Tilburg University

Auction theory for auction design

van Damme, E.E.C.; Börgers, T.

Published in:
Auctioning Public Assets. Analysis and Alternatives

Publication date:
2004

Document Version
Peer reviewed version

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
van Damme, E. E. C., & Börgers, T. (2004). Auction theory for auction design. In M. Janssen (Ed.), *Auctioning Public Assets. Analysis and Alternatives* (pp. 19-64). Cambridge University Press.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Auction Theory for Auction Design^{*}

Tilman Börgers^{**} and Eric van Damme^{***}

First version January 2003

Final Revision, August 2003

Keywords: auctions, market design, government policy

JEL-codes: D44, C72, L50

^{*} Paper to appear as Chapter 1 in the book “Auctioning Public Assets” edited by Maarten Janssen, to be published by Cambridge University Press. We are grateful to Maarten for inviting us to write this survey. Preparing the survey has made us painfully aware of the gaps in our knowledge of the auction literature and we apologize for all errors and omissions. Tilman Börgers gratefully acknowledges financial support from the “Economic and Social Research Council” (UK).

^{**} Prof.dr. Tilman Börgers, Department of Economics and ESRC ELSE Centre, University College London, Gower Street, London WC1E 6BT, U.K. Phone +44-20-7679 5879, Fax +44-20-7916 2774, e-mail: t.borgers@ucl.ac.uk, <http://www.ucl.ac.uk/~uctpa01/borgers.htm>

^{***} Prof.dr. E.E.C. van Damme, Tilburg University, CentER for Economic Research and Tilburg Law and Economics Center (TILEC), P.O. Box 90153, 5000 LE Tilburg, The Netherlands. Phone +31-13-4663045, Fax +31-13-4663266, e-mail Eric.vanDamme@uvt.nl, <http://center.uvt.nl/staff/vdamme/>

1 Introduction

In this introductory survey we review research papers on auction theory that may be of relevance to the design of auctions of government assets in general, and of spectrum license auctions in particular. We focus on the main intuitions emerging from these papers, and refer to the original papers for technical details.

We begin in Section 2 with a discussion of why economists typically favour auctions over other methods for allocating licenses to operate in a market. In Section 3, we have a first discussion on auction design, stressing the fact that a seller will typically face a much more complicated problem than just what auction form to use; he also has to think carefully about what to sell, who to allow as bidders and when to sell. Of course, the solution to these problems will also depend on what goal is to be achieved. Assuming these problems are solved, we turn, in Section 4 to an exposition of auction formats. We start the discussion with the simple case in which the seller has just one indivisible object for sale, for which we describe the four basic auction forms: two open auctions, the English (or ascending) auction and the Dutch (or descending) auction, and two sealed-bid formats, the first price auction and the second price (or Vickrey) auction. In the second part of the section, we then show how these auction formats can be extended to deal with the situation in which the seller has available multiple units of the same object, or multiple objects. In this process, we will encounter a large variety of auction formats. In Section 5 we discuss these various auction formats from the bidders' perspective: what strategies could one expect the competitors to follow and how should one bid oneself? We also discuss the implications of rational bidding strategies for the variables which the seller probably cares about, such as efficiency and revenue. In Section 6 we pull our insights together, and ask which policy lessons our analysis suggests. Section 7 concludes.

Already at the outset, we wish to stress the limitations of this paper: the reader should be aware that this is a theory paper. Theory alone has no policy implications; it needs to be combined with empirical analysis (of field data, or experimental data) before policy recommendations can be derived. In the present paper, empirical or experimental evidence is cited where it is particularly prominent, but it is not

surveyed systematically. Therefore, what we say in this paper does not in itself provide a basis for policy recommendations. Put differently, any policy implication that is derived from the theory exposited in this paper should be prefaced with the qualification: ‘if the theory captures practice well, then policy should be’.

2 Why Auctions?

Governments allocating spectrum licenses to mobile telephony companies or, more generally, licenses to operate on a market, have a variety of methods at their disposal. The traditionally most popular method has been the Beauty Contest where companies are invited to submit business plans, and a government agency selects those companies whose business plans seem most credible, and which are most likely to deliver services that the government believes to be valuable. In recent years, auctions have been the most popular method. What is the rationale for using auctions?

To answer this question, we wish to make a distinction between auctions as used in the private sector and auctions used by the government. We first discuss why a private-sector seller may prefer to dispose of an item by means of an auction. Next, we consider which of these arguments also apply when it is the government that acts as a seller.

A seller of a unique item would typically want to get the best price for the item, hence, the question is what selling mechanism would result in the highest expected price. If the seller would know what each interested buyer would be willing to pay for the item, his problem would be trivial: he would simply make a ‘take-it-or-leave-it’ offer to the buyer with the highest willingness to pay. Of course, in actual practice, the seller does not have the required information, and in these circumstances, he may set the price too low, in which case he does not expropriate what the market can bear, or he may set the price too high, so that he does not succeed in selling the item.

An ascending auction then provides an attractive alternative. In such an auction, each buyer is willing to bid as long as the price is lower than the bidder’s reservation value,

hence, bidding will continue until the second highest reservation value is reached, and the ultimate price will be this second highest value. The seller thus does worse than with complete information, but typically he does better than by making a 'take-it-or-leave-it' offer. Moreover, when the number of bidders is large, the auction performs almost as well as the seller could have performed had he complete information. This is the main reason why auctions are attractive mechanisms for private sellers: they extract good prices even if the seller is poorly informed about individual buyers' willingness to pay.

As a possible selling mechanism, a private-sector seller may also consider to negotiate with potential buyers. He might hope to learn buyers' true willingness to pay by observing their strategic moves in the negotiation. But, of course, buyers will anticipate in a negotiation that they will be closely watched. They will be very weary of giving too much away too early. Bids in an auction might also give information away, but as long as the seller's commitment to the auction mechanism is firm, bidders know in advance how their bids are going to be used in the allocation process. They do not have to worry about concealing information. Therefore, auctions encourage more information revelation by buyers, and it is this information revelation that is needed for a successful sale. Furthermore, auctions may attract more interested parties than negotiation processes, and Bulow and Klemperer (1996) have shown that, under certain assumptions, an auction without a reserve price, as long as it attracts at least one more bidder than a negotiation, raises more expected revenue than any negotiation procedure.

In the context described above, the ascending auction has another very attractive property: it results in an efficient allocation, i.e. the auction allocates the object to that bidder who values it most. It is this property that makes auctions an attractive selling mechanism also for governments. Just as a seller in the private sector, a government seller typically is uncertain about how much bidders are willing to pay for the items that it sells, but, in contrast to private sellers, governments may not be primarily interested in raising revenues, but in achieving an efficient outcome of some sort. (See Section 3 for a brief discussion on the goals of the government and for why also a government might be interested in raising revenues.) The above argument suggests that it still might be a good idea to auction as the auction may produce an efficient

outcome. Indeed, the case for using auctions to sell licenses has usually been based on the twin arguments that an auction is an efficient procedure (i.e. it is quick, transparent, not very susceptible to lobbying, and reasonably proof to legal action) that produces an efficient outcome, see McMillan (1994).

One should point out, however, that the efficiency argument in favour of auctions is not as strong as it might appear at first. First of all, when there are ‘frictions’, the efficiency property need not hold; for example, if the person with the highest value faces a binding budget constraint at a level lower than the second highest value, the bidder with the second highest value will win, see Krishna (2002) for some results on auctions in which bidders are budget constrained. Secondly, and in particular in the case of a government seller, one should be very careful with what one means by ‘efficiency’: one should be aware that ‘economic efficiency’ is not equivalent to ‘the licenses ending up in the hands of those that value them most’. As Janssen and Moldovanu show in detail in chapter 5 in this book, the reason lies in all kinds of externalities that exist in license auctions. The main externality is that a benevolent government will sell the licenses (also) having the consumer welfare in mind; consumers, however, are not participating directly in the auction. As a result, the outcome in which the license is put in the hands of the firm that values it most, may not be the one that consumers prefer. In fact, the preferences of the consumers may be exactly opposite.

As a specific example, based on Gilbert and Newbery (1982), suppose that a government sells a second license to operate in a market in which already one player is active. For a newcomer, the license represents the right to compete, while for the incumbent it offers the opportunity to maintain his monopoly. Since the incumbent’s profit loss from losing the monopoly is typically larger than the entrant’s gain in profit from being allowed to compete, the monopolist will win an auction for the second license. An auction will, hence, allocate the license to the monopolist and will not produce a competitive outcome. As a competitive outcome yields higher economic efficiency (total welfare) than a monopolistic outcome, an ordinary auction will not achieve the efficiency goal.

In order to nevertheless reach an efficient outcome in this asymmetric situation, the government might use an auction variant, for example, the government might simply ban the incumbent from the auction of the second license. In this case, one of the entrants is sure to win and this 'asymmetric auction' might attract more bidders and might result in higher revenue than the auction in which the incumbent is allowed to bid and in which entrants know that they cannot win. More sophisticated 'discriminatory auctions' can have the same effect and we refer to chapter 4 by Maasland et al. for further discussion, in particular about whether such auctions might violate basic EU-principles such as involving discrimination or state aid. The point here, however, is more general: if an ordinary auction does not produce the desired result, then one may adjust the auction rules to obtain an outcome that one likes. Auctions are an extremely flexible allocation mechanism, hence, they allow a government considerable freedom of action.

Just as price setting or negotiations are alternative selling mechanisms for private-sector sellers, the Beauty Contest is typically the alternative selling mechanism considered by governments for the allocation of government assets. In such a Beauty Contest, bidders describe in detail what they plan to do with the license, with the government then selecting the best plan. There are, perhaps, two main concerns which economists have about Beauty Contests. One is that the commitments made by bidders in Beauty Contests are hard to enforce. If bidders anticipate this enforcement problem, then they can promise arbitrary things, and there is no guarantee that the winners are really those who make best use of the objects for sale. The second concern is that, given the discretion and subjective element in Beauty Contest, there might be more potential for corruption of government officials in a Beauty Contest than in an auction. We refer to the Introduction of this book for more details on this issue.

Summarizing the above, we may state that auctions have certain desirable properties that alternative allocation mechanisms do not have and that, therefore, an auction may be preferred whenever allocation by means of an auction is feasible. This, however, does not imply that any auction will do and that auctions do not have any drawbacks. In the remainder of this paper, we will show that the choice of auction may be of great

importance and that ‘side constraints’ in the auction may be needed in order to ensure that a desirable outcome is reached.

3 Pre-Auction Decisions

When a government is selling assets or licenses, a large number of design questions have to be addressed. First of all, the government should be clear about the *goals* that it wants to achieve. For example, should the government try to maximize revenue, or should it aim for market efficiency? One argument for why governments might be concerned about auction revenues is that such revenues may allow governments to reduce more distorting taxes elsewhere in the economy. However, efficiency is typically the dominant goal of governments.

The efficiency goal is sometimes identified with the objective of ‘placing licenses into the hands of those that value them most’. This is not always the same as efficiency, though. For a general discussion on this important point, we refer to chapter 5 by Janssen and Moldovanu. An example was already given in the previous section. As another example, think of a government selling licenses to operate radio stations. Under quite natural and general conditions, stations that broadcast ‘middle of the road music’ will be willing to pay most for these licenses; all stations broadcasting similar music, however, will typically not be an efficient outcome. In such a case, if the government wants to achieve an efficient outcome, it should impose conditions on some of the licenses, which will typically reduce revenue. This example, hence, also shows that the different goals that the government may want to pursue (efficiency and revenue) may be in conflict.

Once government objectives are clear, the next important question is: ‘*What* will be sold?’ An example where this clearly mattered are the recent European UMTS-auctions. There the question was: ‘How large (in terms of spectrum) should a UMTS license be?’ It was not clear how much spectrum a UMTS operator would need, and therefore, how many licenses could be fitted into the available spectrum, hence, how many players there would be in the resulting market. While most countries simply

fixed this number in advance, Germany and Austria dealt with this difficulty in a different way. These countries decided not to auction licenses, but rather abstract blocks of spectrum, and bidders could choose themselves for how many blocks they wanted to bid. The key idea behind these auction designs was that the mechanism not only helped governments to discover which companies should hold licenses, but also how much spectrum was actually needed for third generation spectrum licenses, hence, to discover for how many companies there was space in the spectrum.

An interesting objection has been raised in the academic literature against this innovative approach: it is that companies' bids in these auctions will not primarily reveal to governments how much companies value extra spectrum, and thus what the optimal size of a license is, but how much companies value monopoly power (see Jehiel and Moldovanu (2000b)). This is because bidders will understand that the future market structure emerges endogenously from the auction. By buying up spectrum a bidder can reduce the amount of spectrum available to others, and, in particular, a bidder can prevent others from entering the market. Thus, bids in these auctions might not be related at all to the true value of the spectrum, and instead might indicate which value the bidder attaches to a reduction in the number of competitors in the market. If this argument is accepted, then it appears better to make a possibly imperfect judgment about the optimal size of licenses, and to let the auction only determine who gets which. It should be added that in practice this argument has not appeared to be of much relevance to the German and Austrian auctions. The precise reasons for this are unclear, and it is worth keeping this argument in mind for future auctions.

The above example also indicates that relatively frequently a government may need to build additional regulatory constraints in the auction. This need especially arises in the situation of franchise bidding where the government awards the right to provide a service to that party that is willing to do it for the lowest compensation, and where the auction results in the license winners enjoying market power on the ensuing market. In these situations, part of the compensation is paid before any service is delivered and the government has to ensure that the service is indeed delivered and is of the quality that has been promised and agreed upon. Elaborate contracts and extensive

monitoring may be needed in this case of ‘moral hazard’; Williamson (1976) gives a good overview of the difficulties and the trade-offs involved.

Another issue to be considered before an auction is who should be allowed to participate. For bids in an auction to be credible, bidders must be financially respectable, and most government auctions include an appropriate screening of bidders. Requiring deposits form another safeguard against non-serious bids. In some cases one may go further in restricting the set of admissible bidders. For example, if licenses to operate in a particular industry are auctioned, then one may wish to exclude incumbents from the auction, either to ensure that the post-auction market is more competitive, or simply to attract more entry into the auction.

Next, also the *timing* of auctions is important. Consider again the experience with the European UMTS-auctions. Governments that were earlier in auctioning their licenses have typically earned (much) higher revenue per capita than countries that were later. The UK was the first country to auction its licenses and in effect, therefore, the UK was not only auctioning a license to operate in the UK, but the option to construct a pan-European network. This option attached to the UK-license might have made the UK-license more valuable, it might have attracted more bidders to the UK-auction, with higher revenues as a natural consequence. Similarly, if the German UMTS-auction would have taken place later in time, the tide might have turned and the Sonera/Telefonica consortium might have realised that a 6-player German market was not viable and not profitable for them; in that case, German revenue could have been much lower. While it might have been beneficial for revenues to hold auctions earlier, it might have been beneficial for efficiency to hold them later. As time progressed, more information about UMTS technology, and the corresponding handset technology, became available, and thus efficiency became more feasible.

In essence, all of the above arguments amount to saying that the outcome is determined by supply and demand conditions, and that the government can influence both of these. Perhaps less obvious at first is the fact that the outcome will also depend on the market mechanism, the auction format, that is used. Therefore, we now turn to a discussion of auction formats.

4 Auction formats

We now assume that the questions ‘what to sell?’, ‘when to sell?’ and ‘who to allow to bid?’ have been answered, and we focus on the question ‘how to auction?’. While our main interest is in describing auction mechanisms that can be used for selling multiple identical or heterogeneous objects, we start with the simplest case in which there is just one object for sale.

4.1 Selling a single object

Two types of auctions can be distinguished: auctions can be open or sealed bid. In sealed bid formats, bidders simultaneously and independently submit a bid, possibly in a sealed envelope, or perhaps using a more modern communication technique. Then these bids are opened and the auction outcome is determined following some rules that have been announced in advance. In an *open auction procedure*, bidding proceeds in stages in real time. In each round, bidders act simultaneously and independently; at the end of each round, all bidders observe the outcome of that round, and then adjust their bids on the basis of what they have seen so far.

The best-known open procedure is the ascending, or English auction, in which the price is raised until just one bidder is left. This bidder then wins the object at the price at which the ultimate competitor dropped out. In practise, one observes a large diversity of English auction forms: the auctioneer may announce successive prices, or the initiative for calling out prices may lie with the bidders themselves; bidders may know which competitors are still in the race, or they may not have this information, etc.

A second open procedure is the Dutch, or descending auction in which the auctioneer lowers the price until one of the bidders shouts ‘mine’ or pushes a button on his computer terminal. The (first) bidder to stop the auction clock wins the object and pays the price where he stopped the clock. Note the important distinction with the English auction: in the English auction, the winner pays a price that is determined by

his strongest competitor; in the Dutch auction, the winner pays a price determined by himself.

In *sealed bid procedures* bidders bid only once; they simultaneously communicate their bids to the auctioneer. Any reasonable auction format will allocate the object to the bidder who has made the highest bid, however, there is a variety of ways in which the price can be determined, with different corresponding auction formats.

The easiest rule for determining the payment by the winning bidder is, of course, that he has to pay his own bid. This is also the most common sealed-bid procedure, and we will refer to it as the ‘first-price sealed bid auction’. The reader may notice that this procedure bears a strong resemblance to the Dutch auction procedure. After all, in the Dutch auction, each bidder also has to decide on just one number: the price at which he will stop the auction clock. Calling the latter price the player’s ‘bid’, we see that, in the Dutch auction, the highest bidder wins and pays his bid. Consequently, the Dutch auction is equivalent to the first price sealed bid auction.

There is, however, at least one important alternative to the ‘pay your bid’ rule: the successful bidder may be required to pay the highest unsuccessful bid. This sealed bid auction format is called the ‘second price auction’, or the Vickrey auction, after William Vickrey, a winner of the Nobel Prize in Economics, who proposed it; see Vickrey (1961). As in both this auction and in the English auction, the winner pays a price that is determined by his strongest competitor, these two formats are related to each other. The ‘second price sealed-bid’ format, however, is not fully equivalent to the English auctions that are being used in real life; a crucial difference is that the ascending price format allows bidders to observe the drop out points of other bidders, which might be valuable information. Therefore, one needs to study the ascending price auction separately from the second price sealed bid auction.

Of course, open auctions and sealed bid auctions are only two extreme types of auctions and it is easy to conceive of intermediate forms. One important intermediate form is the ‘Anglo-Dutch’ format (see Binmore and Klemperer (2002)). Under this format, an open ascending auction takes place first, until the number of remaining bidders reaches a certain threshold. Then a sealed bid ‘first price’ auction is conducted

with the remaining bidders. This auction format bears a certain resemblance to the way real estate is auctioned in the Netherlands. Usually this is done by means of a pair of auctions: an English auction followed (a week or so later) by a Dutch auction. In contrast to the Anglo-Dutch format, however, the first auction in this case only stops when one bidder is left and everybody can participate in the second auction. The price resulting from the first auction determines the reserve price of the second auction and, if the price resulting in this second auction is higher than in the first, the winner of the first auction receives a certain percentage of the winning bid.

4.2 Selling multiple units

When selling multiple units of the same object, a first choice to be made is whether the units will be sold sequentially, i.e. one after the other, or simultaneously, i.e. all at the same time. When using a sequential auction, one has to decide which auction form will be used at each stage. This might be any of the auction forms that have been discussed above. For example, in the Dutch flower auction in Aalsmeer, flowers are sold by means of a sequence of Dutch auctions.

Our emphasis here will be on simultaneous auctions. As before, one may distinguish between open and sealed bid auctions. Two prominent open formats are the descending price format and the ascending price format. An ascending price format involves a gradually increasing price, with bidders indicating how many units they want at each price, and the auction closing once the number of units requested by the remaining bidders is equal to the number of available units. All bidders then have to pay the price at which the auction closed. As in the single unit case, the price where a bidder reduces his demand may reveal important information to the competing bidders.

Formally, in the ascending price, or English, auction, the auctioneer gradually and continuously raises the price. At each price p , each bidder i indicates his demand $d_i(p)$, i.e. he informs the auctioneer about how many units he would like to have at this price. The auctioneer then calculates total demand

$$d(p) = \sum_j d_j(p) \quad (1)$$

and compares total demand with total supply s . Prices are increased until a price p^* is reached where $d(p^*) = s$ and each bidder i then is allocated $d_i(p^*)$ units for a price p^* for each. Hence, all units sell at the same price. In actual practise, different variants may be distinguished: bidders may, or may not, know the demand as expressed by their competitors; they may, or may not, be allowed to increase their demand again after they have first reduced it, etc.

In the descending price, or Dutch, auction, the price starts at a relatively high level and is then gradually lowered. At each price p , bidders will be informed about the supply $s(p)$ that is still left and they have to indicate when the price has reached a level at which they are willing to buy one or more units. The auction closes when as many bidders have indicated their willingness to bid as there are items available, i.e. when $s(p) = 0$. Each bidder has to pay the price at which he indicated that he was willing to buy. In this case, when bidder i buys three units, say at prices p_1 , p_2 and p_3 , he pays a price p_1 for the first unit, p_2 for the second unit and p_3 for the third unit. Hence, this auction form is discriminatory: different units (might) sell for different prices.

Each of the above auction formats has a related sealed bid version. In sealed bid auction formats, bids take the form of demand curves: bidders indicate separately how much they are willing to pay for the first unit they acquire, how much they are willing to pay for the second unit, etc. Typically the outcome of the auction is determined by finding first the price at which demand equals supply. All bids made above this price are satisfied, with a tie-breaking rule specifying which bids at the market-clearing price will be satisfied as well. Different sealed bid auction formats differ with respect to the precise rules that determine bidders' payments. In a 'uniform price auction' the market-clearing price is also the price that all bidders have to pay for all units that they have been allocated. In a 'discriminatory price auction' bidders have to pay for each unit exactly how much they bid.

Formally, in the uniform price auction, each bidder i communicates directly his entire demand curve $d_i(\cdot)$ to the auctioneer. The auctioneer then computes total demand $d(\cdot)$, as well as the market clearing price p^* for which $d(p^*) = s$. Each bidder i is then allocated $d_i(p^*)$ units for which he pays $p^* d_i(p^*)$ in total. When the number of units is an integer, say n , two variants may be distinguished: the market clearing price may be the lowest one of the accepted bids, or it may be the highest one of the rejected bids, i.e. in the latter case is the highest price p for which $d(p) = n+1$. In the former case, the uniform price auction is related to the ascending price open auction.

Also in the discriminatory auction, the bidders communicate entire demand functions to the auctioneer. The auctioneer calculates the market-clearing price just as before, but now each bidder pays his bid for each unit that he is awarded. For example, if bidder i indicated that he wanted 5 units and that he was willing to pay $p_1, p_2, \dots, p_5$ for these respective units with $p_1 > p_2 > p_3 > p_4 > p_5$ and the market clearing price p^* satisfying $p_3 > p^* > p_4$, then bidder i will be awarded 3 units and he will be requested to pay $p_1 + p_2 + p_3$ in total. Obviously, this discriminatory auction is closely related to the descending price auction. However, in contrast to the single object case, there is now one important difference. It is that, in the descending auction, all bidders except the first one to bid can observe some bids by previous bidders. This additional information may be useful to them.

In his seminal 1961-article, Vickrey noted that, in the case where bidders are interested in buying multiple units, both the uniform and the discriminatory auction have important drawbacks and he proposed an auction form that does not suffer from these drawbacks. In a multi-unit 'Vickrey auction' also the highest bids are accepted, but the pricing rule is more complicated: bidders have to pay for the k -th unit which they gain the value of the k -th highest losing bid placed by the other bidders. This pricing rule is a direct generalization of the one-unit Vickrey-rule and it has a clear economic interpretation. In the one-unit case, the winner of the auction pays the value that the strongest competitor expresses for the item. To phrase this slightly differently, the winner pays the externality that he exerts on the competing bidders, that is, the value that they could have generated had he not been present in the auction. In the multi-unit case, the units are allocated to those bidders that express the highest values,

and each winner pays the value the other bidders could have generated had he not been present.

An example may illustrate this. Suppose six identical units are for sale, and there are three bidders, who each are interested in at most four units. The bidders' marginal values are given in Table 1.

	1 st	2 nd	3 rd	4 th
1	50*	47*	40*	32
2	42*	28	20	12
3	45*	35*	24	14

Table 1: An example to illustrate the Vickrey auction.

(The table should be read as follows: bidder 1 expresses a value (bid) of 50 for the first unit, 47 for the second unit that he gets, etc.). The Vickrey auction allocates three units to bidder 1, one to bidder 2 and two to bidder 3, as indicated by the entries with * in the table. In this way the highest possible total value is realised. How much should bidder 1 pay for his units? If he were not there, we could allocate 3 units more to the players 2 and 3. Of these we would give 2 units to player 2 (values 28 and 20) and 1 unit to player 3 (value 24). Consequently, player 1 should pay 28, 24 and 20 for his units, a total of 72. Similarly, player 2 should pay the externality that he exerts on the bidders 1 and 3, i.e. he should pay 32. Finally, player 3 receives two units and he should pay 32 for the second and 28 for the first, or a total of 60.

The reader may now wonder whether this Vickrey auction has an equivalent open variant. The answer is affirmative, as has recently been shown in Ausubel (2003). In Ausubel's auction, as bidding progresses, bidders 'clinch' units sequentially. The price to be paid for each unit is the price at which the auction stood at the time the unit was clinched. More formally, the price is gradually increased from 0. At each price p , each player expresses his demand $d_i(p)$ and we compute $d(p)$ just as before. In addition, for each price, we calculate the total demand of the opponents

$$d_{-i}(p) = \sum_{j \neq i} d_j(p) \quad (2)$$

as well as the supply that is available to satisfy the demand of player i after his competitors have satisfied all their demand:

$$s_i(p) = s(p) - d_{-i}(p). \quad (3)$$

As we increase p , total demand $d(p)$ will fall and at a certain p we will have

$$d_{-i}(p) < n \quad (4)$$

where n is the total number of units that is available. Let (p_1, i) be the first combination where this happens. At this price, the competitors of i demand one unit less than is available, hence, i has ‘clinched’ one unit, and the Ausubel-auction indeed allocates one unit to bidder i at this price p_1 . We thereby reduce supply by one unit (hence $s(p) = n-1$ for $p > p_1$), we also reduce the demand of player 1 by one unit and we continue the process. We repeat this process, always allocating one unit to a player k as soon as the residual supply that is available for this player $s_k(p)$ is strictly positive, until total residual supply becomes zero.

We can illustrate the Ausubel-auction by means of the values given in Table 1. If one increases p , one sees that residual demand remains at least 7 as long as $p < 20$. When $p = 20$, the total demand of the bidders 2 and 3 drops to 5 and bidder 1 can be allocated his first unit at this price. We now cross out 50 from the first row in the table and reduce the supply to 5. Next, at $p = 24$, bidder 3 drops a unit and we have $s_1(p) = 1$ so that bidder 1 can be awarded a second unit at price 24. And so on.

4.3 Multi-object auctions

We now allow for the possibility that the objects on offer are non-identical. For example, spectrum licenses sold by auction may differ in size, or in their location in the electromagnetic spectrum. These objects may have different values, hence, they will fetch different prices.

When heterogeneous objects are sold, again sequential or simultaneous sales are a possibility. When the choice is for a sequential auction, an important decision is the order in which the objects are sold: should the object with the highest expected price be sold first or last? Or is it preferable to adopt a random order? The sequencing may also be determined endogenously, i.e. the buyers may determine which object is sold first. For example, the seller can initially auction the right to choose first from the set of all objects; the highest bidder wins and chooses an object from the set. The bidders are then informed which objects are still left, and the process repeats itself.

When the FCC planned to sell multiple, non-identical spectrum licenses in the beginning of the 1990s, the auction theorists McAfee, Milgrom and Wilson devised the 'simultaneous ascending auction' by means of which the licences could be sold simultaneously, see Milgrom (2000). In this auction, all objects are sold simultaneously using an English auction procedure in which prices on each object are increased until there is no more bidding for any of the objects. At that point, the auction ends and the bidders that have made the highest bids receive the objects. As always, variants are possible, prices can be raised continuously or in discrete steps, for example, and bidders may receive full or incomplete information about which bidders are standing high at a certain point in time. We now describe one variant in more detail.

Label the available objects as $A_1, A_2, \dots, A_n$ and let there be m bidders, $i = 1, \dots, m$. The auction will proceed in a number of rounds and, in each round, it will be in a certain state. The state of the auction includes a description of (i) who has made the highest bid on each item up to that round, (ii) the value of that bid and (iii) the minimum that has to be bid on each object in the next round in order for the bid to be valid. Hence, the state of the auction at time t includes a table of the following type:

A_1	A_2	$\dots$	A_n
B_1^t	B_2^t	$\dots$	B_n^t
b_1^t	b_2^t	$\dots$	b_n^t
m_1^{t+1}	m_2^{t+1}	$\dots$	m_n^{t+1}

Table 2: Description of the state in the simultaneous ascending auction.

The columns of this table correspond to the various lots; B_j^t denotes the bidder that is standing high on lot j at the end of round t and b_j^t is the corresponding highest bid; m_j^{t+1} is the minimum bid that has to be made in round $t+1$. The auction starts in round 1 with the minimum bids m_j^1 having been chosen by the auctioneer. In each new round, the auctioneer sets new minimum prices, which typically are a certain percentage increment, say 5% or 10%, above the previous highest bids.

In addition to information on the lots, bidders also have information about the number of ‘bidding rights’, R_i^t , that each bidder i still has in round t . The bidding rights provide an upper boundary for the number of objects for which bidder i may seek to become the leading bidder in round t . Thus, if bidder i has R_i^t bidding rights in round t and this bidder is currently having the highest bids on k lots, then, in round $t+1$, this bidder is allowed to bid on at most

$$\max(0, R_i^t - k) \quad (5)$$

lots on which he is not standing high. The auction rules will determine how the bidding rights evolve, hence, in addition to Table 2 above, in each round also the table with remaining bidding rights will be available to players:

1	2	$\dots$	m
R_1^t	R_2^t	$\dots$	R_m^t

Table 3: Player’s bidding rights in round t in the SAA.

The rules may, for example, reflect concerns about competition on the aftermarket, so that bidders are not allowed to acquire more than certain maximum number of objects. On the other hand, in order to speed up the auction, if a bidder would like to receive k objects, then we would like to force him to bid on k units, or at least, we would not want him to bid for too long a time on a substantially smaller number of objects. The rules may then say that a bidder loses bidding rights if he doesn't bid for sufficiently many objects.

Let us give one example. Suppose that we want bidders to bid seriously from the start and that each bidder could possibly acquire all n objects. In that case we will have $R_i^1 = n$ for each player i . Secondly, the number of bids that player i will make in this round will determine his number of bidding rights in round 2: if bidder i bids on only l lots, then $R_i^2 = l$. Subsequently, if in round t bidder i is standing high on l_1 lots and he bids on l_2 lots on which he currently is not standing high, then in round $t+1$, we will have $R_i^{t+1} = l_1 + l_2$. Note that, as a consequence, $R_i^{t+1} \leq R_i^t$ for all i and t .

In each round, bidders, having access to the above two tables, will simultaneously decide on which lots to bid and how much to bid. Of course, bidders will have to take into account the restrictions on the minimum bids and the bidding rights. As a result of the bidding, the auctioneer will adjust the 'bid table' and the 'activity table' and provide the updated information to the bidders. The process will continue until a round t^* is reached in which no more bids are made. The bidders that are standing high at t^* receive the lots and pay the price they have bid, hence, lot j is sold to bidder $B_j^{t^*}$ for the price $b_j^{t^*}$. Note that all auctions close simultaneously; as long as there is bidding on at least one lot, it is (theoretically) possible that in some future round there might still be bidding on other lots. Also note that the simultaneous auction allows bidders a lot of flexibility: a bidder who is bidding only on lot j at first, might switch to a different lot j' if he has been overbid on j , and if he finds that j is getting too expensive. Because of this flexibility, one may expect that, in this auction, similar objects will be sold at similar prices. This property is not guaranteed when the objects

are sold in a sequential auction, and this is one of the reasons why a simultaneous format is preferred to a sequential one.

Finally, remark that, in this simultaneous ascending auction, bidders bid on individual lots; there is no possibility to directly bid on packages. As we will see in the next section, when different objects are complements, i.e. when the value of a pair of objects together is larger than the sum of the individual values, allowing such package bidding might improve the efficiency properties of the auction. In that section, we will also briefly discuss how package bids can be included and whether allowing for package bidding has drawbacks as well.

5 Bidding behaviour

To find out which auction format is optimal for the seller, one first has to ask how bidders will bid under different auction formats. In this section, we will describe and explain some aspects of bidding behaviour, and we will examine their implications for the choice of auction format. We will not provide a full overview of the results that are available, but limit ourselves to a couple of salient features with high practical relevance. As in the previous section, we move from the simplest to the more complicated situations.

5.1 Single object, own value is known

Let us write v_i for the value that bidder i assigns to the object that is for sale. Consequently, if player i wins the object for a price p , then his net gain is $v_i - p$; if i does not win the object, he does not have to pay and his utility is normalized to 0.

In the English auction, as long as the price is below one's own value, it is optimal to stay in the auction: if one quits one is sure to lose, while one might make a positive profit if one stays in. On the other hand, if the price is above the personal value, it is optimal to drop out, since winning would confer a loss. We can conclude that rational

bidders will remain in the auction until their value is reached and that the bidder with the highest value will win the auction: the auction outcome is efficient.

A similar conclusion is reached in the Vickrey auction: bidders should submit bids that are equal to their true valuation of the object (Vickrey, 1961, 1962). The reason is that under the second price rule the bid only determines *if* the bidder wins the object, but not *how much* he has to pay when he wins. A bid that is exactly equal to the true value ensures that a bidder wins whenever the price determined by the auction is below the bidder's value, and that he loses otherwise. Formally, for each bidder it is a (weakly) dominant strategy to bid truthfully: if my value is v_i , then, for any possible combination of bids of my opponents, bidding $b_i = v_i$ yields at least as much profit as any alternative bid, and sometimes the truthful bid yields strictly more.

Note that the above conclusions do not depend on the risk attitudes of the players, nor on the information that they have about their competitors' values. The simplicity of the optimal bidding strategy in the English and in the Vickrey auction can be regarded as one important advantage of these formats. However, it turns out that student subjects in experiments often do not discover the optimal bidding strategy in the Vickrey auction, even if they are given the opportunity to gather experience and learn, see Kagel (1995). Thus, it seems that, perhaps, not too much weight should be attached to the strategic simplicity of the Vickrey auction.

The situation is fundamentally different in the Dutch and first-price auctions. Under such a format, the only way for a bidder to achieve a positive surplus is for him to bid less than his true value. The issue now is by how much bidders will shade their bids, and this is a difficult problem: the longer a bidder waits, the more profit he makes if he wins, but the larger the risk that he will lose the auction. Hence, a bidder is facing a risk-return trade-off and his decision will depend on his beliefs about the competitors' values and his risk attitude. The more risk averse he is, or the more intense he expects the competition to be, the higher he will bid.

Let us assume that bidders are risk neutral, so that they only care about expected gains, an assumption that will be maintained throughout most of this paper. Suppose also for a moment that each bidder not only knows his own value, but also the values

of all competitors. In that case, the bidder with the highest value knows that he can safely wait until the clock reaches the second highest value: no competitor will bid at such a price since he would make a loss when winning at that price. Consequently, in this case, the bidder with the highest value will win and he will pay (approximately) the second highest value, just as in the English auction.

One of the results derived in Vickrey (1961) was that this equivalence of auction forms generalizes to certain settings in which bidders are uncertain about their opponents' values. Consider the so called 'symmetric independent private values' (SIPV)-model, in which bidders are risk-neutral, and consider their values as independent draws from the same distribution. If the seller does not impose a minimum bid, then, in an equilibrium each bidder will bid the value that he expects his toughest competitor to have, conditional on his own value being the highest:

$$B_i(v_i) = E(\max_{j \neq i} v_j \mid \max_{j \neq i} v_j \leq v_i) \quad (6)$$

As a consequence, in this benchmark case, the bidder with the highest value will win the object, hence, the auction outcome is efficient. Furthermore, the above equation shows that bidders will shade their bids exactly so that on average the payment will be equal to the second highest value and, therefore, the expected price is equal to the expectation of the price paid in the equilibrium of the Vickrey auction. It, therefore, also follows that a risk-neutral seller will be fully indifferent between any of the four auction forms (without minimum bids) that have been discussed: they all yield an efficient allocation and the same expected revenue.

Let us briefly illustrate how an equilibrium as in (6) can be derived. Imagine that there are two bidders, that each bidder i knows his own value v_i , but that he considers his competitor's value v_j to be an (independent) draw from the uniform distribution on $[0,1]$ and that the first price auction is used. Since the situation is symmetric, a strategy $B(\cdot)$ (a map that translates values into bids) that is good for one player should also be good for the opponent. We are looking for a bidding strategy $B(\cdot)$ such that $\langle B(\cdot), B(\cdot) \rangle$ is a symmetric Nash equilibrium, i.e. given that my opponent bids according to $B(\cdot)$, it is in my best interest to bid according to $B(\cdot)$ as well. Bidders with

higher values are more eager to win the object, hence, they will be willing to bid more, and, consequently, we will assume that $B(\cdot)$ is an increasing function. Assuming that player 2 bids according to $B(\cdot)$, let us check under which conditions player 1 finds it optimal to bid $B(x)$ for any possible value x that he might have. If player 1 would bid $B(y)$ instead, then, if his competitor bids according to $B(\cdot)$, his payoff would be

$$u(y | x) = \begin{cases} x - B(y) & \text{if } v_2 < y \\ 0 & \text{if } v_2 > y \end{cases} \quad (7)$$

which would yield the expected payoff

$$Eu(y|x) = [x - B(y)]y. \quad (8)$$

Here we have used, first of all, that $B(\cdot)$ is increasing, so that the bid $B(y)$ is winning if and only if $y > v_2$ and secondly that v_2 is uniform on $[0,1]$ so that $y = \text{Prob}[v_2 \leq y]$. Player 1 wants to maximize his payoff, hence, he wants to choose y such that $Eu(y|x)$ is maximal. The first order condition is

$$\frac{\partial Eu(y | x)}{\partial y} = x - B(y) - B'(y)y = 0 \quad (9)$$

and, to have an equilibrium, this condition should be satisfied for $y = x$, or

$$B(x) + xB'(x) = x \quad (10)$$

We can conclude that the equilibrium strategy $B(\cdot)$ should be a solution to this differential equation. Fortunately, the differential equation is simple to solve, yielding

$$B(x) = x/2 + C/x \quad (11)$$

for some constant C . This integration constant is determined by the minimum bid that the seller requires in the auction. If there is no minimum bid, then a buyer will participate no matter what his value is and we will have $B(0) = 0$. In this case $B(x) =$

$x/2$, and the result confirms equation (6): assuming that player 2's valuation v_2 is less than x , v_2 is uniformly distributed between 0 and x , hence, the conditional expected value from the right hand side of (6) is just the midpoint between 0 and x , that is $x/2$.

We now generalize these observations to an SIPV model with n bidders where values are independent and identically distributed with distribution function F . Consider any symmetric equilibrium of any symmetric auction format. Given his value x , a bidder can calculate upfront what is his probability of winning the auction, $P(x)$, as well as what is the expected transfer, $T(x)$, that he will have to make to the seller. Furthermore, the buyer can calculate the corresponding quantities resulting from him pretending that his value would be y . If a bidder would play as if his value was y , his expected payoff would be

$$U(y|x) = xP(y) - T(y) \quad (12)$$

In equilibrium, pretending to have a different value does not pay, because otherwise a bidder with value x would prefer the bid of a bidder with value y to his own bid, and we wouldn't have an equilibrium. Hence, we must have

$$\frac{\partial U(y|x)}{\partial y} = 0 \quad \text{for } y = x \quad (13)$$

If we write $U(x) = U(x|x)$ for the equilibrium expected utility for a bidder with value x , we therefore have $U'(x) = P(x)$, hence

$$U(x) = U(0) + \int_0^x P(z) dz \quad (14)$$

where we have assumed, without loss of generality, that 0 is the lowest possible value of x . From this it follows that any two auction mechanisms that have the same $P(\cdot)$ -function and that both satisfy $U(0) = 0$ have the same expected utility for the buyers. Moreover, we have that the seller's expected revenue is given by

$$R = n \int T(x) dF(x) \quad (15)$$

and since $T(x) = xP(x) - U(x)$, it follows that also the seller must be indifferent between any two auctions that have the same $P(\cdot)$ -function and that satisfy $U(0) = 0$. In summary, the seller, and all the buyers, are indifferent between auction formats which imply the same rule for allocating the object (the $P(\cdot)$ function) and which imply the same utility for a bidder with the lowest conceivable type. This result is known as the *Revenue Equivalence Theorem*.

Without reserve price, the four standard auction formats defined above imply that, in equilibrium, the object is allocated to the bidder with the highest value, hence, they have the same $P(\cdot)$ function, and that the bidder with the lowest value has zero expected utility, i.e. $U(0)=0$. Therefore, the *Revenue Equivalence Theorem* implies that all players are indifferent between these auction formats.

Let us now ask the question: Which auction format should the seller choose? The *Revenue Equivalence Theorem* implies that this boils down to the question which function $P(\cdot)$ to choose, and which value for $U(0)$. If the seller is only interested in efficiency of the allocation rule, then the four auction formats discussed above, with zero reserve price, are obviously optimal. For the case that the seller wishes to maximize expected returns, Myerson (1981) has solved the problem. He has shown that the seller will optimally set $U(0)=0$, and will allocate the object to the bidder with the highest value, except if this highest value is below some reserve price m^* , in which case the seller doesn't sell at all. The optimal value of m^* turns out to be independent of the number of bidders. It is equal to the 'take-it-or-leave-it' price that the seller would ask when faced with one bidder, with a value drawn from the distribution $F(\cdot)$, hence, if the seller's value of the object is zero, m^* is found by solving

$$\max_m m(1-F(m)) \quad (16)$$

Which auction rules implement this format? Analyzing equilibria of auctions with reserve prices along the lines indicated at the beginning of this subsection, one finds that any of the standard auction formats, with optimal reserve price m^* implies the

desired allocation rule, and yields zero expected utility for the bidder with the lowest conceivable type. Therefore, any such auction is optimal. The seller is indifferent between all four standard auction formats with this reserve price.

The results that we have described in this subsection are famous, but they do not directly apply to license auctions because the circumstances in which these auctions are conducted differ from those assumed in the theorems. The next subsections discuss several ways in which spectrum auctions deviate from the assumptions underlying the *Revenue Equivalence Theorem* and Myerson's optimal auction analysis, and why, given these deviations, the choice of auction format matters.

5.2 Asymmetries

In the previous subsection we assumed that all bidders regarded their valuations as independent draws from the same distribution. This means that bidders regard all competitors ex ante as equally strong. In many practical situations, however, some bidder might be known in advance to be stronger than another. We can describe such a situation formally by assuming that some bidders' valuations are drawn from a distribution which typically yields higher values.

As a simple example, consider the extreme case in which there are only two bidders, and bidder 1's value is drawn from a uniform distribution between 0 and 1, while bidder 2's value is drawn from a uniform distribution between 2 and 3. Thus, it is known in advance that bidder 2 can make much better use of the object than bidder 1. What will happen if one of the four standard formats is used to auction the object? The analysis of the English auction, and the second price auction, won't change. As we already noted above, that analysis is independent of the symmetry assumption. Thus, bidders will bid their true values, bidder 2 will always win, and he will pay bidder 1's value. In the first price and Dutch auction, by contrast, one equilibrium will be that bidder 1 bids his own value, perhaps recognizing that he has no chance of winning. Bidder 2 will then find it optimal to bid the maximum value of bidder 1, that is to make a bid equal to 1, and thus not to risk any probability of not winning the auction. This holds for any value which bidder 2 might have. If this equilibrium is

played, then a first price, or Dutch, auction guarantees the seller a revenue of 1, whereas a second price, or English, auction will only give a revenue equal to the true value of bidder 1. *Revenue Equivalence* between the standard auction formats no longer holds, and the seller has a preference for a first price auction. Vickrey (1961) already showed that in general no unambiguous statement about the revenue ranking of the standard auction formats is possible if bidders are asymmetric. (Also see Maskin and Riley (2000) for a recent investigation.)

5.3 The winner's curse and the linkage principle

Bidding becomes more difficult when a player does not know his own value, which will be the case in many real life auctions. In particular, bidders will typically have different pieces of information about the true value of licenses, and each bidder would revise his own valuation of the licenses if he knew not only his own information, but also the information of the other bidders. While the information of competitors is not directly available, a bidder might be able to infer it from their behaviour. For example, if bidder 1 sees that bidder 2 bids very aggressively on one particular license, he might infer that this license is more valuable than he thought before. Alternatively, if a bidder sees other bidders drop out early, then he might revise his own valuation downwards. In open auctions, one can, hence, learn from the bidding behavior of other players: their bids may reveal some of their information and may allow a bidder to make a better estimate of his own value. The *Revenue Equivalence Theorem* abstracts from such informational issues.

If informational considerations of this sort play a role, then successful bidders run the risk of suffering from the *winner's curse*. This refers to the fact that the winner of an auction is the bidder who has the highest estimate of the value of the license, and that this estimate and the corresponding bid may be overly optimistic. Other bidders apparently have had information that gave reason for more caution, and had the winning bidder had access to other bidders' information at the time of bidding, he would probably have revised his own valuation of licenses downwards. A winner who does not think these issues through in advance will suffer from the winner's curse, i.e. he will pay more than the licenses are worth.

To illustrate the possibility of the winner's curse, let us consider the following question: in the case in which one is not sure about the value and the Vickrey (2nd price) auction is used, should one bid the expected value of the object? Let us assume that the situation is symmetric and that the value of the object is the same for each player. As a specific example, think of bidders bidding for the right to drill for oil in a certain location; the amount of oil that can be extracted is (to a first approximation) independent of the winner of the auction, hence, the right has the same value for all bidders. Indeed, the study of the winner's curse originates from analysing these situations (see Capen, Clapp and Campbell (1971)). If bidding the expected value would be the optimal strategy for one bidder, then it would also be followed by the other players and in that case the winner of the auction would be the one who has estimated the expected value being highest. Being told that he has won the object, this bidder is thus, in fact, told that all other bidders have estimated the value to be lower than he has, which is bad news. Furthermore, the winner has to pay the estimate of the second most optimistic bidder and this may be above the actual value as well. In short, if one bids the expected value, one risks having to pay more for the object than it is actually worth, and thus falling prey to the winner's curse.

A rational strategic bidder in a second price auction will anticipate that a winner's curse might arise, and will adjust his bid downwards, hence, he will bid conservatively. The result of such downward adjustment of bids, will, of course, be lower average revenue for the auctioneer. If a bidder wins, he is told that he has the highest estimate; hence, a conservative way of bidding would be to bid the expected value of the object, conditional on all opponents estimating this value to be lower. It turns out that this is too conservative; one may allow for the possibility that at least one of the opponents is as optimistic as one is oneself. If we write X_j for the stochastic signal received by player j , it can be shown that, in the Vickrey auction, player i should bid

$$B_i(x_i) = E(v_i \mid x_i = \max_{j \neq i} X_j) \quad (17)$$

To illustrate our discussion, suppose there are three bidders who bid in an ascending auction for an object that is worth

$$v=(x_1+ x_2+ x_3)/3 \tag{18}$$

to each of them. Here x_i is a signal that player i has received about the value; player i , however, does not know the signal x_j of the competitor, which he considers to be drawn from the uniform distribution on $[0,1]$. A naïve bidder would bid as if the value of his own signal: $b_i(x_i)= x_i$. Such a bidder would suffer from the winner's curse. The naïve correction for the winner's curse, which conditions on the event that both other signals are below one's own, results in a bid of: $b_i(x_i)=(1/3) x_i + (2/3) (x_i/2)$. This is because the conditional expected value of a signal, conditional on being less than x_i , is $x_i/2$. However, the equilibrium bid is larger than this, and conditions on the event that one of the other bids is the same as x_i . Thus, it is: $b_i(x_i)=(2/3) x_i + (1/3) (x_i/2)$.

Efficiency of auctions where bidders don't know their own value becomes an issue when, unlike in our example, valuations have a common as well as an idiosyncratic component. Under symmetry conditions, the bidder with the highest signal will have the highest value and will make the highest bid, so that the Vickrey auction will also result in an efficient allocation in this case. Under these same conditions, the Dutch (first price) auction also yields an efficient outcome, and the equilibrium bidding strategy is given by a similar, but slightly more complicated, formula. If individual signals are, unlike in our example, affiliated, however, most frequently the seller is not indifferent between these auctions: the Dutch auction results in lower (or at least not higher) expected revenue than the Vickrey auction, which in turn yields no higher expected revenue than the English auction. In other words, the loss that the auctioneer suffers because bidders adjust their bids in anticipation of the winner's curse is typically lower in second price auctions than in first price auctions, and it is lower in the English auction than in the Vickrey auction.

A popular intuition for these results, which is, however, only partially correct, is that, the more information a bidder has, the smaller the chance of falling prey to the winner's curse and the more aggressive a bidder can bid. The result is an application of a more general idea (originally derived in Milgrom and Weber (1982)), which in auction theory is known as the 'linkage principle': it works, at least on average, to the auctioneer's advantage if he can link the price paid by winning bidders to signals that

are correlated with the signals of winning bidders. An excellent exposition of the linkage principle is in Chapter 7 of Krishna (2003). His Proposition 7.1 formulates the linkage principle for a general setting in which bidders' valuations are not necessarily common. The linkage principle is often related to the winner's curse, but the connection is only loose. In particular, there are common value settings where individual signals are independent, and therefore the linkage principle does not apply, even though the winner's curse is clearly relevant.

In a first price auction, conditional on being told that he has won, a bidder does not have more information than the information he used when he made his bid, in particular, the auction price just depends on his own bid. In contrast, in the Vickrey auction, the payment made by the winner is linked to the information of one of the losers of the auction. In a common value environment, if the signal of the loser is correlated to the signal of the winner, the linkage principle implies that the second-price auction offers the auctioneer higher average revenue than the first-price auction.

As a further application of the 'linkage principle', we obtain the classic argument in favour of open auction formats. These formats allow bidders to observe other bidders' decisions, such as these bidders' decisions to exit from the auction. If other bidders' signals contain information about the value of the object, and a bidder's own signal does not include this information, then bidders will have an incentive to learn from others' decisions, and to revise their own plans continuously. This will tighten the link between other bidders' signals, and the price paid by the winning bidder. If signals are correlated, then the linkage principle implies that an open format leads to a higher expected revenue for the seller than a closed format.

To conclude, under natural assumptions, the four basic auction forms still all generate an efficient outcome, but they typically do not yield the same expected revenue for the seller. The English auction yields at least as much expected revenue as the Vickrey auction, and this Vickrey auction yields at least as much as the Dutch auction. If the seller can avoid collusion among the bidders (see below), he is thus advised to organize an open ascending procedure as this provides the tightest link between the price paid by the winning bidder and the signals observed by other bidders.

5.4 Bidding in multiple unit auctions

We emphasize from the outset that the literature on auctions with multi-unit demand is much less developed than the literature on auctions with single unit demand. Multi-unit demand is, in fact, one of the areas on which current research in auction theory concentrates. At this stage, though, the question on which we focused in previous subsections, i.e. ‘What is/are the equilibrium/a of a given auction format under certain assumptions about values and information?’, has not been answered at any level of generality for multiple unit auctions. We can thus only point out some intuitions, but we can’t present any general results.

Since a sequential auction is easy to organize, a seller will find it tempting to sell multiple units sequentially. For a bidder, a sequential auction presents considerable strategic complexity, however. For example, there might be a reason to hide true values in early auctions so as to induce other bidders to bid lower in later auctions. In addition, supply/demand conditions change during the auction. On the one hand there is an incentive to bid *more* aggressively in later auctions because fewer items remain for sale. On the other hand, some bidders, possibly those with the highest valuations, have already won a license and have left the auction. Thus, there is an incentive to bid *less* aggressively. How these intuitions interact, and what is optimal bidding behaviour in sequential auctions, has only been resolved in some special cases (see Chapter 15 of Krishna (2002)). For sequential auctions in which each bidder is interested in just one item, it turns out that the effects exactly cancel each other out, independent of whether the auction format is first price or second price, sealed bid: the expected price for which item $l+1$ will be sold is exactly equal to the price for which item l is sold. However, it is also well-known that in practice bidders’ behavior deviates from the predictions derived for these special cases: in many real life auctions, prices display a downward drift, see Ashenfelter (1989).

The situation gets even more complicated if bidders are interested in multiple units, and especially so when there are complementarities and players’ values are superadditive ($v_1 + v_2 = v_3$). For example, if bidder 1 has already won an item, he may need a second one in order to generate value also from the first. He may, therefore,

need to bid very aggressively on a second unit, which has two effects. On the one hand, his competitors know that it is unlikely that they can win the second unit, which may discourage them from bidding. On the other hand, they know that bidding is relatively riskless for them, hence, they may bid to drive up the price for bidder 1 and to weaken him in that way. To put it differently, bidder 1 is liable to a ‘hold-up’ problem in this case. Of course, if a player foresees this, and considers the risks to be too large, he may decide not to participate in the first place. As a consequence, it does not seem a good idea to use a sequential auction in situations like these, and, overall, the consensus view is that it is better not to confront bidders with the complications of a sequential format.

Moving to simultaneous auctions, we report three interesting intuitions emerging from the literature. The first two relate specifically to uniform price auctions, to which we will mainly restrict ourselves.

We note, first of all, that the uniform $(n+1)$ -auction (in which all winning bidders pay the highest losing bid) does not share the nice properties of the second-price auction from the 1-unit case, at least not in the situation where bidders are interested in more than one unit. The reason is simple; if a bidder can demand more than one unit, he can influence the market price and he will profit from a lower market price on all the units that he gets. Thus, there is the possibility of players engaging in *demand reduction*: bidders understate their true value of different units (Ausubel and Cramton (1996) and Engelbrecht-Wiggans and Kahn (1998)). While demand reduction implies that a bidder will not win some units that he would have liked to win, it is advantageous because it reduces the price which the bidder has to pay for all those units which he will win. When demand goes down, the equilibrium price at which the market clears goes down, too. Demand reduction is a source of inefficiency in uniform price auctions, and it reduces sellers’ expected revenues.

To give a very simple example, suppose that there are two units and two bidders, with bidder 1 wanting to have two units and player 2 interested in just one unit. Bidder 1 will realize that his demand for the second unit might set the price: if he pretends not to value this second unit (bids 0 for it), then he is guaranteed to get one unit for free, and this will frequently be better than to compete head on with bidder 2 in an attempt

to acquire two units. As a second example, assume that there are two bidders that each want two units and that supply is also equal to two units. Furthermore, suppose that each bidder knows his values and that it is known that all (marginal) values are in the order of 50. If bidders bid truthfully and compete head on, price will rise to approximately 50 and each bidder will have utility close to zero. If the ascending English auction is used, it is, however, much better to reduce demand to one unit immediately. If one bidder does this, the other notices immediately that he can stop the auction by also reducing his demand. In this case, each bidder will get one item and the price will be close to zero. In equilibrium, therefore, bidding will stop immediately at price 0. Grimm, Riedel and Wolfstetter (2002) illustrate that a demand reduction strategy was successfully followed in the German DCS-1800 auction. In a simplified model of that auction, these authors also show that the unique 'sensible' equilibrium will result in demand reduction.

The second intuition that emerges from the literature on sealed bid auctions with multi-unit demand is that uniform price auctions offer an opportunity for a particular form of implicit collusion. The strategies adopted by bidders for this form of collusion involve overstating the willingness to bid for the first few units, and to understate the willingness to bid for later units. Values are overstated for those units that the bidder is sure to win. What the bidder bids for these units is certain not to influence the market price directly. Understatement occurs for those units that the bidder regards as 'marginal', i.e. he might win them, or he might not win them. Understatement of values for these units goes beyond the demand reduction effect described in the previous paragraphs, see Binmore and Swierzbinski (2000).

The logic behind these strategies, first pointed out in Wilson (1979), is that each bidder's demand function is very steep around the equilibrium price. The effect of this is that each bidder also faces a very steep residual demand. That will mean that an increase in demand will lead to a very sharp rise in the market-clearing price. This in turn deters all bidders from raising their demand. Note that this form of implicit collusion assumes that bidders can predict the market-clearing price relatively well. If there is large uncertainty regarding this price, then this form of collusion cannot be implemented.

Discriminatory price auctions create other, more subtle inefficiencies, but we will not discuss these here. One advantage of the more complicated pricing rule adopted in the Vickrey auction is that similar problems do not arise there. Vickrey showed that the auction that he proposed inherits the nice properties of the Vickrey-auction in the one-unit case: it is a dominant strategy for each player to truthfully report his values to the auctioneer. Consequently, the Vickrey auction is easy to play, honesty is the best policy, and bids coincide with actual values. As a result, therefore, this Vickrey auction is a robust mechanism that produces an efficient outcome: it is impossible to generate higher total surplus. Note that this does not imply that this auction format also raises the highest possible revenue for the seller; on the contrary, the revenue in the Vickrey auction can be quite low. Of course, since demand reduction also harms revenue, it is not clear that a uniform auction will result in higher revenue.

The third intuition emerging from the literature concerns the case in which bidders don't know their own value, and each bidder believes that other bidders hold private information that is potentially relevant to their own valuation. We explained above that in the case of single unit auctions the winner's curse arises. In the case of multiple unit auctions, there is also a loser's curse. This refers to the fact that losing at the margin now implies good news about the value of the good: the winners of the non-marginal units must have had better information (Pesendorfer and Swinkels (1997)). This leads bidders to bid more aggressively.

The linkage principle which we explained above for single unit auctions fails in the case of multi-unit auctions (Perry and Reny (1999)). There is no clear ranking of different auction formats in terms of their expected revenue. An intuitive explanation of this may be seen in the presence of the loser's curse.

We conclude this subsection by pointing out that additional problems arise when the bidders' values are super-additive, that is, when marginal values are increasing over a certain range. To illustrate, consider the situation represented in Table 4.

	0	1	2	3
1	0	0	60	99
2	0	0	80	90

Table 4: Increasing marginal values (for player 1).

There are two bidders and three units, each bidder requires at least two units to generate value, with bidder 1 valuing two units at 60 and three units at 99, etc. The efficient allocation is that bidder 1 obtains all 3 units. Assume that the ascending English auction is used to allocate these units. Notice that the demand of players 1 is given by

$$d_1(p) = \begin{cases} 3 & \text{if } p < 33 \\ 0 & \text{if } p > 33 \end{cases} \quad (19)$$

whereas the demand of player 2 is given by

$$d_2(p) = \begin{cases} 3 & \text{if } p < 10 \\ 2 & \text{if } 10 < p < 40 \\ 0 & \text{if } p > 40 \end{cases} \quad (20)$$

These demands are such that there is no market-clearing price, i.e. $d_1(p) + d_2(p) \neq 3$ for all p . If the English auction is used to sell the units, then the auction will stop at $p = 34$, where player 2 will buy two units. This English auction does not produce an efficient outcome.

In this example, the problem is caused by the non-convexity; to give player 1 a chance at winning and at reaching an efficient outcome, player 1 should be allowed to make package bids. If bidder 1 could make package bids, he would bid 91 in total for 3 units (without having to specify a specific price per unit) and he would win the auction, thus inducing an efficient outcome. We shall discuss package bidding again in the next subsection.

5.5 Multi-object auctions

It must be emphasized that the literature on this subject is incomplete. The best-understood auction format is the simultaneous ascending auction format (see Milgrom (2000)). In the simple case in which each bidder can buy at most one license, and in which informational considerations don't play a crucial role, this simultaneous ascending auction is reasonably well understood. A simple extension of truthful bidding, called straightforward bidding, can be shown to be rational. The key idea here is that each bidder bids for that license which currently offers the highest surplus, i.e. the highest difference between value and price. A proof that straightforward bidding is an equilibrium strategy can be constructed using older results of Leonard (1983) and Demange et. al. (1986).

We now turn to multi-object auctions in which bidders may demand more than one unit. In this case, several of the problems that we have already encountered in the previous subsection appear as well. In this subsection, we discuss two of these in greater detail: the exposure problem, and the free rider problem.

5.5.1 The exposure problem

In the previous subsection, we already mentioned that a bidder may be liable to hold-up in a sequential auction in which he needs to acquire multiple units in order to generate value. The same happens in multi-object auctions. Imagine that two objects are for sale through English auction and that bidder 1 needs to acquire both to generate value. If he has already bought the first object, for price p_1 , then, in the second auction, this price that has been paid is a sunk cost and it will be completely irrelevant for bidding. If the bidder attaches value v to the pair of objects, then in the second auction he is willing to bid up to v and, if the competition for the second object is intense, for example, since some bidders are only interested in this second object, the total price paid, $p_1 + p_2$, may well exceed the value. Of course, a player will be aware of the risks involved and he may decide that they are so large that it is better not to participate in the auction at all. Sequential auctions, hence, may attract few bidders and may generate low revenue. It may, therefore, not be a good idea to set up

the auction as a sequential one. In Van Damme (2002b), this argument was used to criticize the decision of the government to use a sequential auction for selling licenses to operate gasoline stations, of which the aim was to bring new entrants to the market.

The exposure problem is not avoided by the simultaneous ascending auction (SAA-format). Consider the values as in Table 5, where the lots A and B are complements for bidder 1, but substitutes for player 2. Hence, player 2 is satisfied with one unit, and he is indifferent between which one he gets, while player 1 is willing to pay more for the second object if he has already bought the first one.

	A	B	AB
1	4	4	10
2	V	V	V

Table 5: The exposure problem.

If the SAA-format is used, prices on both lots will rise to 4 at which point player 1 is facing a difficult decision: he is willing to pay up to 5 for each object, provided that he obtains both of them, but he cannot be sure that he will win these two items for a total cost less than 10. If player 1 continues bidding and $V > 5$, then the price of each object will rise above 5 and player 1 is sure to make a loss. Several possibilities now can arise. If player 1 does not take the risk, he quits at 4, each bidder wins one license and, if $V < 6$, the outcome is inefficient. If player 1 continues bidding at $p = 4$, then he might win both lots and force an efficient outcome, but he cannot avoid making losses.

The literature contains only very few results about rational bidding if an exposure problem is present. See Krishna and Rosenthal (1996) for an example. For a practical example of the relevance of the exposure problem, see the discussion in Van Damme (1999) on the Dutch DCS-1800 auction.

5.5.2 Package bidding and the free rider problem

Bidding in multi-object auctions becomes more complicated when there are complementarities among licenses. Complementarities exist if a bidder values some particular license A more if he already holds another license B. In this case, bidders would like to submit two separate bids for A, one applies if the bidder wins A only, and another applies if the bidder wins A and B at the same time. In such cases, it is desirable to expand bidders' strategy sets and to allow them to make mutually exclusive bids for different packages of licenses without the bidder having to specify prices for individual items in the package. Auctions allowing for these possibilities are called 'package auctions' or 'combinatorial auctions' and their great advantage is that they avoid the exposure problem. For a long time economists have regarded such auctions with some scepticism because of the complexity involved. The recently increased popularity of combinatorial auctions derives to a significant extent from experimental research with student subjects in which such auctions have been shown to perform well. The relevant experiments have been undertaken at the University of Arizona, and are documented in a report for the Federal Communications Commission of the United States, see Cybernomics (2000).

The successful experimental work with combinatorial auctions has now also triggered additional theoretical research in this area. As is the case with many other auction formats, there are static and dynamic versions of such auctions. In fact, the Vickrey auction can easily be extended to include package bids: each bidder expresses a value for each possible package, the auctioneer determines which partition of the objects maximizes total value and he asks each bidder to pay the externality that he imposes on others, i.e. the loss in value that they incur because of him getting some of the objects. Formally, let N be the set of items and let S denote any subset. A bid of player i now specifies the amount $b_i(S)$ that i would want to pay for each set of items that he might want to have. Given such bids for all players, the auctioneer can calculate which partition of N maximizes value (as expressed by the bids) and he can allocate items accordingly. If each bidder has to pay the opportunity cost of the items that are allocated to him, i.e. the loss in value to the competitors since these objects are then not available to them, it is a weakly dominant strategy of each bidder to report the value truthfully, and the auction generates an efficient allocation. While this

efficiency property is nice, an obvious drawback is that bidders have to communicate a lot of information to the auctioneer and that the computational burden may be considerable. (If there are 10 objects, then a bid can specify up to 1024 numbers for each bidder.)

Recently, Ausubel and Milgrom (2002) have proposed a dynamic combinatorial auction which allows very flexible package bidding. They show that an extension of honest, ‘straightforward’ bidding creates efficiency in their design, and achieves efficient outcomes. We will not discuss that auction form in detail, but confine ourselves to showing that allowing package bids is not a panacea. Suppose there are three bidders and two objects, with values as in Table 6. Here, bidder 3 is only interested in the pair, while bidders 1 and 2 each are interested in only one of these items

	A	B	AB
1	x_1	0	x_1
2	0	x_2	x_2
3	0	0	10

Table 6: The “free rider” problem.

Suppose that bidder 3 has made a package bid of 10 on the combination AB. Suppose also that players 1 and 2 know that the structure of the values is as in the table (hence, player 1 knows that player 2 is only interested in B), but that the actual values are private information. In this case, the players 1 and 2 face a coordination problem: they have to bid up the prices and they will win only if $p_1 + p_2 > 10$; if they jointly outbid player 3, then the net revenue of player i is $x_i - p_i$, hence, each player wants the total bid to be higher than 10, but each wants to contribute as little as possible to the common good of winning. Consequently, there is a coordination problem (which is acerbated by the fact that the values are private) and the players may not be able to solve this problem. As a result, even though it may be that $x_1 + x_2 > 10$ and that it would be efficient for the players 1 and 2 to win the items, the outcome may be that prices are not raised sufficiently to outbid player 3.

5.6 Collusion

Up to now, in this Section, we have assumed that the bidders behave non-cooperatively. It is, however, easy to see that, in auctions, the incentives to collude are very strong. Suppose there are two bidders for one item and the values are v and V , with $v < V$. If both players participate in an English auction, the price will be v , but if players collude, identify the bidder with the highest value before the auction, and agree that the weaker bidder will not compete, then the price will drop to 0. Consequently, the gains to collusion are v . As these gains can be considerable, there are strong incentives for bid rigging. A simple method for collusion that aims at lowering the price that the winning bidders have to pay is a so-called 'bidding ring' (see Graham and Marshall (1987)). A bidding ring first establishes internally, in a preliminary auction, who should win the object. The bidding ring then only lets the winners of the preliminary internal auction submit serious bids in the official auction. While bidding rings are typically illegal, they are not always easy to detect. Therefore, the possibility of bidding rings needs to be taken seriously.

The issues that arise in this context are similar to the issues that arise in general cartel behaviour: what market structures (auction forms) are most conducive to cartels? How can the cartel identify the efficient outcome? Can it agree on division of the gains achieved from bid-rigging? How can the cartel agreement be enforced? We now discuss these questions in some greater detail.

Because bidding rings are typically illegal, the underlying agreements among bidders cannot be enforced through courts. It is therefore important whether bidders can enter into bidding ring agreements that are self-enforcing in the sense that once the agreement has been established, no party has an incentive to deviate provided that they believe that the other parties to the agreement stick to it. A general theoretical insight which is relevant in this context is that cartel enforcement is simpler in more transparent markets in which players interact repeatedly. The argument is simply that, in a more transparent market, a deviation from the collusive agreement can be detected more easily, while the repeated interaction provides the opportunity to punish those that deviate from the agreement, thus making deviating less attractive.

As an application of this insight, we can see that collusive agreements can most easily be enforced in English ascending auctions. Indeed, in the above example, suppose that the two bidders have reached an agreement and that the weaker bidder is supposed to stay out. Does this weaker player have an incentive to deviate? Well, if he does and participates in the auction, the opponent will notice immediately and counter with a higher bid; bidding will then enter a competitive phase, which the weaker bidder cannot win, as he has the lower value in the first place. Consequently, there is no incentive whatsoever to deviate and the collusive agreement is stable.

Bidding rings in second price auctions are typically also self-enforcing. Suppose bidders have formed a bidding ring and they have identified the bidder that has the 'right' to win the auction. In a second price auction, this bidder can simply enter his value as a bid, and the price will be the maximum of the seller's reservation value and the highest bid of the bidders that are not member of the ring. The ring successfully lowers what the lead bidder of the ring has to pay since typically the highest losing bid will be lower. If the cartel works efficiently, another ring member has no incentive to bid: he has a lower value than the ring member that has the right to bid, hence, he can never win, unless he overbids his value. Non-lead bidders could increase the price that lead bidders have to pay, but, of course, they have no incentive to do so.

A first conclusion, therefore, is that English and Vickrey auctions are even more susceptible to collusion than other markets. In other industrial markets, cartels operate by restricting supply, which confers a positive externality on non-cartel members: these benefit from the higher price, but, unlike cartel members, they are not hindered by production quota. In industrial markets, there, hence, is a free rider problem: one prefers to have a cartel in the market, but one also prefers not to be a cartel member. In an English or second price auction, outsiders, however, do not directly benefit from the existence of an efficient cartel: such a cartel will be represented by the most aggressive member and any non-cartel member will face exactly the same competition as when there was no cartel. Consequently, in such auctions, there is an incentive to form all-inclusive cartels.

In a first price sealed bid auction, the situation is somewhat different. In this case, the cartel bidder will shade his bid, and he will shade it more, the larger the coverage of

the cartel, and this offers cartel members the opportunity to outbid the designated bidder. As a simple example, suppose the cartel is all-inclusive and the seller sets no reserve price. In this case, the designated bidder will bid almost nothing and bidding a small amount suffices to outbid him. As a consequence, bidding rings can be expected to be more stable in second price auctions than in first price auctions.

A second conclusion, therefore, is that open auctions are more susceptible to collusion than sealed-bid auctions and that second price auctions are more vulnerable to collusion than first price auctions. It follows that, when we take into account the possibility of collusion, the seller's ranking of auctions is exactly opposite to the ranking we derived under non-cooperative behavior. While under non-cooperative behavior, the English auction is preferred to the Vickrey auction and the Vickrey auction is preferred to the Dutch auction (this on the basis of the seller's expected revenue), we now see that the English auction is more susceptible to collusion than the Vickrey auction and that this Vickrey auction in turn is more susceptible than the Dutch auction.

The above paragraphs discussed the issue of enforcing the cartel agreement, but will a cartel be able to conclude such an agreement and will it be able to identify the most efficient bidder in the cartel? An efficient cartel must allocate the right to bid to the strongest bidder in the cartel and it must adequately compensate the other cartel members for giving up their right to bid. One way of achieving these objectives is to organize a pre-auction knockout in which this right is sold. McAfee and McMillan (1992) analyze such pre-auction knockouts. They distinguish between the situation in which the cartel participates in a sequence of auctions so that the books of the cartel office have to balance only on average and the situation in which transfers among the members have to balance in each possible instance. In the former case, the additional degree of freedom makes the problem easier to solve. Recently, in 2002, a Dutch parliamentary investigation uncovered a bidding ring in the construction sector, which, allegedly, operated in just this way.

Market transparency also helps the cartel in reaching an efficient agreement: the better the information about the players' values, the easier it is to see which bidder should

win and by how much the others should be compensated. In this respect, the Dutch UMTS-auction provides interesting lessons, see Van Damme (2002a) and Janssen et al. (2001) for further discussion on this point.

6 Which auction form to adopt?

We now return to the question: Which auction form will best serve the seller's interests? As already indicated in the Introduction of this paper, the answer will depend not only on theoretical insights, but also on empirical evidence. Here, we confine ourselves to some general theoretical considerations.

In this paper, we have mainly limited ourselves to standard auction formats. In Section 5.1, we briefly discussed 'optimal' auctions, i.e. auctions which maximize the seller's expected revenue. There we have seen that, if the SIPV-assumptions apply, any of the standard auctions is optimal, provided that the reserve price is chosen appropriately. Once the SIPV-assumptions are relaxed, optimally designed auctions are often highly implausible in practice, and require detailed prior knowledge of the agents' subjective beliefs as well as strong precommitment power by the auctioneer, see McAfee, McMillan and Reny (1989). Consequently, we will here restrict ourselves to standard auctions. We discuss various design issues that arise and point out the trade-offs that exist in resolving them.

6.1 Sequential or simultaneous auctions?

This issue arises when the seller has multiple units for sale. It was extensively discussed in the 1990s when the US was working towards the design of the spectrum auctions. We have also given it ample attention in this paper. The only argument we have given for why a seller might choose to adopt a sequential auction is that such an auction appears easier to organize. We have also seen, however, that frequently this argument is not convincing: this simple solution for the seller results in a very complicated problem for the bidders. We have given several arguments for why a simultaneous auction should be the preferred choice, and we have not identified

drawbacks, at least no drawbacks that could not be remedied. The advice definitely is to sell related licenses simultaneously as much as possible.

6.2 Open ascending auctions or sealed bid auctions?

Open auctions sometimes appear to be strategically simpler than sealed bid auctions. Consider the comparison between English auctions, and the related second price, sealed bid (Vickrey) auction. We have seen that, in the SIPV-model, staying in the auction until the value is reached is the optimal strategy in the English auction, and that similarly, in the Vickrey auction, it is optimal to bid one's own value. Experimental research with student subjects as bidders has shown that bidders often do not realize the simple logic behind this result if the auction is conducted in a sealed bid format. By contrast, if the auction is conducted as an open ascending auction, bidders easily understand that they should drop out once bidding has reached their reservation value.

Why is strategic simplicity desirable? Bidders are more likely to play equilibrium strategies, and if the choice of an auction format is based on equilibrium predictions, then it is more likely to be successful if bidders recognize equilibrium. Transparent auctions are also less liable to legal challenge, and they reduce the chance that bidders place bids on which they later have to default.

A different argument, this time in favour of sealed bid auctions, can be constructed from our discussion of asymmetries in Section 5.2. If some bidders are known ex ante to be 'weak', then these bidders have a better chance of winning in first price sealed bid formats than in English ascending format. Thus, weaker bidders may have a stronger incentive to participate in the first place (Klemperer (2002)). This might raise revenue. However, this gain comes at the expense of efficiency.

As explained in Section 5.3, winner's curse effects may imply that open formats lead to higher expected revenue. But, as shown in Section 5.4, in multi-unit settings the effect is no longer clear.

An argument against ascending auctions is that these are more vulnerable against collusion than sealed and descending formats. In open auctions it is easier to detect whether members of a cartel ring deviated from the agreed bidding strategy, which makes it easier to implement some form of punishment for such deviators. From the auctioneer's point of view, this is undesirable, as it might reduce revenues.

Overall, thus, there is no unambiguous theoretical case, either in favour of sealed bid auctions, or in favour of ascending auctions.

6.3 Reserve prices

Any auction will typically involve a reserve price, i.e. only bids above a certain minimum are allowed. In a sealed bid auction, a reserve price can be implemented by simply ignoring bids that are below the reserve price. In open ascending price auctions bidding can simply start at the reserve price. In descending price formats the auction can close at the reserve price even if the number of bids is still below the number of units for sale.

A seller might be tempted to set the reserve price simply equal to the value that unsold licenses have for him. However, revenue considerations suggest that a higher reserve price should be set, as in Equation (13). This may appear at first paradoxical because it implies that the seller risks not selling all objects even though there are some bidders whose value is above the auctioneer's own value. However, the existence of the reserve price encourages more aggressive bidding, and this more than compensates for the risk of not trading when trade would be efficient. Equation (13) shows that a revenue-maximizing seller should always set a reserve price that exceeds his value.

If a reserve price is set, it is important that it is credible and that the auctioneer does not later reduce it. Sellers, of course, face a temptation to lower their reserve price, if they see that at the reserve price they cannot sell everything that they want to sell. It is important that sellers do not give in to this temptation. If bidders believe that the auctioneer might later lower the reserve price, then they will bid lower, and the crucial

advantage of the reserve price will be lost. Related to this issue is the question of whether the reserve price is secret or known. A secret reserve price (i.e. bidders know that there is a limit, but they do not know its value) offers the seller an easier opportunity to renege and to sell even if the actual price is below his reserve price. Hence, a secret reserve price is less credible than a public reserve price, and bidders will take this into account in their strategies. Because of these considerations, the literature advises a revenue-maximizing seller to publicly announce the reserve price.

A reserve price can make it more difficult for bidders to successfully collude in an auction, especially if the reserve price is secret. With a secret reserve price, the situation is less transparent, which makes it more difficult for bidders to identify the gains from collusion, hence, what side payments the strongest bidder should make to induce his competitors to stay out of the auction. If the seller adopts a public reserve price, this should be higher the more bidders he expects to participate in the cartel, see Krishna (2002, Chapter 11).

Reserve prices can be problematic, however, when entry is an issue. In general, it appears to be that entry issues should be given more weight than the effects described in the previous paragraphs. If additional bidders enter the auction, then all bidders will raise their bids, and this effect seems to be more important than the incentive to bid higher which reserve bids provide (see Bulow and Klemperer (1996)). In the symmetric independent private value model, when the number of bidders is endogenous, it is optimal not to set any reserve price at all (see McAfee and McMillan (1987), and Levin and Smith (1994)).

6.4 What should bids look like?

Our discussion of super-additive valuations and complementarities indicated that auction formats which don't give bidders the opportunity to express these by placing package bids create complicated strategic problems for bidders, and inefficiencies. Thus, it seems desirable to offer bidders an opportunity to express in their bids such key aspects of their preferences. On the other hand, excessive flexibility seems to make the strategy space too large. For example, when a multi-object auction calls for

package bids for arbitrary combinations of objects, then the strategy space may very soon become too large to be manageable for bidders.

6.5 How much information should be revealed?

In the case of single-unit auctions the linkage principle suggests that it is good for revenues to reveal at each stage of a multi-round auction each bid, and the bidder who placed it. On the other hand, the more public bids are, the easier it becomes for bidders to collude in multi-unit or multi-object auctions.

6.6 What should be the pricing rule?

It is usually a good rule that the highest bidders win an auction. It is less clear how prices should be determined. However, in general a good rule for determining prices seems to be that each winning bidder should pay for the externality that he imposes on other bidders by winning. The price should therefore be the highest valuation of the other bidders. This idea underlies not only the auction formats proposed by Vickrey, but also the Ausubel auction in the multi-unit case, and the Ausubel-Milgrom auction in the multi-object case. In this context, it is important that the seller can commit himself to not use for other purposes the information that is revealed in such Vickrey auctions. Indeed, bidders in these auctions are willing to bid truthfully only since this information will not be used for pricing purposes. Rothkopf et al (1990) have argued that, in practise, Vickrey auctions are rare since governments will not be able to enter in such commitments. In particular, bidders will fear that the information they reveal in the auction might be used against them after the auction. Furthermore, truthful revelation might also present the government with a problem; when, in New Zealand, the Vickrey auction was used for selling telecommunications licenses, the government got in political trouble because it was revealed that the winning bidder was willing to pay much more than the price he had to pay (see McMillan (1994)).

6.7 Risk aversion

Bidders are often averse to risk. This leads to different bidding behaviour from the one discussed so far in some auction formats, and therefore some of the results explained above no longer hold. Consider the comparison between the first price, sealed bid auction and the Vickrey auction in the case that bidders have symmetric private values but are risk averse. Bidding behaviour in the Vickrey auction will not change: bidding one's true value remains the best strategy, by the same argument that was explained in Section 5.1. However, equilibrium bids in the first price, sealed bid auction will increase. The reason can be explained as follows. By raising his bid, a bidder reduces the uncertainty to which he is exposed in return for an additional payment that is to be paid in case of winning the auction. The risk reduction is more valuable to a risk-averse bidder than to a risk-neutral bidder. The marginal incentives to raise one's bid are therefore higher for a risk-averse bidder, and the risk-averse bidders will bid more in equilibrium than risk-neutral bidders. Because with risk-neutral bidders the Vickrey and the first price, sealed bid auction yield identical expected revenue, it follows that with risk-averse bidders expected revenues are higher under a first-price sealed-bid format than under the Vickrey format.

Often it is also the case that the seller is risk averse. In the case of spectrum auctions, for example, it seems plausible that government agencies which sell spectrum would like to avoid as much as possible any uncertainty about the revenues which they receive. In this case, too, the first price auction is preferable. The intuition is related to the intuition explained above. The first price auction removes part of the uncertainty about the returns. Conditional on the identity of the winner, and the winner's true valuation of the object, there is no further uncertainty about returns in the first price auction, but there is such uncertainty in the second price auction.

6.8 Externalities

It has been argued that, in spectrum auctions, bidders' valuations of licenses reflect so-called 'allocative externalities', i.e. that a bidder's valuation of a license depends on which other bidders get a license. If bidders' valuation structure incorporates

locative externalities, it becomes much harder to analyse optimal bidding behavior in standard auctions. Important contributions to this literature are due to Jehiel and Moldovanu (1996, 2000a, 2001). Situations in which such externalities are present can represent formidable strategic complexity for bidders, in fact, equilibria in the standard game-theoretic sense need not exist.

Revenue maximising mechanisms in the case of positive externalities turn out to be very sophisticated. They require substantial precommitment power by the auctioneer, and seem in practice not plausible. If attention is restricted to more simple mechanisms, however, then Jehiel and Moldovanu (2000a, Section 4.1) obtain an interesting result which is of potential practical relevance: If there are strong negative externalities, and bidders are afraid that licenses fall into the hands of particular competitors, then it might be in the interest of the auctioneer to set a very low reserve price, or to pay for participation. This will intensify bidders' fears that the competitors obtain licenses, and will therefore induce them to bid more aggressively.

Maasland and Onderstal (2002a/b) consider the case of financial externalities, in which auction losers prefer the auction winner to pay more. One reason for such preference might be that losers meet winners in other markets and that, because of budget constraints, the winner might be a less fierce competitor in this other market if he has paid more in the auction. The authors show that, if such externalities are present, bidders will bid more aggressively in the first price auction than in the second price auction, hence, the seller's expected revenue is higher in the former.

6.9 Fighting Collusion

Bidding rings lower the expected revenue of the auctioneer and, from his point of view, they are, of course, undesirable. The arguments from Section 5.6 indicate that second price auctions are more in danger of being manipulated by bidding rings than first price auctions, and that English auctions are even more vulnerable to collusion. This general argument is also relevant in the case when multiple objects are for sale, but simultaneous ascending auctions are vulnerable to further types of collusion (see Brusco and Lopomo (2002) and chapter 3 in this volume by Salmon). Bidders can agree to share licenses in a particular way, and to place very low bids without

challenging each other. If any bidder deviates from this agreement, then bidders revert to a more competitive equilibrium. These bidding strategies are self-enforcing. The envisaged collusion relates to market sharing and differs from the collusion which a bidding ring practices in that all bidders who are participating in the collusion are present in the auction. By contrast, in a bidding ring, only certain members of the bidding ring enter the auction.

The vulnerability of the simultaneous ascending auction to collusion strengthens the case for a sealed bid format. However, it is unclear how a sealed bid format could address all the complex issues that this case raises.

6.10 Asymmetries between bidders

Which auction format should the auctioneer choose when he knows that bidders are asymmetric? The issue is complicated and arguments can be constructed which go both ways, in favour of a first-price format, but also in favour of a second-price format (see Krishna (2002, Chapter 4.3)). However, probably the most important consideration in the context of spectrum auctions is entry. If weaker bidders know in advance that they have a low chance of winning a license, then they won't be willing to participate in the auction. There are a variety of costs associated with the participation in the auction. Weak bidders will give these costs more weight than the prospect of winning a license.

A further interesting intuition regarding asymmetries among bidders is that optimal auction designs will typically seek to favour the weaker bidders, for example by offering them 'bidding credits' (see Myerson (1981) and Maskin and Riley (2000)).

7 Conclusion

What lessons can be drawn from the material in this paper? Frequently, people complain about the advice that economists give, arguing that, if one asks 10 economists for advice, one will get 11 different opinions. This complaint is

understandable, but probably not always justified. In this specific case, the questions ‘What auction should I use as a seller?’ and ‘How should I bid in this auction?’ do not allow for a unique answer, i.e. the answer will be context dependent and cannot be determined by theory alone.

Economic theory offers two types of general theorems that are very useful. The first type of ‘useful’ theorems are ‘equivalence theorems’ that inform us that it really does not matter what one does. The Revenue Equivalence Theorem falls within this class, just as the Modigliani-Miller Theorem from the area of finance. Such theorems are useful as a theoretical benchmark, as a starting point: if certain conditions hold, it does not really matter what one does. In practice, the conditions underlying these theorems need not hold and this invites further theoretical development about what happens when one of the maintained assumptions is violated. A detailed investigation then may show what is the best thing to do. Since Vickrey started the theoretical study of auctions, this is exactly what has happened in the academic literature on auctions.

This study has shown that the answer of what to do does not admit an easy and uniform answer: the ‘optimal’ auction depends on the details of the situation. As Paul Klemperer has said: ‘Auction design is a matter of horses for courses, not one size fits all’. This brings us to the second type of ‘useful’ theorems that economics offers. These are the so-called impossibility theorems that inform us that it is impossible to find a mechanism that satisfies all the properties that one wants. The most well-known impossibility theorem is Arrow’s Theorem about the impossibility of aggregating individual preferences into a consistent welfare function. In the field of auctions, there is a similar theorem: there is no single auction that is always best. The auction to use depends on the circumstances of the case and, as several examples in this paper have shown, getting the details of the auction right may influence whether or not the auction will be a success. The material from the previous sections may be helpful in putting some structure on these details, so that the auction designer can see the forest from the trees, allowing him to implement a better design.

References

- Ashenfelter, O. 1989, 'How Auctions Work for Wine and Art', *Journal of Economic Perspectives* 3: 23-36.
- Ausubel, L. 2003, 'A Generalized Vickrey Auction', *Econometrica*, forthcoming.
- Ausubel, L. and P. Cramton 1996, 'Demand Reduction and Inefficiency in Multi-Unit Auctions', Mimeo, University of Maryland.
- Ausubel, L. and P. Milgrom 2002, 'Ascending Auctions with Package Bidding', *Frontiers of Theoretical Economics* 1 1, Article 1, Berkeley Electronic Press.
- Binmore, K. and P. Klemperer 2002, 'The Biggest Auction Ever', *The Economic Journal* 112: C74-C96.
- Binmore, K. and J. Swierzbinski 2000, 'Treasury Auctions: Uniform or Discriminatory', *Review of Economic Design* 5: 387 - 410.
- Brusco, S. and G. Lopomo 2002, 'Collusion via Signalling in Simultaneous Ascending Bid Auctions with Multiple Objects and Complementarities', *Review of Economic Studies*, forthcoming.
- Bulow, J. and P. Klemperer 1996, 'Auctions Versus Negotiations', *American Economic Review* 86: 180-94.
- Capen, E., R. Clapp and W. Campbell 1971, 'Competitive Bidding in High-Risk Situations', *Journal of Petroleum Technology* 23, 641-53.
- Cybernomics 2000, 'An Experimental Comparison of the Simultaneous Multiple Round Auction and the CRA Combinatorial Auction', Report to the Federal Communications Commission.

- Damme, E. van 1999, 'The Dutch DCS-1800 Auction', in Fioravante Patrone, Ignacio García-Jurado, Stef Tijs (eds.), *Game practise: Contributions from applied game theory*, Kluwer Academic Publishers.
- Damme, E. van 2002a, 'The Dutch UMTS-auction', *Ifo Studien, Zeitschrift für Wirtschaftsforschung* 1/2002: 175-200
- Damme, E. van 2002b, 'Pompen en verzuipen' (Pumping and drowning), *Economisch Statistische Berichten* 87-4354: 271-73
- Demange, G., D. Gale and M. Sotomayor 1986, 'Multi-item auctions', *Journal of political economy* 94: 863-72.
- Engelbrecht-Wiggans, E. and C. Kahn 1998, 'Multi-Unit Auctions with Uniform Prices', *Economic Theory* 12: 227-58.
- Gilbert, R. and D. Newbery 1982, 'Preemptive Patenting and the Persistence of Monopoly', *American Economic Review*, 72: 514-26
- Graham, D. and R. Marshall 1987, 'Collusive Behavior at Single-Object Second-Price and English Auctions', *Journal of Political Economy* 95: 1217-39.
- Grimm, V. F. Riedel and E. Wolfstetter 2002, 'Low Price equilibrium in multi-unit auctions: the GSM spectrum auction in Germany', Discussion paper SFB 373, Humboldt University, Berlin, June.
- Janssen, M., A. Ros and N. van der Windt 2001, 'De draad kwijt?', Onderzoek naar de gang van zaken rond de Nederlandse UMTS-veiling (Research into the proceedings of the Dutch UMTS-auction), Erasmus University Rotterdam.
- Jehiel, P. and B. Moldovanu 1996, 'Strategic Non-Participation', *RAND Journal of Economics* 27: 84-98.

- Jehiel, P. and B. Moldovanu 2000a, 'Auctions with Downstream Interaction Among Buyers', *RAND Journal of Economics* 31: 768-91.
- Jehiel, P. and B. Moldovanu 2000b, 'A Critique of the Planned Rules for the German UMTS/IMT-2000 Licence Auction', Mimeo, University College London and University of Mannheim
- Jehiel, P. and B. Moldovanu 2001, 'The European UMTS/ IMT-2000 Licence Auctions', Mimeo, University College London and Universität Mannheim.
- Kagel, J. 1995, 'Auctions: A Survey of Experimental Research' in J. Kagel and A. Roth (eds.), *The Handbook of Experimental Economics*, Princeton University Press, 501-85.
- Klemperer, P. 2002, 'What Really Matters in Auction Design', *Journal of Economic Perspectives* 16: 169-89.
- Krishna, V. and R. Rosenthal 1996, 'Simultaneous Auctions with Synergies', *Games and Economic Behavior* 17: 1-31.
- Krishna, V. 2002, '*Auction Theory*', Academic Press.
- Leonard, H. 1983, 'Elicitation of Honest Preferences for the Assignment of Individuals to Positions', *Journal of Political Economy* 91: 461-79.
- Levin, D. And J. Smith 1994, 'Equilibrium in Auctions with Entry', *American Economic Review* 84: 585-99.
- Maasland, E. and S. Onderstal 2002a, 'Auctions with financial externalities', CentER Discussion Paper 21, Tilburg University.

Maasland, E., and S. Onderstal 2002b, 'Optimal auctions with financial externalities', CentER Discussion Paper 22, Tilburg University.

Maskin, E. and J. Riley 2000, 'Asymmetric Auctions', *Review of Economic Studies* 67: 413-38.

McAfee, P. and J. McMillan 1987, 'Auctions with Entry', *Economics Letters* 23: 343-47.

McAfee, P. and J. McMillan 1992, 'Bidding Rings', *American Economic Review* 82: 579-99.

McAfee, P., J. McMillan and P. Reny 1989, 'Extracting the Surplus in a Common Value Auction', *Econometrica* 57: 1451-60.

McMillan, J. 1994, 'Selling Spectrum Rights', *J. Economic Perspectives* 8: 145-62.

Milgrom, P. 2000, 'Putting Auction Theory to Work: The Simultaneous Ascending Auction', *Journal of Political Economy* 108: 245-72.

Milgrom, P. and R. Weber 1982, 'A Theory of Auctions and Competitive Bidding', *Econometrica* 50: 1089-122.

Myerson, R. 1981, 'Optimal Auction Design', *Mathematics of Operations Research* 6: 58-73.

Perry, M. and P. Reny 1999, 'On the Failure of the Linkage Principle in Multi-Unit Auctions', *Econometrica* 67: 895-900.

Pesendorfer, W. and J. Swinkels 1997, 'The Loser's Curse and Information Aggregation in Common Value Auctions', *Econometrica* 65: 1247-81.

Rothkopf, M., T. Teisberg and E. Kahn, 1990 'Why are Vickrey auctions rare?', *Journal of Political Economy* 98: 94-110.

Vickrey, W. 1961, 'Counterspeculation, Auctions and Competitive Sealed Tenders', *Journal of Finance* 16: 8-37.

Vickrey, W. 1962, 'Auctions and Bidding Games', in *Recent Advances in Game Theory*, Princeton University Press, 1962.

Williamson, O. 1976, 'Franchise bidding for natural monopolies: In general and with respect to CATV', *Bell J. Economics* 7: 73-104.

Wilson, R. 1979, 'Auctions of Shares', *Quarterly Journal of Economics* 94: 675-89.