-hard optimisation problems such as the generalised assignment problem, the multiple knapsack problem and the bin packing problem.

Despite the higher complexity of their clearing algorithms, multi-attribute call markets are said to count on liquidity and efficiency advantages over multi-attribute CDAs [Economides and Schwartz 1995]. Motivated by the higher complexity of call markets, Engel *et al.* [Engel et al. 2006] provide an expressive bidding language to subsequently develop network flow models that make it possible 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 [Giovannucci et al. 2004; Suyama and Yokoo 2005]. 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 [Suyama and Yokoo 2005] neatly describes 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 sub-tasks. One constructor might be able to handle multiple sub-tasks, while another company is specialised to a particular sub-task. 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 theoretical side, the work in [Suyama and Yokoo 2005] focuses on efficient protocols for combinatorial multi-attribute auctions. Therefore, this work must be regarded as a step beyond the work by [Che 1993], which did not treat multiple tasks (items). On the application side, Giovannucci *et al.* [2004] 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 [Klemperer 2002a]). 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 upon a price at which to trade.²⁵ Typically, of course, stock markets are multi-unit auctions since few such deals are over a 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 [Friedman 1993]. In addition, prices at the start of trading periods are set by a series of call markets, one for each stock.

Looking more closely [Hasbrouck et al. 1993], the situation in the New York Stock Exchange (NYSE) is a little more complex. At the time that [Hasbrouck et al. 1993] 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 [Rust et al. 1994] describe The Chicago Board of Trade operating (the latter being a market for commodities rather than stocks), with traders literally calling out bids and asks until they identified a trading partner. In addition 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 format that has always been assumed in the experimental double markets studied in the economics literature, for example in [Smith 1962], where trade takes place directly between buyers and sellers.²⁶

A slightly different mechanism, that of the Wunsch auction, is used by the Arizona stock market [McCabe et al. 1993] for all trading, and by the NYSE for trading securities after the exchange has closed [Friedman 1993]. 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 [McCabe et al. 1993].²⁷ Of course [Friedman and Rust 1993b], the NYSE operated as a call market before the 1860s, and only switched to the continuous form of the auction in response to rising trade volumes.

Another kind of auction that has become very common is the Internet auction.

²⁵[Cassady 1967, 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.

²⁶Note 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.

²⁷An auction that is similar to both the Wunsch auction and the call market is that which was used to determine the offer price of Google shares [Hansell 2004] 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.

This, 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 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 [2000] 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 the growth in such sites 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.²⁸ At the time of writing onSale has become a simple retailer of consumer electronics, but Priceline is an example of a site that still 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 still surprising. Lucking-Reiley [2000] 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 [Shoham 2001a] 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 merchandise 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 susceptibility to sniping,²⁹ which, as discussed in Section 3.7, 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 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

²⁸One might argue that all on-line retailers, especially in a world of smart consumers who make use of shop-bots [Greenwald and Kephart 1999], are engaged in a sell-side auction in which they compete for buyers with other retailers.

²⁹The practise 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.

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 [2000].

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 [Lucking-Reiley 2000] 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 [GE 2000], 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 [PricewaterhouseCoopers 2001], many of which use auction mechanisms to match buyers with sellers. These figures are backed up by Lucking-Reiley's survey [2000], 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 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 [Klemperer 2002c]. 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 [Klemperer 2002b], 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 [Kahn et al. 2001], 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. As we already anticipated in section 2, there are two concepts that are central to auction design: *efficiency* and *optimality*. The notion of efficiency we are concerned with in auction design is *allocative efficiency*, namely with achieving an allocation

of goods and money to the participants in the auction that maximises the total value over all participants. Hence, the goods are allocated to those who value them most. Optimality is concerned with maximising the expected revenue of the bid-taker. Therefore, from the perspective of the bid-taker, the revenue, or the expected selling price, is what matters. However, from the perspective of the market as a whole, efficiency—that the object ends up in the hand of the person who values it most—is more important. Both concepts will guide our analysis throughout 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 (and still meet the budgetary constraints of the seller)?” — both of which are relevant in different circumstances.³⁰

Since the outcome of any auction depends upon the way in which participants bid, any analysis will have to model this behaviour. Each bidder employs a bid function converting her valuation about the object in sale, and previous bids observed, into a bid. At this point, there are several issues to take into account influencing bidding behaviour. Firstly, we have to deal with the *uncertainty* regarding valuations faced by bidders, which is an inherent feature of auctions³¹. Here we will consider four models that each bidder can employ to assess the value of the object in sale at the time of bidding: the *independent private values* model, the *common values* model, Milgrom and Weber’s *correlated values* model [Milgrom and Weber 1982], and Klemperer’s *almost common values* model [Klemperer 1998]. Secondly, the bidders respond to uncertainty depending on their attitudes toward risk. Moreover, there is a further issue regarding whether bidders handle different information, and hence are recognisably different from each other³².

In the following we compare different auctions mainly based on the valuation model of their bidders. 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*. Our analysis assumes rational, *game-theoretic* agents. Hence, each agent has a private valuation of the object at auction, each agent has some private utility function (of the valuation) she intends to maximise, and agents play a best-response strategy to each other.³³ Therefore, from a game-theoretic perspective, a bidder’s strategy is a bid function that depends on the bidder’s own valuation and other bidders’ behaviours. Hence, such strategic behaviour determines how much a bidder shades or reveals when taking part in an auction.

Our aim is to provide an intuitive understanding of these results and to keep

³⁰The first question is clearly of interest to the seller, and to any auction institution that charges a fraction of the sale price. The second question is of interest to an auction institution that makes a flat rate commission on each good sold.

³¹A seller uses an auction because she is unsure about the values that bidders attach to the object in sale.

³²Technically speaking whether bidders are *symmetric* or not

³³This game-theoretic model of bidding behaviour is different from the *price-taking*, or *myopic best-response*, model. The latter one considers an agent playing a best-response to the current auction price and allocation, without modelling either the strategies of other agents or the effect of its own actions on the future state of the auction.

the discussion non-technical. Given that this is a highly technical area, anyone who wants to understand this material in detail will have to follow up with the references we provide. We also limit the scope of the discussion by focusing, unless otherwise stated, on results for sell-side auctions, and only those auctions that involve a single, indivisible, good, and which only involve risk-neutral buyers. By this latter we mean that buyers will be ready to bid as much as the expected value that they place on the good.

We start with the independent private value model, and our description of it and the results that follow from it are drawn from [Milgrom and Weber 1982] which gives a deeper discussion of many of the points made below.

3.1 Independent private values model

The name “independent private values” describes several of the important features of this model. First, the model assumes that each buyer has a value that she places on the good for sale, and this value is private to that buyer — that is buyers do not know the values of other buyers. The second assumption relates to the distribution of these values. Considering the whole set of buyers, each can be thought of as having been assigned a private value which is independently drawn from some distribution of values³⁴, a distribution that is known to them. 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 in which the value is independent. Finally, the independent private values model assumes that each buyer has accurate knowledge of her value.³⁵ As [McAfee and McMillan 1987] suggests, examples of auctions for which this model is appropriate are the auction of antiques where buyers are buying for their own use and not for resale, and the auction of government contracts provided each bidder knows the cost of completing the contract. 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 [1961] points out, in a first-price auction, each bidder picks a price and offers to buy the object at that price. Therefore, the bidder who bids the highest price wins. Now consider a Dutch auction. Although it may seem quite different from the first-price, a bidder will find that her problem is identical to that facing a bidder in a first-price auction. Thus, in a Dutch auction a bidder’s strategy must be to select a single number representing the first price at which the bidder will claim the object (her bid). Under the rules of the Dutch auction, the object will be awarded to the highest bidder at a price equal to her bid. This is also exactly what happens in a first-price auction. Therefore, either in a first-price or a Dutch auction the only real choice a bidder has is to pick a price to bid.

The second conclusion is that there is another kind of equivalence³⁶ between

³⁴If bidders are symmetric, the distribution is common to all bidders

³⁵It might seem odd that a buyer can have a private value, but not know what it is, but see the discussion of the common values model below.

³⁶The kind of equivalence that Milgrom and Weber [1982] state holds here is weaker than that between Dutch and first-price sealed bid auctions.

English auctions and second-price sealed bid auctions. In a second-price sealed bid auction, if a bidder bids her value for the good and 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 there will only be one more bid once the price reaches the value placed on the object by the bidder who values it second-highest. This will be the final bid by the bidder with the highest valuation for the object, and it need only be a negligible amount more than the penultimate bid. Thus the two auctions will have effectively the same outcome (though in practise there may be some non-negligible increment by which the winner has to pay over the second-highest bid).

The third result is that the outcomes of English and second-price sealed bid auctions are Pareto optimal because the winner is the bidder who has the highest value for the object. Under the independent private values model the outcomes of Dutch and first-price sealed bid auctions are Pareto optimal as well³⁷. A fourth result is that all four auctions have equal expected revenue for the seller, and so from the seller's perspective are equally good choices as a means of disposing of the good in question.

3.2 Common values model

Now, several of the results that we have just given rely on the assumption that each bidder's private value is known to her. This is reasonable in many situations, for example in the auctions of antiques or government contracts mentioned above, but is much more likely to be violated in auctions such as those for mineral or spectrum rights. To see why this is the case, consider the case in which the auction is for the right to extract minerals from some piece of land. The amount of the mineral, the cost of extraction, and the market price of the extracted mineral, will all determine the value of the mineral rights. Now, this means that while the value of the land to a given buyer may well otherwise meet the assumptions of the independent private values model, it is, at least partially, uncertain. Thus, while the buyer can estimate the value of the land, she won't have a precise value for it and so the independent private values model does not hold.

In such situations, it seems reasonable to assume that the value of the object being auctioned has more or less the same value to every bidder — the reason that this is called the *common values* or *mineral rights* [Milgrom and Weber 1982] model — but that the estimates that participants make of some parts of the value vary (for example, in terms of the example we gave above, estimates of the amount of the mineral present may 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.³⁸ This phenomenon is known as the *winner's curse*, reflecting the fact

³⁷Though this will not be if the model is not symmetric (see below), that is if the private values of different buyers are drawn from different distributions.

³⁸For the winner to have underestimated the value, all other bidders must also have underestimated it and by more.

that the winner is typically paying too much.

As Milgrom [1989] 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, their bid should be lowered. If all bidders know this, then auctions will typically attract lower highest bids than those in which bids are not shaded, and so there are incentives for the holder of the auction to come up with a mechanism in which bidders do not benefit from bid shading.

Finally, notice that Milgrom and Weber propose in [Milgrom and Weber 1982] a general model based on the notion of *affiliation* that supports the correlation among bidders' valuations and includes both the independent private values model and the common values model as special cases. Thus, bidders' valuations are said to be *affiliated* if they are positively correlated (if, roughly, some valuations being large makes it likely that the other valuations are large). We discuss this model next.

3.3 The correlated values model

One way to think about the common values model is that it addresses situations in which the assumption, made by the independent private values model, that the buyers all have an accurate value for the good being auctioned, is violated. As Milgrom and Weber [1982] argue, many auctions also violate the assumption that values are independently drawn from some underlying distribution. To take the mineral rights auction as an example, independence means 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. It seems clear that this will not always be the case.

To allow for an analysis which handles those cases where the independence assumption is violated, Milgrom and Weber developed a model in which bidders have values that are positively correlated³⁹ or *affiliated*, 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 common values model as a special case.

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 rather 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

³⁹Actually the necessary notion to get this effect is somewhat stronger than correlation [Milgrom 1989].

sealed bid auction, it is always better to reveal information than not to 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 the 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 [Milgrom 1989] 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 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 values⁴⁰), and higher in English auctions than in second-price sealed bid auctions.

3.4 Almost common values

The common values model and the correlated values model provides two ways to relax the assumptions that the independent private values model makes about the values that bidders in an auction have for a good. Another approach is Klemperer’s [1998] *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

⁴⁰Milgrom [Milgrom 2000] 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 practise people do not analyse the auction and determine strategy.

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 [1998] discusses, offers a convincing explanation of the results of a number of auctions that fit the almost common values mold. One is the case of the Los Angeles license in one spectrum auction, in which the license can be assumed to have had 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 2002c] points out that Glaxo’s assertion that it “would almost certainly top a rival bid” when acquiring Welcome in 1995, could 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 [2002] give a theoretical analysis of such situations.

3.5 Revenue equivalence

The result discussed in Section 3.1, 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. This result stems from auction theory’s most celebrated theorem, namely the *revenue equivalence theorem* (RET).⁴¹ The importance of the RET is that it establishes the conditions under which different auction forms yield the same expected revenue; and it makes it possible to rank auctions based on their revenues when its conditions do not hold. It is thus the cornerstone of theoretical attempts to relate auctions to one another, and of any attempt to choose between auctions based upon the conditions under which they will be employed.

The motivation underlying the RET is to help the seller choose the auction type that leads to the highest revenue. One statement of the theorem is as follows:

THEOREM 3.1 REVENUE EQUIVALENCE. *Assume each of a given number of risk-neutral potential buyers has a privately known valuation independently drawn from a strictly increasing atom-less distribution, and that no buyer wants more than one of the k identical indivisible objects. Then any mechanism in which (i) the objects always go to the k buyers with the highest valuations and (ii) any bidder with the lowest feasible valuation expects zero surplus, yields the same expected revenue*

⁴¹The theorem was first presented by Vickrey [1961] and subsequently generalised, roughly simultaneously, by Myerson [1981], and Riley and Samuelson [1981].

(and results in each bidder making the same expected payment as a function of her valuation).

This version of the theorem is taken from [Klemperer 2004, p. 17]. Appendix 1.A of [Klemperer 2004] includes a relatively simple proof of the theorem, and shows how it can be generalised to cases such as the common values model provided that buyers values are still independent.

According to the RET, then, when the stated conditions hold, (i) all the standard auction mechanisms, the English, Dutch, first-price sealed-bid and second-price sealed-bid (as well as many non-standard auction types), are equally profitable for the seller; and (ii) buyers are indifferent between these mechanisms. However, if the conditions in the RET do not hold, auctions may differ in expected revenue. If we just consider the four standard auctions, this is exactly the first result that we discussed under the independent private values model in Section 3.1.

A second, very influential result in auction theory is obtained when the conditions for the RET do not hold. If the assumption that bidders have independent private information is relaxed to assume that bidders' private information is affiliated,⁴² Milgrom and Weber show that ordinary ascending English auctions yield higher revenue than first-price sealed-bid auctions [Milgrom and Weber 1982]. (Again we have already met this result under a different guise — it is the second result discussed in Section 3.3). Nonetheless, despite the influence of this theoretical result, as discussed in [Klemperer 2004], there is no empirical evidence that argues that the affiliation effect is important.

McAfee and McMillan [McAfee and McMillan 1987] shed further light on the relationship between auctions with their analysis of what they call the *benchmark model*. This is basically the model we started with, having the four assumptions that:

- (1) Bidders are risk neutral;
- (2) Bidders have known, independent, private values.
- (3) Bidders are symmetric (meaning that each bidder has a private value drawn from the same distribution)⁴³; and
- (4) Payment is a function of bids alone (so the auction itself determines the exchange of goods and money and there are no additional payments between bidders).

Looking at the situation in which these assumptions hold, 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

⁴²Although the requirement is weaker than affiliation, affiliation is considered to be a satisfactory approximation.

⁴³McAfee and McMillan [1987] 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) as an example of a case in which symmetry does not hold.

given bidder so that she can choose her best bid given her guess of the decision rules being followed by the other bidders. 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 [Vickrey 1961] 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, but it says nothing specific about the strategies taken by the bidders. The dominant strategy is easy to see (and state) for English and second-price auctions, and it is easy to imagine that bidders will follow this strategy. In contrast, the Nash equilibrium strategy for bidders in first-price and Dutch auctions is more complex and it is correspondingly harder for bidders to see what to do. In turn, this may make it less likely that bidders will follow the equilibrium strategy, and, as [McAfee and McMillan 1987] points out, in turn this might be a reason for a non risk-neutral seller choosing one kind of auction over another.

If the assumption about symmetry fails, 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 who values it the most (which is where revenue equivalence fails), and neither the first-price sealed-bid, nor the English auction are fully efficient with asymmetric bidders.

The assumption about payments is violated whenever bidders receive incentive payments or have to pay royalties,⁴⁴ 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%.

Finally, consider what happens if the assumption about risk neutrality is violated because bidders are risk averse. Risk aversion makes bidders bid more aggressively in first-price auctions. Therefore, a seller faced by risk-averse bidders prefers the first-price auction to second-price sealed-bid or ascending auctions because it produces more revenue.

Although we now know how to compare the auctions analysed above in terms of revenue, it is natural to wonder about their optimality. The bad news is that, in these auctions, the seller's expected utility is not optimal and in fact can be arbitrarily far from optimal (which is the case of the Vickrey auction in general, and was clearly the case in the NZ license auctions discussed above). However, there is an auction, called the *Myerson auction* after its inventor, that yields the highest expected revenue to the seller compared to any other allocation mechanism [Myerson 1981]. Although theoretically elegant, there are several reasons that hinder the

⁴⁴In 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.

use of the Myerson auction, as discussed in [Sandholm and Gilpin 2006]. Firstly, the fact that buyers are required to fully reveal their true valuations may have a negative strategic long-term impact (e.g. when negotiating with the very same seller)⁴⁵. Secondly, the rules of the Myerson are complex and counter-intuitive, discouraging buyers from participating. Thirdly, since bid shading is rational in many commonly used auction mechanisms, buyers are likely to shade their bids in Myerson auctions, eventually leading to suboptimal allocations.

3.6 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 and computer science. In this section we will give a brief, and largely non-technical, overview of some of the basic results in the area. For a detailed technical discussion, the interested reader is referred to the excellent introductions in [Jackson 2003] (from which this section is largely drawn), [Parkes 2001, Chapter 2], [Maskin and Sjöström 2002] and the most recent survey on mechanism design for computer scientists in [Nisan 2007].

We will start by roughly sketching the basic framework in which one typically thinks about mechanism design and introducing some of the terminology.⁴⁶ 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 said to be *efficient* (again in the sense of allocative efficiency), 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 auction

⁴⁵Note, however, that this criticism is not exclusive to the Myerson auction, but applies to any auction requiring to fully reveal true valuations.

⁴⁶The 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.

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 in its favour by, for example, bidding below its private value. To encourage truthful reporting of type, it is possible to tax or subsidise 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 not to penalise 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.⁴⁷ 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 once agents directly report their types.

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 truthfully report (through their messages) their types. In the language of mechanism design, such mechanisms are *incentive compatible*. Given an incentive compatible mechanism, participants can do no better than if they truthfully report their types so we can rely on their self-interest to ensure that they do so.

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 agent does [Osborne and Rubinstein 1994, 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

⁴⁷A 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.

simply *strategy-proof*. In other words, in such mechanisms, there is no reason 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 result⁴⁸ states that under quite weak conditions any mechanism can be transformed into an equivalent incentive-compatible direct mechanism such that it implements the same social choice function. In other words, to find the kind of mechanism we want, we only have to look in the space of incentive-compatible direct mechanisms. Although the revelation principle is a powerful theoretical tool, from a computational perspective: (i) it is very expensive for agents because it places very high demands on information revelation; and (ii) it does not indicate how to construct a mechanism.

The revelation principle is one of the positive results of mechanism design. A more negative result is the Gibbard-Satterthwaite impossibility theorem [Gibbard 1973; Satterthwaite 1975]. 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.

Further negative results that are directly applicable to auctions are the impossibility results of Hurwicz and others [Green and Laffont 1979; Hurwicz 1972; 1975; Hurwicz and Walker 1990], 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 [Clarke 1971; Groves 1973; Vickrey 1961]. The form of these mechanisms [Groves 1973] 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 each agent i . This tax is assessed by composing two terms. A first term states that each agent is paid an amount equal to the sum of the values obtained by all other agents. Notice that adding this to agent i 's own value becomes exactly the total value. Therefore, a VCG mechanism aligns all agents' incentives with the social goal of maximising the social welfare. The second term in the transfer function is an arbitrary function $h_i(\cdot)$ of the reported types of every agent except i . Notice that this function has no strategic implications for agent i because it does not depend its own report. From agent i 's point of view it is just a constant that determines

⁴⁸There is another version of the revelation principle which holds for mechanisms that are implemented in Bayesian-Nash equilibrium [Osborne and Rubinstein 1994, page 26], but this is beyond our scope.

how much money is paid and in which direction. This freedom in selecting $h_i(\cdot)$ leads to a “family” of mechanisms. Different choices make different tradeoffs across the properties of the mechanism.

A particularly interesting type of transfer function is represented by Clarke payments. These identify a tax for agent i by considering how the total value maximised by the choice rule would differ had i not taken part in the mechanism. In other words, each agent i is asked to pay an amount equal to the damage that it causes to the other agents.

If the agents in a VCG mechanism are competing in the purchase of a single good, then we get the Vickrey auction [Vickrey 1961]. Agents report their valuation for the good, and the social choice rule allocates it to the agent who reports the highest value for it. After agents have reported their valuations, 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 bid. As already discussed, the Vickrey auction can be generalised [Mackie-Mason and Varian 1994; Varian 1995], 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 [Clarke 1971]. 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 the references we gave at the start of this section. 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 [Parkes 2001] and [Conitzer and Sandholm 2004] point out with respect to the revelation principle, knowing that any mechanism can be implemented as an incentive-compatible direct mechanism does not actually tell you *how* to do the implementation, and certainly does not tell you how to do the implementation *efficiently* (in a computational, rather than an economic, sense).

From the perspective of computationally efficient implementation, we find that the contribution by Nisan and Ronen [2007] to be highly significant. They focus on VCG-based mechanisms, specifically mechanisms that use a sub-optimal polynomial time algorithm for computing the outcome, and calculate the payments by applying the VCG payment rule to the underlying algorithm. VCG-based mechanisms are the computationally-tractable way of developing a mechanism when the VCG method cannot be applied because of computational intractability. Although Nisan and

Ronen observe that “essentially all reasonable VCG-based mechanisms are not truthful”, they provide the means to construct truthful mechanisms. Thus, given *any* algorithm for the corresponding optimisation problem of a VCG-based mechanism, they define the *second chance* mechanism based on it. When the agents behave truthfully, the welfare obtained by the mechanism is at least as good as the one obtained by the algorithm’s output. Nisan *et al.* show that there is a strong rationale for truth-telling behaviour and so shed light on how to derive truthful mechanisms from (computationally tractable) VCG-based mechanisms.

3.7 Collusion, lying, and other sharp practise

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 practise, 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 [1989] cites [Graham and Marshall 1987]’s analysis of *rings* of bidders (also called “pies” [Cassady 1967, page 187] and “kippers” [Cassady 1967, 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, subsequently re-auctioning them as a way of deciding which of the group actually get the items, and dividing 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 [Cassady 1967, 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* [Leyton-Brown *et al.* 2002] in which buyers aggregate to increase their market power.⁴⁹ 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

⁴⁹Though 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, [Cassady 1967, 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.

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 the silent auctions⁵⁰. 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 [Lucking-Reiley 2000], 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, as we have seen in certain settings will generate 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 practise 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. Swiss auctions [Ungern-Sternberg 1991], which run on a first-price sealed-bid basis, try to cope with the ill effects of retractions. Once the auction is over, the winner is not allowed to change her bid, but instead she can choose whether to withdraw it or not.⁵¹ 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 rules for the Swiss auction allow the auctioneer to 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 [Cramton and Schwartz 2000], 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.⁵² However, in some auctions, the signalling was more collusive in nature. For example [Cramton and Schwartz 2000] describes signaling 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

⁵⁰The growth of sites like <http://www.ezsniper.com/> and <http://www.auctionsniper.com/> which sell sniping services, suggests that sniping is rife in internet auctions.

⁵¹The practical reason for this form of auction is that it is used 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 over-committed.

⁵²In the United States, each key on the keypad of a phone has letters associated with it, just as mobile phone keypads do.

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. Though this form of bidding pushes up prices temporarily, it counts as collusion because its overall effect is to divide up the spectrum between the code-bidders in a way that reduces the amount that they pay overall. Cramton and Schwartz [1998] exhaustively analysed the DEF-block auctions held August 1996–January 1997, and found that although only a small fraction of bidders employed code-bidding tactics, the bidders that did won more than 40% of the available spectrum in terms of population covered (476 of the 1,479 licenses 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 [2002c] reports a similar exchange between Mannesman and T-Mobil⁵³ 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 bid 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 [1996] in his analysis of multi-unit Japanese auctions. His analysis suggests that such auctions will end with all buyers bidding at the reservation price,⁵⁴ and that this might 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 [Porter and Zona 1993], the distribution of school milk [Pesendorfer 2000], utility procurement and other auction markets. The problem with collusion is that it is extremely hard to detect, a fact that is widely admitted in the economics literature. The general trend in collusion detection is summarised in the work of Porter and Zona [1993] and Bajari and Ye [Bajari and Ye 2003] — they create models for competitive and collusive behaviour to subsequently assess, over recorded auction data, whether there was competition or collusion. And yet things get even more difficult in combinatorial auction settings as studied by Conitzer and Sandholm [2006] when using the VCG-mechanism. VCG is the canonical mechanism for motivating bidders to bid truthfully in combinatorial auctions and exchanges. Unfortunately, VCG is not collusion-resilient. In [Conitzer and Sandholm 2006], the authors show that collusion is much worse in combinatorial auctions

⁵³Then the name for the company that now calls itself T-Mobile.

⁵⁴Which 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.

and exchanges than in single-item settings. They provide necessary and sufficient conditions for when these highly undesirable scenarios can occur (in some scenarios colluders can even get all the goods for free!), and show how computationally hard it is to verify these conditions. The good news is that it is also hard for colluders to computationally assess a beneficial 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 price⁵⁵ 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 [Milgrom 1989], 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 [BBC News 2001]. 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 led 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 CEOs 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 in the United States.

Shills, briefly mentioned in Section 2, are one way for sellers or auctioneers to manipulate the price in their favour. For example, in an English auction, a *shill* — also called a “puffer” [Cassady 1967, page 212] — 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. English auctions are not the only kind of auction that can suffer from shilling. They can also, for example, push up the price in a Vickrey auction by increasing the second price, though this can be harder to implement than shilling in an open-cry auction because there is less information about what other bids are. In any kind of auction there is no set of shilling tactics that are guaranteed to work.⁵⁶ Such practise can be very hard to detect, and it is complicated by some common, legitimate, practises by auctioneers. First, [Cassady 1967, 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 [Cassady 1967, page 105] points out, it is quite possible for the auctioneer to bid the price up even when these conditions are not met. Second, bidders in an auction may place “book bids” [Cassady 1967, page 152], where a potential buyer who cannot attend the auction instructs the auctioneer the maximum that she is prepared to pay, and

⁵⁵According to [Lucking-Reiley 2000], 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%.

⁵⁶The 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 auctioneer then bids on her behalf.⁵⁷ It is unclear how to tell an auctioneer who is executing book 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 practises on eBay [Walton 2006] — 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 over-bidding 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 in [Yokoo et al. 2001]. Yokoo et al. coin the notion of *false-name-proof* mechanism for those mechanisms whose bidders have no incentive to use multiple identifiers. Unfortunately, as discussed in [Yokoo et al. 2004] no mechanism that allocates items efficiently is false-name-proof. An alternative, more empirical, approach is taken in [Gerding et al. 2007] which analyses the ability of different auction fees to deter shill bidding and quantify their impact on market efficiency. This paper shows that in a market with competing sellers auction fees based on the difference between the payment and the reserve price are more effective than the more commonly used fees with respect to deterring shill bidding and increasing market efficiency.

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 [Lee 2002] 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 [McMillan 1994], 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

⁵⁷As already mentioned, this is an idea that has been adopted by eBay, which offers bidders the services of an automated bidding agent that will maintain the highest bid for an item up to some specified limit.

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 [2002c] points out, not setting a high cost for defaulting means that the bidding is for options on goods (and options for which there is no default penalty) rather than the goods themselves, and this changes the game somewhat. Furthermore, companies like the high bidders in the Australian auction, which were small and had few resources [McMillan 1994] are actually favoured in this kind of market. If their option is over-valued 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 [Friedman 1993] (and our own variation on that [Parsons et al. 2008]), and Smith’s classification of micro-economic models in general [Smith 1982]. 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 [Wurman et al. 2001], a paper that stems from the authors’ work on the Michigan Internet AuctionBot [Wurman et al. 1998]. 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,⁵⁸ clearing (that is

⁵⁸Wurman *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 that a trader

determining what gets traded at 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 [Nisan 2006] 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 [Nisan 2000]).

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 block of stock increase, so that one trader bids affects the next bid from *all* traders⁵⁹. 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 over-bid 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).

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.

⁵⁹This so-called “NYSE rule”, which is imposed by a number of double auction markets, does not restrict bids on the next block of stock to be traded — the restriction is very local in time and not intended not to continually force up the trading price, but just to ensure that bids quickly converge on a price.

Turning to the mechanism the auction uses for clearing, Wurman *et al.* [2001] 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*,⁶⁰ 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 practise in buy-side multi-round procurement auctions to signal to 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 [Wurman et al. 2001] 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,⁶¹ and different classes of matching function can be distinguished on the 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

⁶⁰To use the terminology of, for example, the New York Stock Exchange.

⁶¹To be precise, they consider the matching function has two parts — decide which agents will trade, then determine the exact terms of the trade.

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 \bar{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 $[\bar{p}, \underline{p}]$. The transaction price is then set, for example using the k -double auction rule [Satterthwaite and Williams 1993]:

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

$k \in [0, 1]$. The two extreme values of k produce the $(m + 1)$ th and m th 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 [Wurman et al. 2001], 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.⁶² 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).

⁶²Some of these are considered by Wurman *et al.* [2001] so probably the best way to think of the abstract process model is as a refinement of some components of the parametric model.

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 [Menezes 1996] 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 accommodate the retraction of asks.

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 [Strobel and Weinhardt 2003] they may involve several *stages*.⁶³ 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 defense against last-minute bidding [Roth and Ockenfels 2000].

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)⁶⁴ released by the Object Management Group [Group 2006].

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. 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 accommodate 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 explicitly modelled in the abstract process model.
- (3) Some multi-round auctions, such as the Elimination/Survival auction [Fujishima et al. 1999] and the Anglo-Dutch [Klemperer 2002c], 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

⁶³Indeed it is a common procurement practise to carry out negotiations through multi-stage auctions, and current commercial tools support such practises [Cerquides et al. 2007].

⁶⁴BPMN 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.




Fig. 3. Abstract auction process model

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.⁶⁵ 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 Table I 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 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

⁶⁵An 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 [Lucking-Reiley 2000].

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.

Table I. The abstract process model instantiated for a single-unit sell-side English auction.

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 Table II (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 that change are:

⁶⁶Typically there is only one bid in a single-unit Dutch auction, but in the case of a tie and subsequent tie-break, there will be at least three.

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 bids⁶⁶ 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).

Table II. The abstract process model instantiated for a single-unit sell-side Dutch auction.

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.

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 Table III.

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.⁶⁷

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).

Table III. The abstract process model instantiated for a single-unit sell-side Japanese auction.

The multi-unit version of this auction is discussed in detail by Menezes [1996], 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 to the number being auctioned, the winner determination and clearing processes are invoked.

⁶⁷Cassady [Cassady 1967, page 89], says that “it is not an uncommon sight” to see two bidders in Japan break a tie by playing a game of Jan Ken Pon (or “paper, rock, scissors”). Given the nature of the bidding in a true Japanese auction (see above), this is not so surprising — the simultaneous nature of the bidding process makes it rather like a round of the game.

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 [Parsons 2002]. Here we briefly touch upon the specification of a continuous double auction. The process starts with the auctioneer announcing the start of the auction along with a closing time. Once the closing time is reached and the market clears for the last time, the auctioneer announces that the auction is over. Between the opening and closing there is a sequence of auction rounds that operate as follows. First the auctioneer asks for bids and asks. Then buyers and sellers transmit bids and asks, and bid and ask retractions, to the auctioneer, then the clearing time arrives and the auctioneer determines clearing prices and informs buyers and sellers of what goods they will trade. This process then repeats. Table IV summarises the process.

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 at 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.

In Table VI 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 V. 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. We should point out

Bid call. The auctioneer broadcasts a call for bids for a particular type of commodity and declares when the market will close.

Ask call. The auctioneer broadcasts a call for asks for a particular type of commodity and declares when the market will close.

Bid collection. Buyers transmit bids to the auctioneer, and the auctioneer collates the bids. Bids may be subject to rules such as the “NYSE rule” which stipulates that any new bid for an item that has already been bid on, must beat any standing bid.

Ask collection. Sellers transmit asks to the auctioneer, and the auctioneer collates the asks, subject to any rules on asks.

Bid retraction. In most models of the continuous double auction, there is no explicit retraction. Once placed, a bid can only be superseded by another bid, subject to any rules on placing bids. However, there is nothing to prevent a specific continuous double auction to allow bid retraction.

Ask retraction. As for bid retractions.

Winner determination. As each new bid or ask is collected, the auctioneer determines if the highest bid and the lowest ask “cross” (so that the bid is higher than the ask), and if so, determines the clearing price and the number of units that the buyer and seller will receive.

Clearing. The auctioneer informs the buyers and sellers of the price and number of units that they will be trading and updates the lists of bids and asks with unmet units.

Information revelation. In some continuous double auctions, such as that on the NYSE trading floor, there is no explicit information revelation process — every trader is assumed to be able to hear all the bids and asks. In other CDAs, there is an explicit posting of market quotes.

Tie breaking. There is no explicit tie-breaking; this becomes merged into the winner-determination process.

Closing. The auctioneer declares that the auction is closed.

Table IV. The abstract process model instantiated for a continuous double auction.

$x.y$	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 V. Operators for regular expressions (x and y stand for process names)

that double-sided, single-round and multi-round auctions may all incorporate the interleaving of bid and ask call, bid and ask collection, and (eventually) bid and ask retraction. Finally, note that the continuous version of the double-sided auction is slightly different from all of the above since clearing may occur at the end of a bidding round. This latter feature may be captured by our process model.

Table VII 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 VI. Recall that

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	$Opening.((BidCall AskCall).(BidCollection AskCollection).[BidRetraction AskRetraction])^+.[Clearing].Closing$
continuous double-sided	$Opening.((BidCall AskCall).(BidCollection AskCollection).[BidRetraction AskRetraction]).[Clearing].[InfoRevelation])^+.[Clearing].Closing$

Table VI. Instantiation of auction models

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 protocol	Process specification
Dutch	$Opening.((BidCall.BidCollection)^+.[TieBreaking].[InfoRevelation])^+.[Clearing].Closing$
English	$Opening.(BidCall.BidCollection.[BidRetraction].WinnerDetermination.InfoRevelation)^+.[Clearing].Closing$
Japanese	$Opening.(BidCall.BidCollection.[BidRetraction].WinnerDetermination.TieBreaking.InfoRevelation)^+.[Clearing].Closing$
Continuous Double	$Opening.((BidCall AskCall).(BidCollection AskCollection).[BidRetraction AskRetraction]).[Clearing])^+.[Clearing].Closing$

Table VII. Instantiation of common auctions

5. BRIDGING COMPUTING AND AUCTIONS

In Section 2 we surveyed the menagerie of auctions, and in Section 4 we focused on two models from the computer science literature that are helpful to implement auctions. Therefore, we have been concerned with the computational realisation of auctions as well as with introducing the theoretical analysis of them (Section 3), covering topics that are interesting both to computer scientists and to economists. However, the relationship between computer science and auctions is even stronger and, the purpose of this section is to further explore the connection. On the one hand, we will survey issues that research in computer science has tackled in order to realise electronic auctions.⁶⁸ On the other hand, we shall discuss how auctions can serve to solve computation problems.

Firstly, we focus on the research issues computer science has tackled to enact electronic auctions.

—*Human-computer interaction.* Firstly, there has been some interest in developing visualisation methods using 3D spaces to facilitate bidders choose interesting

⁶⁸Notice that some of the issues have been anticipated in Section 2.

items out of many different items in auction services such as eBay or Yahoo [Morita et al. 2005]. Other research has focused on providing facilities to explore the bid space instead by means of various analysis facilities [Lee et al. 2000]. Secondly, some research has focused on the development of visualisation assistants aimed at helping auction operators identify the most salient features in a market to undertake analysis of the market activity [Healey et al. 2001]. Healey et al. [2001] develop visualisation assistants for the Trading Agent Competition (TAC).⁶⁹ Finally, [Bogdanovych 2008] proposes developing electronic auctions as regulated virtual worlds (virtual institutions) to provide users with an immersive experience, closer to reality.

- *Preference representation and elicitation.* The issue of preference elicitation has become particularly relevant in multi-attribute auctions and combinatorial auctions. In multi-attribute auctions, the issue is how to learn the auctioneer’s and buyers’ preferences concerning the rule used for scoring and ranking bids [Bichler et al. 2001] (typically in the form of the weights in a multi-attribute utility function). Regarding combinatorial auctions, on the one hand, there is the issue of providing agents with a bidding language to compactly encode bids [Nisan 2006], and, on the other hand, there is the issue of designing query strategies for the auctioneer (elicitor) to request bidders for limited, but relevant, information about their valuations (instead of requesting all their bids) [Sandholm and Boutilier 2006]. The auctioneer employs information received by bidders to build models of their valuations until it is feasible to assess an optimal allocation. Notice that the preference elicitation problem in combinatorial auctions requires that the auctioneer deploys a querying strategy asking bidders, whereas in multi-attribute auctions the auctioneer is queried about their preferences.
- *Winner determination algorithms.* The winner determination problems for the double and multi-dimensional auctions described in section 2 lead to hard combinatorial optimisation problems. The literature has proposed either: (i) exact algorithms based either on mathematical programming tools (e.g. [Andersson et al. 2000]) or on the specific structure of the problem (e.g. [Sandholm 2006]); or (ii) approximate (local) algorithms (e.g. [Hoos and Boutilier 2000; Zurel and Nisan 2001]). Recent contributions on computationally efficient WDP solvers for different auction types — namely, [Sandholm 2006] for combinatorial auctions, [Giovannucci et al. 2007] for MMUCAS, and [Engel et al. 2006] for multi-attribute double auctions — suggest that a careful, formal analysis of the structure of WDPs can provide guidance for developing efficient winner determination solvers. Notice that most of the existing winner determination algorithms in combinatorial auctions are centralised, with very few exceptions distributing the solving of the WDP amongst the bidders [Kelly and Stenberg 2000; Narumanchi and Vidal 2006; García and Vidal 2007]. Interestingly, when distributing the winner determination, the computational problem shifts the WDP to a bid generation problem. However, there has not been much research in this direction. Finally, notice that in order to evaluate and compare the performance of winner determination algorithms there has been a significant effort on the development of

⁶⁹<http://www.sics.se/tac/>.

test suites such as [Leyton-Brown and Shoham 2006] and [Vinyals et al. 2008]. In addition, [Leyton-Brown et al. 2006] gives a methodology for constructing hardness models for winner determination algorithms so that these models can predict the running time of a given algorithm on previously unseen winner determination problem instances. Interestingly, the very same methodology has been reported to be valuable in a completely different combinatorial optimisation problem. Thus, the work in [Xu et al. 2008] describes an automated approach for constructing per-instance algorithm portfolios for SAT that use empirical hardness models to choose among their constituent solvers.

- Automated bidding.* Since the continuous double auction (CDA) is the dominant market institution for real-world trading of equities, commodities, derivatives, etc., the design of bidding algorithms for CDA has been an active focus of research (e.g. [Cliff 1997; Gjerstad and Dickhaut 1998; Das et al. 2001]). Moreover, there has been also interest in evaluating whether automated bidders can indeed outperform humans [Tesauro and Das 2001]. The interest in the design of algorithms for automated bidding led to the Trading Agent Competition, where bidding agents are pitched against one another either in a travel agents' market or in a supply chain scenario. As to combinatorial auctions, we distinguish between contributions on the bid generation problem (namely, how to bid in a combinatorial auction given that evaluating and submitting all possible bundles is not practical for the bidders and the auctioneer) [An et al. 2005; Wang and Xia 2005; Song and Regan 2002a], and on decision support for bidders to help them enter the winning set of bidders during iterative combinatorial auctions [Leskelä et al. 2007] (by generating recommendations for bids that would be among the current winners of the auction). Finally, the automated generation of multi-attribute bids in procurement auctions [Cerquides et al. 2007] has been given little attention.
- Complexity.* The hardness of the winner determination problem has been also an intense matter of research, particularly in the realm of combinatorial auctions [Lehmann et al. 2006]. An additional piece of significant work is offered in [Müller 2006], where Müller studies theoretically tractable cases of the winner determination problem. His main result is the establishment of patterns of the winner determination problem that allow to draw an equivalence between the winner determination problem and combinatorial optimisation problems.

Now it is time to consider how auctions have contributed and are contributing to solve computational problems. Computer science, and in particular artificial intelligence, has paid attention to the notion of *markets-as-computation*, namely the use of a market-based method, such as an auction, to compute the outcome to a distributed problem. Along this line the work on the *market-oriented programming* (MOP) [Wellman 1993] paradigm is based on the adoption of economic equilibrium principles as a technique to solve particular problems of distributed resource allocation. In MOP, distributed computation is implemented as a market price system wherein agents offer to buy or sell quantities of commodities. When the system reaches equilibrium, the market has assessed the allocation of resources. Besides MOP, auctions have been largely, and successfully, employed to solve a wealth of distributed control (e.g. air-conditioning control [Huberman and Clearwater 1995],

production control [Bussmann and Schild 2000]), distributed task allocation (e.g. collaborative planning [Hunsberger and Grosz 2000], formation of virtual organisations [Patel et al. 2005], coordination for robot navigation [Sierra et al. 2000], coordination in disaster management [Ramchurn et al. 2008]), and distributed resource allocation (resource allocation in sensor networks [Ostwald et al. 2005]) problems. All these market-based distributed systems share the commonality of being *flexible* and *robust* to rapidly adapt to changes and respond to failures. From a designer’s point of view, choosing (or defining) a market price system (e.g. an auction) and defining the bidding functions for the agents involved in the computational market become central issues. The methodology described in [Jennings and Bussmann 2003; Bussmann et al. 2004] provide valuable guidelines as to how to tackle this endeavour.

As noticed in [Parkes 2001, Chapter 3], the motivation for market-based approaches has been to produce efficient solutions to distributed problems. Hence computational aspects were investigated before game-theoretical properties of market mechanisms. Nonetheless, more recently we observe that there is an ongoing effort to integrate game-theoretic concerns and computational concerns in the emerging field of *computational mechanism design* (CMD) (refer to [Dash et al. 2003], [Parkes 2001] and [Conitzer 2008] for excellent introductions to the field). The challenge of CMD is to make mechanisms computationally feasible while maintaining useful game-theoretical properties such as efficiency and strategy-proofness. Notice that CMD still follows the *markets-as-computation* view, thus aiming to compute social-welfare maximising outcomes (namely outcomes that maximise the total utility of the agents in the system) for distributed problems. However, CMD does not assume that the computational entities composing a distributed system can cooperate to find a good system-wide solution. In other words, CMD assumes that the designer might not be able to enforce bidding strategies on the market participants. As noted in [Dash et al. 2003], this is the case for computer systems that are open (individual components can enter and leave at will), containing components representing distinct stakeholders with different aims and objectives (e.g. grid computing, pervasive computing, e-commerce, mobile computing, P2P systems). In such a setting, the designer can impose a market protocol but not the bidding functions (strategies) that individual stakeholders adopt.

Next, we summarise some of the research challenges CMD is tackling to design computational mechanisms that reconcile game-theoretic and computational properties:

- *Outcome determination.* Depending on the mechanism’s social choice function, computing the optimal outcome may amount to solving an NP-hard combinatorial optimisation problem.
- *Computationally efficient approximation mechanisms.* How to find good outcomes, close to the optimal one but easier to compute than the optimal ones, while satisfying incentive properties [Archer et al. 2003]; or how to develop less (computationally) complex mechanisms [Nisan and Ronen 2007].
- *Preference representation and elicitation.* As noticed above, we remind the reader that preference representation has to do with the issue of providing agents with a bidding language to compactly encode their bids, whereas preference elicitation

has to do with designing querying strategies for agents to progressively reveal their preferences till it is feasible to assess the optimal outcome. When using querying strategies the computational burden is shared between the agent and the centre computing the optimal allocation. As noted in [Dash et al. 2003], this poses interesting computational issues on the agent side: the computational burden to shift to agents (depending on their resources) and the trustworthiness on agents performing their computations.

- Dynamic mechanisms.* How to implement mechanisms that allow individuals to provide incremental information about their preferences (instead of complete information revelation) so that the centre/auctioneer can solve easier instances of the winner determination problem [Parkes 1999].
- Distributed implementation.* How to implement a model of decentralised computation that assesses the outcome of the mechanism [Petcu et al. 2006].
- Designing mechanisms for computationally bounded agents* to mitigate the computational burden of solving the strategy selection problem (from the information about the mechanism and the other agents in the mechanism) [Parkes 1999].
- Robustness.* There are several issues here. Firstly, consider that in many real-world applications agents can, and often do, fail in completing their allocated tasks. Therefore, the issue is how to develop mechanisms that can take into account measures of the probability of an agent succeeding at a given task and produce allocations that maximise social utility [Porter et al. 2002; Dash et al. 2004; Ramchurn et al. 2009]. Secondly, there is the issue of developing mechanisms such that bidders never have an incentive to cheat by using multiple identifiers (false-name-proof) [Suyama and Yokoo 2005]. Thirdly, there is the issue of developing mechanisms that prevent shill bidding [Gerding et al. 2007].
- On-line mechanisms.* How to design mechanisms that dynamically make decisions (e.g. resource allocation) as individuals arrive and leave the system [Friedman and Parkes 2003; Parkes and Duong 2007].
- Automated market mechanism design.* How to solve the mechanism design problem automatically for some particular setting at hand [Conitzer and Sandholm 2002; Phelps 2008; Phelps et al. 2008; Sandholm 2003]. In this direction, the CAT tournament [Gerding et al. 2007; Niu et al. 2008] appears as a very interesting initiative with complementary objectives to the TAC tournament. The objective of the tournament is to encourage research in the design and application of computational market mechanisms, particularly robust mechanisms capable of automatically adapting to changing environmental conditions.

6. 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 [Cassady 1967], 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 [1980], writing at the start of the 1980s, surveyed most of the results in auction

theory that were known at the time, and Stark and Rothkopf [1979] provided a comprehensive bibliography from the same period. McAfee and McMillan [1987] gave a very useful guide to the literature in the late 1980s, which centres around the revenue equivalence theorem, the assumptions behind it, and the implications of relaxing those assumptions. Section 3.5 includes a very brief precis of this paper, though we hope we have done more than just summarise McAfee and McMillan. Milgrom [1989], 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 [Maskin and Riley 1985]) in [Klemperer 2002a], 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,. Klemperer [Klemperer 2002a] does, however, deal with the double auction as well (along with a number of more exotic types of auction), and the double auction is the main subject of [Friedman 1993]. This latter predates the huge recent interest in double auctions in computer science, some of which is reflected in [Parsons et al. 2008] (a paper that is intended as a companion to this one), while [Cliff 1997] 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.

Another exception is represented by the excellent compilation of works on combinatorial auctions in [Cramton et al. 2006]. The book offers contributions to a wide range of issues in combinatorial auctions, namely mechanism design, complexity and algorithmic issues, implementation, empirical evaluation, and applications. Thus, it is an invaluable resource to computer scientists interested in combinatorial auctions.

Finally, note that we have only given a brief introduction to both mechanism design and the emerging field of computational mechanism design — our aim was to give an outline of the area and put the work in context vis a vis work on auctions. As mentioned above, there are a number of references we recommend the reader who wants to deepen their knowledge on these topics. For an introduction to classic mechanism design, we recommend both [Parkes 2001, Chapter 2] and [Jackson 2003], though Nisan offers an introduction to the field more biased towards computer scientists in [Nisan 2007]. As to computational mechanism design, a good starting point is [Parkes 2001, Chapter 3], though we definitely recommend the reader to dive into [Shoham and Leyton-Brown 2008, Chapter 10] for further discussions into implementation issues as well as for a description of several computational applications of mechanism design. Finally, one can always opt for a more gentle introduction to computational mechanism design by going through [Dash

et al. 2003].

7. 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, largely written for and by economists. Much of this work is in the sub-fields of game theory called “auction theory” and “mechanism design”, and is mainly theoretical (though, following Smith [Smith 1982] 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 and the equilibrium behaviour of markets — 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 four 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. Third, 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. Fourth, the paper points out connections between research in computer science and research into auctions, both ways in which the study of auctions is extending computer science, and ways in which computer science is extending the study of auctions.

ACKNOWLEDGMENTS

This work was funded in part by: HP under the “Always on” grant, NSF IIS-0329037 “Tools and Techniques for Automated Mechanism Design”, IEA (TIN2006-15662-C02-01), Agreement Technologies (CONSOLIDER CSD2007-0022, INGENIO 2010), EVE (EVE (TIN2009-14702-C02-01), and the Jose Castillejo programme (JC2008-00337) of the Spanish Ministry of Science and Innovation. JAR also thanks Generalitat de Catalunya (grant 2009-SGR-1434). This work began when SP and JAR worked at the Center for Coordination Science at MIT — we both learnt a lot in our time there. Finally, we are very grateful to Kevin Leyton-Brown for kindly providing us with an electronic draft version of his book, to Joel Kammet for his many comments, and to the reviewers for their many suggestions on how to improve

the paper.

REFERENCES

- AN, N., ELMAGHRABY, W., AND KESKINOCAK, P. 2005. Bidding strategies and their impact on revenues in combinatorial auctions. *Journal of Revenue and Pricing Management* 3, 4, 337–357.
- ANDERSSON, A., TENHUNEN, M., AND YGGE, F. 2000. Integer programming for combinatorial auction winner determination. In *Proceedings of the 4th International Conference on Multiagent Systems*, E. Durfee, S. Kraus, H. Nakashima, and M. Tambe, Eds. IEEE Computer Society, 39–48.
- ARCHER, A., PAPADIMITRIOU, C., TALWAR, K., AND TARDOS, E. 2003. An approximate truthful mechanism for combinatorial auctions with single parameter agent. *Internet Mathematics* 1, 2, 129–150.
- AUSUBEL, L. M. 1997. An efficient ascending-bid auction for multiple objects. Working Paper 97-06, Department of Economics, University of Maryland.
- AUSUBEL, L. M. AND CRAMTON, P. 1998. Demand reduction and inefficiency in multi-unit auctions. Working Paper 96-07, Department of Economics, University of Maryland. First draft published November 1995, most recent draft published March 1998.
- AUSUBEL, L. M. AND MILGROM, P. 2001. Ascending auctions with package bidding. Working paper, Department of Economics, University of Maryland.
- BABAIOFF, M. AND WALSH, W. 2005. Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. *Decision Support Systems* 39, 123–149.
- BAJARI, P. AND YE, L. 2003. Deciding between competition and collusion. *The Review of Economics and Statistics* 85, 4, 971–989.
- BBC NEWS. 2001. Sotheby's and Christie's chiefs charged. BBC News. Available on: http://news.bbc.co.uk/onthisday/hi/dates/stories/may/2/newsid_2480000/2480711.stm.
- BICHLER, M. 2000. An experimental analysis of multi-attribute auctions. *Decision Support Systems* 29, 3, 249–268.
- BICHLER, M. 2001. *The Future of e-Markets: Multi-dimensional Market Mechanisms*. Cambridge University Press, New York.
- BICHLER, M. AND KALAGNANAM, J. 2005. Configurable offers and winner determination in multi-attribute auctions. *European Journal of Operational Research* 160, 380–394.
- BICHLER, M., LEE, J., KIM, C., AND LEE, H. 2001. Design and implementation of an intelligent decision analysis system for e-sourcing. In *International Conference on Artificial Intelligence*. Las Vegas, Nevada.
- BOGDANOVYCH, A. 2008. Virtual institutions. Ph.D. thesis, University Technology Sydney.
- BOUTILIER, C. AND HOOS, H. H. 2001. Bidding languages for combinatorial auctions. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, B. Nebel, Ed. Morgan Kaufmann, San Francisco, CA, 1211–1217.
- BRANCO, F. 1997. The design of multidimensional auctions. *RAND Journal of Economics* 28, 1, 63–81.
- BULOW, J. AND KLEMPERER, P. 2002. Prices and the winner's curse. *RAND Journal of Economics* 33, 1, 1–21.
- BUSSMANN, S., JENNINGS, N., AND WOOLDRIDGE, M. 2004. *Multiagent systems for manufacturing control: A design methodology*. Series on Agent Technology. Springer-Verlag, Berlin, Germany.
- BUSSMANN, S. AND SCHILD, K. 2000. Self-organizing manufacturing control: An industrial application of agent technology. In *Proceedings of the 4th International Conference on Multiagent Systems*, E. Durfee, S. Kraus, H. Nakashima, and M. Tambe, Eds. IEEE Computer Society, 87–94.
- CARNEY, T. F. 1971. *The Economies of Antiquity*. Coronado Press, Lawrence, KS.
- CASSADY, JR., R. 1967. *Auctions and auctioneering*. University of California Press, Berkeley, CA.

- CERQUIDES, J., ENDRISS, U., GIOVANNUCCI, A., AND RODRIGUEZ-AGUILAR, J. A. 2007. Bidding languages and winner determination for mixed multi-unit combinatorial auctions. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, M. M. Veloso, Ed. Hyderabad, India, 1221–1227.
- CERQUIDES, J., REYES-MORO, A., RODRIGUEZ-AGUILAR, J. A., AND LÓPEZ-SÁNCHEZ, M. 2007. Enabling assisted strategic negotiations in actual-world procurement scenarios. *Electronic Commerce Research* 7, 3/4, 189–220.
- CHE, Y. K. 1993. Design competition through multidimensional auctions. *RAND Journal of Economics* 24, 4, 668–680.
- CLARKE, E. H. 1971. Multipart pricing of public goods. *Public Choice* 11, 17–33.
- CLIFF, D. 1997. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical Report HP-97-91, Hewlett-Packard Research Laboratories, Bristol, England.
- CONITZER, V. 2008. Mechanism design for MAS. Course at the Dubai Agents and Multiagent Systems School.
- CONITZER, V. AND SANDHOLM, T. 2002. Complexity of mechanism design. In *Proceedings of the 18th Conference in Uncertainty in Artificial Intelligence*, A. Darwiche and N. Friedman, Eds. Morgan Kaufmann, San Francisco, CA.
- CONITZER, V. AND SANDHOLM, T. 2004. Computational criticisms of the revelation principle. In *Proceedings of the 5th ACM Conference on Electronic Commerce*. New York.
- CONITZER, V. AND SANDHOLM, T. 2006. Failures of the VCG mechanism in combinatorial auctions and exchanges. In *Proceedings of the 5th International Conference on Autonomous Agents and Multiagent Systems*, H. Nakashima, M. P. Wellman, G. Weiss, and P. Stone, Eds. ACM Press, New York, 521–528.
- CRAMTON, P. 1997. The FCC spectrum auctions: An early assessment. *Journal of Economics and Management Strategy* 6, 3, 431–495.
- CRAMTON, P. 2002. Spectrum auctions. In *Handbook of Telecommunications Economics*, M. Cave, S. Majumdar, and I. Vogelsang, Eds. Elsevier, Amsterdam, Chapter 14, 605–639.
- CRAMTON, P. AND SCHWARTZ, J. 1998. Collusive bidding in the FCC spectrum auctions. Working paper, University of Maryland.
- CRAMTON, P. AND SCHWARTZ, J. 2000. Collusive bidding: Lessons from the FCC spectrum auctions. *Journal of Regulatory Economics* 17, 3 (May), 229–252.
- CRAMTON, P., SHOHAM, Y., AND STEINBERG, R., Eds. 2006. *Combinatorial Auctions*. The MIT Press, Cambridge, MA.
- DAS, R., HANSON, J. E., KEPHART, J. O., AND TESAURO, G. 2001. Agent-human interactions in the continuous double auction. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, B. Nebel, Ed. Morgan Kaufmann, San Francisco, CA, 1169–1187.
- DASH, R. K., PARKES, D. C., AND JENNINGS, N. R. 2003. Computational mechanism design: A call to arms. *IEEE Intelligent Systems* 18, 6, 40–47.
- DASH, R. K., RAMCHURN, S. D., AND JENNINGS, N. R. 2004. Trust-based mechanism design. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems*, N. R. Jennings, C. Sierra, L. Sonenberg, and M. Tambe, Eds. ACM Press, New York, 748–755.
- DE VRIES, S. AND VOHRA, R. 2003. Combinatorial auctions: A survey. *INFORMS Journal of Computing* 15, 3, 284–309.
- ECONOMIDES, N. AND SCHWARTZ, R. A. 1995. Electronic call market trading. *Journal of Portfolio Management* 21, 3, 10–18.
- ENGEL, Y., WELLMAN, M. P., AND LOCHNER, K. M. 2006. Bid expressiveness and clearing algorithms in multiattribute double auctions. In *Proceedings of the 7th ACM Conference on Electronic Commerce*. ACM Press, New York, 110–119.
- ENGLBRECHT-WIGGANS, R. 1980. Auctions and bidding models: A survey. *Management Science* 26, 119–142.
- FINK, E., JOHNSON, J., AND HU, J. 2004. Exchange market for complex goods: Theory and experiments. *Netnomics* 6, 1, 21–42.
- ACM Journal Name, Vol. V, No. N, Month 20YY.

- FRIEDMAN, D. 1993. The double auction institution: A survey. See Friedman and Rust [1993a], Chapter 1, 3–25.
- FRIEDMAN, D. AND RUST, J., Eds. 1993a. *The Double Auction Market: Institutions, Theories and Evidence*. Santa Fe Institute Studies in the Sciences of Complexity. Perseus Publishing, Cambridge, MA.
- FRIEDMAN, D. AND RUST, J. 1993b. Preface. See Friedman and Rust [1993a], 199–219.
- FRIEDMAN, E. J. AND PARKES, D. C. 2003. Pricing wifi at Starbucks: Issues in online mechanism design. In *Proceedings of the 4th ACM Conference on Electronic Commerce*. ACM Press, New York, 240–241.
- FUJISHIMA, Y., MCADAMS, D., AND SHOHAM, Y. 1999. Speeding up ascending bid auctions. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, T. Dean, Ed. Stockholm, Sweden, 548–553.
- GARCÍA, B. M. AND VIDAL, J. M. 2007. Bidding algorithms for a distributed combinatorial auction. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, E. H. Durfee, M. Yokoo, M. N. Huhns, and O. Shehory, Eds. IFAAMAS, 694–701.
- GE. 2000. Letter to Share Owners. Annual report, General Electric Corporation, USA.
- GERDING, E., MCBURNEY, P., NIU, J., PARSONS, S., AND PHELPS, S. 2007. Overview of CAT: A market design competition. Tech. Rep. ULCS-07-006, University of Liverpool, Department of Computer Science.
- GERDING, E. H., ROGERS, A., DASH, R. K., AND JENNINGS, N. R. 2007. Sellers competing for buyers in online markets: Reserve prices, shill bids, and auction fees. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, M. M. Veloso, Ed. Hyderabad, India, 1287–1293.
- GIBBARD, A. 1973. Manipulation of voting schemes: A general result. *Econometrica* 41, 587–602.
- GIOVANNUCCI, A., RODRÍGUEZ-AGUILAR, J. A., CERQUIDES, J., AND ENDRISS, U. 2007. 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*, E. H. Durfee, M. Yokoo, M. N. Huhns, and O. Shehory, Eds. IFAAMAS, 710–717.
- GIOVANNUCCI, A., RODRÍGUEZ-AGUILAR, J. A., REYES, A., NORIA, F. X., AND CERQUIDES, J. 2004. Towards automated procurement via agent-aware negotiation support. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM Press, New York, 244–251.
- GIOVANNUCCI, A., VINYALS, M., RODRÍGUEZ-AGUILAR, J. A., AND CERQUIDES, J. 2008. Computationally-efficient winner determination for mixed multi-unit combinatorial auctions. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi-agent Systems*, L. Padgham, D. Parkes, J. Müller, and S. Parsons, Eds. IFAAMAS, 1071–1078.
- GJERSTAD, S. AND DICKHAUT, J. 1998. Price formation in double auctions. *Games and Economic Behaviour* 22, 1–19.
- GOEREE, J. K. AND OFFERMAN, T. 2004. The Amsterdam auction. *Econometrica* 72, 1, 281–294.
- GOITEN, S. D. 1967. *A Mediterranean Society: The Jewish Communities of the Arab World as Portrayed in the Documents of the Cairo Geniza*. Vol. 1, Economic Foundations. University of California Press, Berkeley and Los Angeles.
- GONEN, R. AND LEHMANN, D. 2000. Optimal solutions for multi-unit combinatorial auctions: Branch and bound heuristics. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*. ACM Press, New York, 13–20.
- GONG, J. 2002. Exchanges for complex commodities: Search for optimal matches. M.S. thesis, University of South Florida.
- GRAHAM, D. A. AND MARSHALL, R. C. 1987. Collusive bidder behaviour at single-object second-price and English auctions. *Journal of Political Economy* 95, 579–599.
- GREEN, J. R. AND LAFFONT, J.-J. 1979. On coalition incentive compatibility. *Review of Economic Studies* 46, 243–254.
- GREENWALD, A. AND STONE, P. 2001. Autonomous bidding agents in the Trading Agent Competition. *IEEE Internet Computing* 5, 2, 52–60.

- GREENWALD, A. R. AND KEPHART, J. O. 1999. Shopbots and pricebots. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, T. Dean, Ed. Stockholm, Sweden, 506–511.
- GROUP, O. M. 2006. BPM 1.0: BPM business process modeling notation specification. Tech. rep., Object Management Group. February. Final adopted specification.
- GROVES, T. 1973. Incentives in teams. *Econometrica* 41, 617–631.
- HANSELL, S. 2004. Google’s slow search for a good share price. New York Times: Business Day.
- HASBROUCK, J., SOFIANOS, G., AND SOSEBESS, D. 1993. New York Stock Exchange systems and trading procedures. Working Paper 93-01, New York Stock Exchange. April.
- HEALEY, C. G., AMANT, R. S., AND CHANG, J. 2001. Assisted visualization of e-commerce auction agents. In *Proceedings of Graphics Interface 2001*. Ottawa, Canada, 201–208.
- HOOS, H. H. AND BOUTILIER, C. 2000. Solving combinatorial auctions using stochastic local search. In *Proceedings of the 17th National Conference on Artificial Intelligence*. AAAI Press, 22–29.
- HUBERMAN, B. A. AND CLEARWATER, S. H. 1995. A multi-agent system for controlling building environments. In *Proceedings of the 1st International Conference on Multiagent Systems*, V. R. Lesser and L. Gasser, Eds. The MIT Press, Cambridge, MA, 171–176.
- HUNSBERGER, L. AND GROSZ, B. 2000. A combinatorial auction for collaborative planning. In *Proceedings of the 4th International Conference on Multiagent Systems*, E. Durfee, S. Kraus, H. Nakashima, and M. Tambe, Eds. IEEE Computer Society, 151–158.
- HURWICZ, L. 1972. On informationally decentralized systems. In *Decision and Organisation: A Volume in Honour of Jacob Marchak*, C. McGuire and R. Radner, Eds. North-Holland.
- HURWICZ, L. 1975. On the existence of allocation systems whose manipulative Nash equilibria are Pareto optimal. Unpublished manuscript.
- HURWICZ, L. AND WALKER, M. 1990. On the generic nonoptimality of dominant-strategy allocation mechanisms: A general theorem that includes pure exchange economies. *Econometrica* 58, 683–704.
- JACKSON, M. O. 2003. Mechanism theory. In *Optimization and Operations Research*, U. Devigs, Ed. The Encyclopedia of Life Support Science. EOLSS Publishers, Oxford, UK.
- JENNINGS, N. AND BUSSMANN, S. 2003. Agent-based control systems. *IEEE Control Systems Magazine* 23, 3, 61–74.
- KAHN, A. E., CRAMTON, P. C., PORTER, R. H., AND TABORS, R. D. 2001. Pricing in the California power exchange electricity market: Should California switch from uniform pricing to pay-as-bid pricing? Blue Ribbon Panel Report. A study commissioned by the California Power Exchange.
- KALAGNANAM, J., DAVENPORT, A., AND LEE, H. 2000. Computational aspects of clearing continuous call double auctions with assignment constraints and indivisible demand. Tech. Rep. RC21660(97613), IBM. February 2nd.
- KALAGNANAM, J. AND PARKES, D. C. 2004. Auctions, bidding and exchange design. In *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*, D. Simchi-Levi, S. D. Wu, and M. Shen, Eds. International Series in Operations Research and Management Science. Kluwer, Chapter 5.
- KELLY, F. AND STENBERG, R. 2000. A combinatorial auction with multiple winners for universal service. *Management Science* 46, 4, 586–596.
- KLEMPERER, P. 1998. Auctions with almost common values: The “wallet game” and its applications. *European Economic Review* 42, 3–5, 757–769.
- KLEMPERER, P. 2002a. Auction theory: A guide to the literature. *Journal of Economic Surveys* 13, 3, 227–286. Reprinted as Chapter 1 of [Klemperer 2004].
- KLEMPERER, P. 2002b. How (not) to run auctions: The European 3G telecom auctions. *European Economic Review* 46, 4/5 (May), 829–845. Reprinted as Chapter 5 of [Klemperer 2004].
- KLEMPERER, P. 2002c. What really matters in auction design. *Journal of Economic Perspectives* 16, 169–189. Reprinted as Chapter 3 of [Klemperer 2004].
- KLEMPERER, P. 2004. *Auctions: Theory and Practice*. Princeton University Press.
- LEE, J. 2002. Making losers of auction winners. The New York Times.
- ACM Journal Name, Vol. V, No. N, Month 20YY.

- LEE, J., BICHLER, M., VERMA, S., AND LEE, H. 2000. Design and implementation of an interactive visual analysis system for e-sourcing. Tech. Rep. RC22045, IBM.
- LEHMANN, D., MÜLLER, R., AND SANDHOLM, T. 2006. The winner determination problem. See Cramton et al. [2006], Chapter 12, 297–317.
- LESKELÄ, R.-L., TEICH, J. E., WALLENIS, H., AND WALLENIS, J. 2007. Decision support for multi-unit combinatorial bundle auctions. *Decision Support Systems* 43, 2, 420–434.
- LÉVY, J. 1967. *The Economic Life of the Ancient World*. University of Chicago Press, Chicago, IL.
- LEYTON-BROWN, K., NUDELMAN, E., AND SHOHAM, Y. 2006. Empirical hardness models for combinatorial auctions. See Cramton et al. [2006], Chapter 19, 479–504.
- LEYTON-BROWN, K. AND SHOHAM, Y. 2006. A test-suite for combinatorial auctions. See Cramton et al. [2006], Chapter 18, 451–478.
- LEYTON-BROWN, K., SHOHAM, Y., AND TENNENHOLTZ, M. 2002. Bidding clubs in first-price auctions. In *Proceedings of the 18th National Conference on Artificial Intelligence*. AAAI Press, San Matteo, CA, 373–378.
- LUCKING-REILEY, D. 2000. Auctions on the internet: What’s being auctioned, and how? *Journal of Industrial Economics* 48, 227–252.
- MACKIE-MASON, J. K. AND VARIAN, H. R. 1994. Generalized Vickrey auctions. Tech. rep., Department of Economics, University of Michigan.
- MASKIN, E. S. AND RILEY, J. G. 1985. Auction theory with private values. *American Economic Review* 75, 150–155.
- MASKIN, E. S. AND SJÖSTRÖM, T. 2002. Implementation theory. In *Handbook of Social Choice Theory and Welfare*, K. J. Arrow, A. K. Sen, and K. Suzumura, Eds. North-Holland, Amsterdam.
- MCAFEE, R. P. AND MCMILLAN, J. 1987. Auctions and bidding. *Journal of Economic Literature* 25, 2, 699–738.
- MCAFEE, R. P. AND MCMILLAN, J. 1996. Analyzing the airwaves auction. *Journal of Economic Perspectives* 10, 159–176.
- MCCABE, K. A., RASSENTI, S. J., AND SMITH, V. L. 1993. Designing a uniform-price double auction: An experimental evaluation. See Friedman and Rust [1993a], 307–332.
- MCMILLAN, J. 1994. Selling spectrum rights. *Journal of Economic Perspectives* 8, 191–199.
- MENEZES, F. 1996. Multiple-unit English auctions. *European Journal of Political Economy* 12, 4, 671–684.
- MILGROM, P. 1989. Auctions and bidding: A primer. *Journal of Economic Perspectives* 3, 3, 3–22.
- MILGROM, P. 2000. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy* 108, 245–272.
- MILGROM, P. AND WEBER, R. 1982. A theory of auctions and competitive bidding. *Econometrica* 50, 1089–1122.
- MONDERER, D. AND TENNENHOLTZ, M. 1999. Distributed games. *Games and Economic Behavior* 28, 1, 55–72.
- MONDERER, D. AND TENNENHOLTZ, M. 2000. k -price auctions. *Games and Economic Behavior* 31, 1–2, 220–244.
- MORITA, T., YUITOU, M., HIDAKA, T., AND HIRAKAWA, Y. 2005. Visualization methods using three-dimensional space for auctions. In *The 2005 Symposium on Applications and the Internet Workshops*. 396–399.
- MU’ALEM, A. AND NISAN, N. 2002. Truthful approximation mechanisms for restricted combinatorial auctions. In *Proceedings of the 18th National Conference on Artificial Intelligence*. AAAI Press, San Matteo, CA, 379–384.
- MÜLLER, R. 2006. Tractable cases of the winner determination problem. See Cramton et al. [2006], Chapter 13, 319–336.
- MURNIGHAN, J. K. 1992. *Bargaining Games*. William Murrow and Company.
- MYERSON, R. B. 1981. Optimal auction design. *Mathematics of Operation Research* 6, 1, 58–73.

- NARUMANCHI, M. V. AND VIDAL, J. M. 2006. Algorithms for distributed winner determination in combinatorial auctions. In *Agent-Mediated Electronic Commerce. Designing Trading Agents and Mechanisms*. Lecture Notes in Computer Science, vol. 3937. Springer-Verlag, 43–56.
- NISAN, N. 2000. Bidding and allocation in combinatorial auctions. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*. ACM Press, Minneapolis, Minnesota, United States, 1–12.
- NISAN, N. 2006. Bidding languages for combinatorial auctions. See Cramton et al. [2006], Chapter 9, 215–232.
- NISAN, N. 2007. Introduction to mechanism design (for computer scientists). In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, Eds. Cambridge University Press, Cambridge, UK.
- NISAN, N. AND RONEN, A. 2007. Computationally feasible VCG mechanisms. *Journal of Artificial Intelligence Research* 29, 19–47.
- NIU, J., CAI, K., GERDING, E., MCBURNEY, P., AND PARSONS, S. 2008. Characterizing effective auction mechanisms: Insights from the 2007 TAC market design competition. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, L. Padgham, D. Parkes, J. Müller, and S. Parsons, Eds. IFAAMAS, 1079–1086.
- OSBORNE, M. J. AND RUBINSTEIN, A. 1994. *A Course in Game Theory*. MIT Press, Cambridge, MA.
- OSTWALD, J., LESSER, V., AND ABDALLAH, S. 2005. Combinatorial auction for resource allocation in a distributed sensor network. In *RTSS '05: Proceedings of the 26th IEEE International Real-Time Systems Symposium*. IEEE Computer Society, 266–274.
- PARKES, D. 1999. iBundle: An efficient ascending price bundle auction. In *Proceedings of the 1st ACM Conference on Electronic Commerce*. ACM Press, New York, 148–157.
- PARKES, D. 2000. Optimal auction design for agents with hard valuation problems. In *Agent Mediated Electronic Commerce II: Towards Next-Generation Agent-Based Electronic Commerce Systems*, F. Y. A. Moukas, C. Sierra, Ed. LNAI, vol. 1788. Springer Verlag, Berlin, 206–219.
- PARKES, D. 2001. Iterative combinatorial auctions: Achieving economic and computational efficiency. Ph.D. thesis, University of Pennsylvania, Department of Computer and Information Science.
- PARKES, D., CAVALLO, R., ELPRIN, N., JUDA, A., LAHAIE, S., LUBIN, B., MICHAEL, L., SHNEIDMAN, J., AND SULTAN, H. 2005. ICE: An iterative combinatorial exchange. In *Proceedings of the 6th ACM Conference on Electronic Commerce*. ACM Press, New York, 249–258.
- PARKES, D. AND KALAGNANAM, J. 2005. Models for iterative multiattribute procurement auctions. *Management Science* 51, 435–451.
- PARKES, D., KALAGNANAM, J. R., AND ESO, M. 2001. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *Proceedings 17th International Joint Conference on Artificial Intelligence*, B. Nebel, Ed. Seattle, WA, 1161–1168.
- PARKES, D. C. AND DUONG, Q. 2007. An ironing-based approach to adaptive online mechanism design in single-valued domains. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*. AAAI Press, San Matteo, CA, 94–101.
- PARSONS, S. 2002. Exception analysis for double auctions. Research Note, Center for Coordination Science, Sloan School of Management, Massachusetts Institute of Technology.
- PARSONS, S., MARCINKIEWICZ, M., NIU, J., AND PHELPS, S. 2008. Everything you wanted to know about double auctions, but were afraid to (bid or) ask. Technical report, Department of Computer and Information Science, Brooklyn College.
- PATEL, J., TEACY, W. T. L., JENNINGS, N. R., LUCK, M., CHALMERS, S., OREN, N., NORMAN, T. J., PREECE, A. D., GRAY, P. M. D., SHERCLIFF, G., STOCKREISSER, P. J., SHAO, J., GRAY, W. A., FIDDIAN, N. J., AND THOMPSON, S. G. 2005. Agent-based virtual organisations for the grid. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multi-agent Systems*, F. Dignum, V. Dignum, S. Koenig, S. Kraus, M. P. Singh, and M. Wooldridge, Eds. 1151–1152.
- PESENDORFER, M. 2000. A study of collusion in first-price auctions. *Review of Economic Studies* 67, 381–411.
- ACM Journal Name, Vol. V, No. N, Month 20YY.

- PETCU, A., FALTINGS, B., AND PARKES, D. C. 2006. MDPOP: faithful distributed implementation of efficient social choice problems. In *Proceedings of the 5th International Conference on Autonomous Agents and Multiagent Systems*, H. Nakashima, M. P. Wellman, G. Weiss, and P. Stone, Eds. ACM Press, New York, 1397–1404.
- PHELPS, S. 2008. Evolutionary mechanism design. Ph.D. thesis, University of Liverpool, Department of Computer Science.
- PHELPS, S., CAI, K., MCBURNEY, P., NIU, J., PARSONS, S., AND SKLAR, E. 2008. Auctions, evolution, and multi-agent learning. In *Adaptive Agents and Multi-Agent Systems III*, K. Tuyls, A. Nowe, Z. Guessoum, and D. Kudenko, Eds. LNAI, vol. 4865. Springer Verlag, Berlin.
- PORTER, R., RONEN, A., SHOHAM, Y., AND TENNENHOLTZ, M. 2002. Mechanism design with execution uncertainty. In *Proceedings of the 18th Conference in Uncertainty in Artificial Intelligence*, A. Darwiche and N. Friedman, Eds. Morgan Kaufmann, San Francisco, CA.
- PORTER, R. H. AND ZONA, J. D. 1993. Detection of bid rigging in procurement auctions. *Journal of Political Economy* 101, 3, 518–538.
- PRICEWATERHOUSECOOPERS. 2001. E-markets: Realism, not pessimism.
- PRINGSHEIM, F. 1949. The Greek sale by auction. In *Scritti in Onore di Contardo Ferrini Pubblicati in Occasione della sua Beatificazione*. Pubblicazioni dell' Università Cattolica del Sacro Cuore, Nuova Serie, vol. XXVIII. Societa Editrice "Vita e Pensiero", Milan, 284–343.
- RAMCHURN, S. D., MEZZETTI, C., GIOVANNUCCI, A., RODRIGUEZ-AGUILAR, J. A., DASH, R. K., AND JENNINGS, N. R. 2009. Trust-based mechanisms for robust and efficient task allocation in the presence of execution uncertainty. *Journal of Artificial Intelligence Research* 35, 1, 119–159.
- RAMCHURN, S. D., ROGERS, A., MACARTHUR, K., FARINELLI, A., VYTELINGUM, P., VETSIKAS, I., AND JENNINGS, N. R. 2008. Agent-based coordination technologies in disaster management. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, L. Padgham, D. Parkes, J. Müller, and S. Parsons, Eds. IFAAMAS, 1651–1652.
- RILEY, J. G. AND SAMUELSON, W. F. 1981. Optimal auctions. *American Economic Review* 71, 381–392.
- ROTH, A. E. AND OCKENFELS, A. 2000. 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.
- RUST, J., MILLER, J. H., AND PALMER, R. 1994. Characterizing effective trading strategies. *Journal of Economic Dynamics and Control* 18, 61–96.
- SANDHOLM, T. 2002. An algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence* 135, 1–54.
- SANDHOLM, T. 2003. Automated mechanism design: A new application area for search algorithms. In *Principles and Practice of Constraint Programming*, F. Rossi, Ed. Lecture Notes in Computer Science, vol. 2833. Springer, Berlin, Germany, 19–36.
- SANDHOLM, T. 2006. Optimal winner determination algorithms. In *Combinatorial Auctions*, P. Cramton, Y. Shoham, and Steinberg, Eds. MIT Press, Cambridge, MA.
- SANDHOLM, T. AND BOUTILIER, C. 2006. Bidding languages for combinatorial auctions. See Cramton et al. [2006], Chapter 10, 233–263.
- SANDHOLM, T. AND GILPIN, A. 2006. Sequences of take-it-or-leave-it offers: near-optimal auctions without full valuation revelation. In *Proceedings of the 5th International Conference on Autonomous Agents and Multiagent Systems*, H. Nakashima, M. P. Wellman, G. Weiss, and P. Stone, Eds. ACM Press, New York, 1127–1134.
- SANDHOLM, T. AND SURI, S. 2006. Side constraints and non-price attributes in markets. *Games and Economic Behaviour* 55, 321–330.
- SANDHOLM, T., SURI, S., GILPIN, A., AND LEVINE, D. 2002. Winner determination in combinatorial auction generalizations. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM Press, New York, 69–76.
- SATTERTHWAITE, M. A. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10, 187–217.

- SATTERTHWAITE, M. A. AND WILLIAMS, S. R. 1993. The Bayesian theory of the k -double auction. See Friedman and Rust [1993a], Chapter 4, 99–123.
- SHOHAM, Y. 2000. A survey of auction types. Lecture notes: Stanford University CS206—Technical Foundations of Electronic Commerce.
- SHOHAM, Y. 2001a. Auctions on the internet: what’s actually happening. Lecture notes: Stanford University CS206—Technical Foundations of Electronic Commerce.
- SHOHAM, Y. 2001b. Combinatorial auctions. Lecture notes: Stanford University CS206—Technical Foundations of Electronic Commerce.
- SHOHAM, Y. 2001c. The zoology of auctions. Lecture notes: Stanford University CS206—Technical Foundations of Electronic Commerce.
- SHOHAM, Y. 2002. Combinatorial auctions. Lecture notes: Stanford University CS206—Technical Foundations of Electronic Commerce.
- SHOHAM, Y. AND LEYTON-BROWN, K. 2008. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York.
- SIERRA, C., DE MANTARAS, R. L., AND BUSQUETS, D. 2000. Multiagent bidding mechanisms for robot qualitative navigation. In *Intelligent Agents VII*. Lecture Notes in Artificial Intelligence, vol. 1986. 198–212.
- SMITH, V. L. 1962. An experimental study of competitive market behaviour. *The Journal of Political Economy* 70, 2 (April), 111–137.
- SMITH, V. L. 1982. Microeconomic systems as an experimental science. *American Economic Review* 72, 5, 923–955.
- SONG, J. AND REGAN, A. 2002a. Approximation algorithms for the bid construction problem in combinatorial auctions for the procurement of freight transportation contracts. Tech. rep., University of California, Irvine. ITS-CUI Working Paper.
- SONG, J. AND REGAN, A. C. 2002b. Combinatorial auctions for transportation service procurement: the carrier perspective. *Transportation Research Record* 1833, 40–46.
- STARK, R. AND ROTHKOPF, M. H. 1979. Competitive bidding: A comprehensive bibliography. *Operations Research* 27, 364–391.
- STROBEL, M. AND WEINHARDT, C. 2003. The Montreal taxonomy of electronic negotiations. *Group Decision and Negotiation* 12, 143–164.
- SUYAMA, T. AND YOKOO, M. 2005. Strategy/false-name proof protocols for combinatorial multi-attribute procurement auction. *Autonomous Agents and Multi-Agent Systems* 11, 1, 7–21.
- TESAURO, G. AND DAS, R. 2001. High-performance bidding agents for the continuous double auction. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*. ACM Press, New York, 206–209.
- THOMAS, J. A. C. 1957. The auction sale in Roman law. *The Judicial Review*, 42–66.
- UCKELMAN, J. AND ENDRISS, U. 2008. Winner determination in combinatorial auctions with logic-based bidding languages. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, L. Padgham, D. Parkes, J. Müller, and S. Parsons, Eds. IFAA-MAS, 1617–1618.
- UNGERN-STERNBERG, T. V. 1991. Swiss auctions. *Econometrica* 58, 341–357.
- VARIAN, H. R. 1995. Economic mechanism design for computerized agents. In *Proceedings of the USENIX Workshop on Electronic Commerce*. USENIX, New York. Minor typos fixed 3rd March 2000.
- VICKREY, W. 1961. Counterspeculation, auctions, and competitive sealed bids. *Journal of Finance* 16, 1, 8–37.
- VINYALS, M., GIOVANNUCCI, A., CERQUIDES, J., MESEGUER, P., AND RODRIGUEZ-AGUILAR, J. A. 2008. A test suite for the evaluation of mixed multi-unit combinatorial auctions. *Journal of Algorithms* 63, 1-3, 130–150.
- WALTON, K. 2006. *Fake: Forgery, Lies & eBay*. Simon Spotlight Entertainment, New York.
- WANG, X. AND XIA, M. 2005. Combinatorial bid generation problem for transportation service procurement. *Transportation Research Record: Journal of the Transportation Research Board* 1923, 189–198.
- ACM Journal Name, Vol. V, No. N, Month 20YY.

- WELLMAN, M. 1993. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research* 1, 1–23.
- WELLMAN, M. P., GREENWALD, A., STONE, P., AND WURMAN, P. R. 2002. The 2001 trading agent competition. In *Proceedings of the 18th National Conference on Artificial Intelligence*. American Association for Artificial Intelligence, Edmonton, Alberta, Canada, 935–941.
- WELLMAN, M. P., WURMAN, P. R., O'MALLEY, K., BANGERA, R., LIN, S., REEVES, D., AND WALSH, W. E. 2001. Designing the market game for a trading agent competition. *IEEE Internet Computing* 5, 2, 43–51.
- WOLFSTETTER, E. 1996. Auctions: An introduction. *Journal of Economic Surveys* 10, 4, 367–420.
- WURMAN, P. AND WELLMAN, M. 2000. AkBA: a progressive, anonymous-price combinatorial auction. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*. Minneapolis, 21–29.
- WURMAN, P. R., WALSH, W. E., AND WELLMAN, M. P. 1998. Flexible double auctions for electronic commerce: Theory and applications. *Decision Support Systems* 24, 17–27.
- WURMAN, P. R., WELLMAN, M. P., AND WALSH, W. E. 2001. A parameterization of the auction design space. *Games and Economic Behavior* 35, 1/2, 304–338.
- XIA, M., STALLAERT, J., AND WHINSTON, A. B. 2004. Solving the combinatorial double auction problem. *European Journal of Operation Research* 164, 1, 234–251.
- XU, L., HUTTER, F., HOOS, H., AND LEYTON-BROWN, K. 2008. SATzilla: Portfolio-based algorithm selection for SAT. *Journal of Artificial Intelligence Research* 32, 565–606.
- YOKOO, M., SAKURAI, Y., AND MATSUBARA, S. 2001. Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence* 130, 2, 167–181.
- YOKOO, M., SAKURAI, Y., AND MATSUBARA, S. 2004. The effect of false-name bids in combinatorial auctions: New fraud in internet auctions. *Games and Economic Behaviour* 46, 1, 174–188.
- ZUREL, E. AND NISAN, N. 2001. An efficient approximation algorithm for combinatorial auctions. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*. ACM Press, New York, 125–136.