

Auctioning Public Assets

Analysis and Alternatives

Edited by

Maarten Janssen



PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge, CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa
<http://www.cambridge.org>

© Cambridge University Press 2004

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2004

Printed in the United Kingdom at the University Press, Cambridge

Typeface Plantin 10/12 pt. *System* L^AT_EX 2_ε [TB]

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Auctioning public assets: analysis and alternatives / Maarten Janssen (editor).
p. cm.

Includes bibliographical references and index.

ISBN 0 521 83059 1 (hb) – ISBN 0 521 53757 6 (pb)

1. Telecommunication – Great Britain. 2. Telecommunication – Netherlands.

3. Government auctions – Great Britain. 4. Government auctions –

Netherlands. I. Janssen, Maarten Christiaan Wilhelmus, 1962–

HE8094.A93 2003

352.5'54 – dc21 2003055155

ISBN 0 521 83059 1 hardback

ISBN 0 521 53757 6 paperback

The publisher has used its best endeavours to ensure that URLs for external websites referred to in this book are correct and active at the time of going to press. However, the publisher has no responsibility for the websites and can make no guarantee that a site will remain live or that the content is or will remain appropriate.

Contents

<i>List of figures</i>	<i>page</i> vii
<i>List of tables</i>	viii
<i>List of contributors</i>	ix
<i>Preface</i>	xiii

Introduction	1
MAARTEN JANSSEN	

Part I Theory

1 Auction theory for auction design	19
TILMAN BÖRGERS AND ERIC VAN DAMME	
2 Beauty Contest design	64
MAURICE DYKSTRA AND NICO VAN DER WINDT	
3 Preventing collusion among firms in auctions	80
TIMOTHY C. SALMON	
4 Levelling the playing field in auctions and the prohibition of state aid	108
EMIEL MAASLAND, YVES MONTANGIE AND ROGER VAN DEN BERGH	
5 Allocation mechanisms and post-allocation interaction	130
MAARTEN JANSSEN AND BENNY MOLDOVANU	

Part II Case studies

6 Spectrum auctions by the United States Federal Communications Commission	147
TIMOTHY C. SALMON	

vi	Contents	
7	An analysis of the European 3G licensing process EMIEL MAASLAND AND BENNY MOLDOVANU	177
8	Auctions of gas transmission access: the British experience TANGA MORAE McDANIEL AND KARSTEN NEUHOFF	197
9	The design of Treasury bond auctions: some case studies JOSEPH SWIERZBINSKI AND TILMAN BÖRGERS	230
10	Matching markets BENNY MOLDOVANU	257
11	Competitive procurement of reintegration services in the Netherlands MAURICE DYKSTRA AND JAAP DE KONING	273
12	The provision of rail services LUISA AFFUSO AND DAVID NEWBERY	296
	<i>Index</i>	310

Figures

2.1	Different procedures can lead to different outcomes	<i>page 73</i>
6.1	Mobile telephone operator coverage by county	170
8.1	Competition in the gas supply sector, UK	201
8.2	Gas consumption, UK	203
8.3	Shares of UK gas production	203
8.4	Gas pipelines in the UK and the North Sea	205
8.5	St Fergus and Bacton constraint costs	208
8.6	Elf Exploration's excess nominations	209
8.7	Price in secondary capacity auction	210
8.8	NTS charges and floor prices	212
8.9	Energy availability and sales	213
8.10	Percentage by which auction price exceeds reserve price: average over four rounds	220
8.11	Monthly averaged price difference between day-ahead spot markets for gas at St Fergus and National Balancing Point and weighted average prices paid in auction for monthly entry rights	221
8.12	Available entry capacity in monthly auctions for St Fergus	222
8.13	Monthly averages of spot prices at the National Balancing Point and St Fergus	223

Tables

1.1 An example to illustrate the multi-unit Vickrey auction	<i>page 30</i>
1.2 Description of the state of a simultaneous ascending auction	32
1.3 Player's bidding rights in round t in a simultaneous ascending auction	33
1.4 Increasing marginal values (for player 1)	47
1.5 The exposure problem	49
1.6 The free-rider problem	50
2.1 Bidding rules	74
2.2 Evaluation of bidding rules	76
3.1 Bidder values	84
4.1 Bandwidth in MHz for each licence	123
5.1 Bidder valuations	139
6.1 Bidder values	148
7.1 Classification of licensing methods	178
8.1 Capacity buy-backs, April 2000–February 2001	224
9.1 Summary of case studies	254
10.1 Reserve prices	268
10.2 Possible matching	268
11.1 Organisations involved in social security and active labour market policy in the Netherlands	276
12.1 The bid-scoring system in Argentina	305

1 Auction theory for auction design

Tilman Börgers and Eric van Damme

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 licence 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 licences 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, whom 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 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,

We are grateful to Maarten Janssen for inviting us to write this survey. Preparing it has made us painfully aware of the gaps in our knowledge of the auction literature and we apologise for all errors and omissions. Tilman Börgers gratefully acknowledges financial support from the Economic and Social Research Council (UK).

such as efficiency and revenue. In section 6 we pull our insights together, and ask which policy lessons our analysis suggests. Section 7 concludes.

At the outset, we wish to stress the limitations of this chapter: the reader should be aware that it is a theoretical study. 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 this chapter, empirical or experimental evidence is cited where it is particularly prominent, but it is not surveyed systematically. Therefore, what we say here does not in itself provide a basis for policy recommendations. Put differently, any policy implication that is derived from the theory expositied in this chapter should be prefaced with the qualification: ‘if the theory captures practice well, then policy should be . . .’.

2 Why auctions?

Governments allocating spectrum licences to mobile telephony companies or, more generally, licences to operate in 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 more 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 knew 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 practice, the seller does not have the required information, and in these circumstances he may either set the price too low, in which case he would not expropriate what the market could bear, or set the price too high, so that he would not succeed in selling the item.

An ascending auction provides an attractive alternative. In such an auction, each potential 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 negotiating 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 that they will be closely watched in the negotiation. They will be very wary 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 the bidder who values it most. It is this property that also makes auctions an attractive selling mechanism for governments. As for 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 a government might also 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 licences 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 against 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. Second, 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 licences 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 licence auctions. The main externality is that a benevolent government will sell the licences (also) having consumer welfare in mind. Consumers, however, are not participating directly in the auction and, as a result, the outcome in which the licence 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 licence to operate in a market in which one player is already active. For a newcomer, the licence represents the right to compete, while for the incumbent it offers the opportunity to maintain a 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 licence. Therefore, an auction will allocate the licence 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.

Thus, in order to 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 licence. 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 (see chapter 4 by Maasland, Montangie and van den Bergh for further discussion, in particular about whether such auctions might violate basic EU principles by 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, and 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 licence, 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 used and the subjective elements in a Beauty Contest, there might be more potential for corruption of government officials in a Beauty Contest than in an auction. See the Introduction to this book for more details on this issue.

Summarising 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 this means 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 chapter, 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 licences, 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 maximise revenue, or should it aim for market efficiency? One argument for suggesting that governments might be concerned about auction revenues is that such revenues might allow governments to reduce more distortionary 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 licences 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, see Janssen and Moldovanu (this volume, chapter 5). One example has already been given in the previous section. As another example, think of a government selling licences 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 licences; however, an

outcome in which all stations broadcast similar music will normally not be efficient. In such a case, if the government wants to achieve an efficient outcome, it should impose conditions on some of the licences, which will typically reduce revenue. This example again 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 one of *what* will be sold. An example where this clearly mattered was the recent European UMTS auctions. There the question was: 'How large (in terms of spectrum) should a UMTS licence be?' It was not clear how much spectrum a UMTS operator would need, and therefore how many licences could be fitted into the available spectrum. Thus, it was not clear 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 licences, but rather abstract blocks of spectrum, and bidders could themselves choose the number of blocks for which 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 licences, but also how much spectrum was actually needed for third-generation spectrum licences, and thereby 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 licence is, but rather 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 what 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 judgement about the optimal size of licences, and to let the auction determine only 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 into the auction.

This need arises especially 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 licence winners enjoying market power in 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 forms another safeguard against non-serious bids. In some cases one may go further in restricting the set of admissible bidders. For example, if licences 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 entries into the auction.

The timing of auctions is also important. Consider again the experience with the European UMTS auctions. Governments that were early in auctioning their licences have typically earned (much) higher revenue per capita than those that were late. The UK was the first country to auction its licences and, therefore, the UK was, in effect, not only auctioning a licence to operate in the UK, but the option to construct a pan-European network. This option might have made the UK licence more valuable, so it might have attracted more bidders to the UK auction, with higher revenues as a natural consequence. Similarly, if the German UMTS auction had taken place later in time, the tide might have turned and the Sonera/Telefonica consortium might have realised that a six-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 ‘whom 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: sealed-bid or open. In *sealed-bid formats*, bidders simultaneously and independently submit a bid, possibly in a sealed envelope, or perhaps using a more modern communication technique. These bids are then opened and the auction outcome is determined following some rules that have been announced in advance. In *open auction procedures*, 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 practice, 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 be denied this information, etc.

A second open procedure is the descending, or Dutch, 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 at which he stopped the clock. Note the important difference from 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 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 auction 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 ‘first-price sealed-bid’ auction is conducted among 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 of one object*

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 at which a bidder reduces his demand may reveal important information to the competing bidders.

Formally, in the ascending, 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 is then allocated $d_i(p^*)$ units at a price p^* for each. Hence, all units sell at the same price. In practice, different variants may be distinguished: bidders may, or may not, know the demand as expressed by their competitors; they may, or may not, be prevented from increasing their demand again after they have previously reduced it, etc.

In the descending, 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. Various 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 the amount they bid.

Formally, in the uniform price auction, each bidder i communicates his entire demand curve $d_i(\cdot)$ directly 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.

In the discriminatory auction, the bidders also 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 indicates that he wants five units and that he is willing to pay p_1, p_2, \dots, p_5 respectively for these units with $p_1 > p_2 > p_3 > p_4 > p_5$ and the market-clearing price p^* satisfies $p_3 > p^* > p_4$, then bidder i will be awarded three 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-unit case, there is now one important difference. It is that 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

Table 1.1. *An example to illustrate the multi-unit Vickrey auction*

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

‘Vickrey auction’ the highest bids are again 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 generalisation of the single-unit Vickrey rule and it has a clear economic interpretation. In the single-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 each of whom is interested in at most four units. The bidders’ marginal values are given in table 1.1. (The table should be read as follows: bidder 1 expresses a value (bid) of 50 for the first unit that he gets, 47 for the second unit, etc.). The Vickrey auction allocates three units to bidder 1, one to bidder 2 and two to bidder 3, as indicated by the entries marked * 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 three units more to bidders 2 and 3. Of these we would give two units to bidder 2 (values 28 and 20) and one unit to bidder 3 (value 24). Consequently, bidder 1 should pay 28, 24 and 20 for his units, a total of 72. Similarly, bidder 2 should pay the externality he exerts on bidders 1 and 3, i.e. he should pay 32. Finally, bidder 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 in the affirmative, as has recently been shown by 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 bidder 1 by one unit and we continue the process. We repeat this process, always allocating one unit to a bidder 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 1.1. 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 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 licences sold by auction may differ in size, or in their location in the electromagnetic spectrum. These objects may have different values, and so will fetch different prices.

When heterogeneous objects are sold, both sequential and simultaneous sales are again possibilities. In the case of 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

Table 1.2. *Description of the state of a simultaneous ascending auction*

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}

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 Federal Communications Commission planned to sell multiple, non-identical spectrum licences at 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 may be represented as shown in table 1.2. The columns of this table correspond to the various lots; B_j^t denotes the bidder that is standing highest 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 per cent, above the previous highest bids.

Table 1.3. *Player's bidding rights in round t in a simultaneous ascending auction*

1	2	...	M
R_1^t	R_2^t	...	R_m^t

In addition to information on the lots, bidders also have information about the number of 'bidding rights', R_j^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_j^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

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

lots on which he is not standing high. The auction rules will determine how the bidding rights evolve, so, in addition to table 1.2, in each round the table of remaining bidding rights will also be available to players (table 1.3). The rules may, for example, reflect concerns about competition in the aftermarket, so that bidders are not allowed to acquire more than a 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 does not bid for a sufficiently large number of objects.

Let us give an 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 bidder i . Second, the number of bids that bidder 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 tables, such as those in tables 1.2 and 1.3, 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, and 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, note that, in this simultaneous ascending auction, bidders bid on individual lots; there is no possibility of bidding directly on packages. As we shall 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 normalised to 0.

In the English auction, as long as the price is below the 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 *whether* 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 auctions 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 chapter. Suppose also for the present that each bidder knows not only 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 by Vickrey (1961) was that this equivalence of auction forms generalises 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

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

As a consequence, in this benchmark case, the bidder with the highest value will win the object, so 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 will be equal to the expectation of the price paid in the equilibrium of the Vickrey auction. It also follows, therefore, 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 what conditions player 1 finds it optimal to bid $B(x)$ for any possible value x that he might have. If player

1 bids $B(y)$ 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 the assumptions, first, that $B(\cdot)$ is increasing, so that the bid $B(y)$ is winning if and only if $y > v_2$ and, second, that v_2 is uniform on $[0, 1]$ so that $y = \text{Prob}[v_2 < y]$. Player 1 wants to maximise his payoff, so 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 , and thus 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 generalise 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 his probability of winning the auction, $P(x)$, as well as the expected transfer, $T(x)$, he will have to make to the seller. Furthermore, the buyer can calculate the corresponding quantities resulting from his pretending that his value

would be y . If a bidder plays as if his value were 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 would not 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 also follows that 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 of the lowest conceivable type. This result is known as the *Revenue Equivalence Theorem*.

Without a 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 among these auction formats.

Let us now ask the question of which auction format the seller should choose? The Revenue Equivalence Theorem implies that this reduces to the question of which function $P(\cdot)$ to choose, and what value for $U(0)$. If the seller is only interested in the efficiency of the allocation rule, then the four auction formats discussed above, with zero reserve price,