Steinberg, Richard (2012). "Auction Pricing", Chapter 27 (pp. 679-712) in: O. Ozer and R. Phillips (eds), *Oxford Handbook of Pricing Management*, Oxford: Oxford University Press.

CHAPTER 27

AUCTION PRICING

RICHARD STEINBERG

27.1 INTRODUCTION

Why auction? The primary reason is that the seller is sufficiently unsure of the valuations of the potential buyers that he requires additional information before he can allow the sale of the items to take place. Two additional reasons are the speed and transparency of the auction process. Formally, we define an *auction* as a procedure to: (1) elicit information from potential buyers regarding their valuations for a specified set of items available for sale; and (2) based on this information, determine an allocation of the items to the potential buyers along with the individual payments required from each.[1]

This chapter is dedicated to the memory of Michael H. Rothkopf, a true leader in the development of auction pricing.

The author would like to thank Lawrence Ausubel, Michael Ball, Octavian Carare, Peter Cramton, Robert Day, George Donohue, Guillermo Gallego, Karla Hoffman, Peter Key, Thomas Kittsteiner, Paul Klemperer, Evan Kwerel, Ailsa Land, Edward Lazear, Preston McAfee, Paul Milgrom, Michael Ostrovsky, David Parkes, Martin Pesendorfer, Stephen Rassenti, Tuomas Sandholm, Itai Sher, Tunay Tunca, and Robert Wilson, as well as the editors, Özalp Özer and Robert Phillips, for comments. However, the accuracy of the material in this chapter, and the views expressed herein, are the responsibility of the author alone. The author would also like to thank Paul Milgrom and the Department of Economics, Stanford University, for inviting him to be Visiting Scholar for Stanford University's Fall term 2008, during which time much of the work on this chapter was completed.

[1] The basic auction terminology is fairly obvious, but for completeness we specify that: the potential buyers who chose to submit their information are called the *bidders*; the information provided by a bidder is his *bid*; the party who elicits the information from the bidders and allocates the items to them is the *auctioneer,* whom we will not usually distinguish from the *seller* (in practice, the former typically works on behalf of the latter); the bidders who are allocated one or more items are the *winners*. A bid includes both a set of items and the

An auction can be used to sell a single item, or multiple units (i.e., a number of identical items), or multiple items (i.e., a number of items, generally nonidentical). Not all auctions are for discrete indivisible goods, like paintings or cattle, but can also be for divisible goods like electricity or the right to emit carbon dioxide. Auctions have been in existence since the fifth century B.C. (if we are to believe Herodotus), and the variety of items sold throughout history is staggering. For a detailed treatment of the institutional aspects of auction markets up until the mid-1960s, see Cassady (1967). For a more informal exposition, but one that is current to the mid-1980s, see Learmount (1985).

We present here a non-technical but rigorous tutorial on the state-of-the-art of auction theory, with an eye toward the real world of pricing. We include some historical details, since knowledge of the evolution of the ideas surrounding auctions is essential to understanding both the modern theory and any current applications. For the most part our presentation is structured by topic, rather than by paper. The significant exception is our discussion in Section 3 of the preeminent paper of Vickrey (1961) that established the field.

Five auction topics are not covered here (except in passing), since each comprises an area that is now so extensive as to require a major review of its own: sequential auctions, online auctions, procurement auctions, experimental methods, and empirical approaches. For sequential auctions, see Krishna (2010, Chapter 15); Krishna's presentation is mathematical, however his chapter notes constitute an excellent non-technical overview. For online auctions, see Ockenfels, Reiley, and Sadrieh (2006); for procurement auctions, see Bichler, Davenport, Hohner, and Kalagnanam (2006); and for online procurement auctions, see the Special Issue of *Production and Operations Management* edited by Bichler and Steinberg (2007). For experimental methods, see Kagel and Levin (2013); and for empirical approaches, see Athey and Haile (2006) and Hendricks and Porter (2007).

---

associated *bid price* (which may in fact be a set of prices), but we follow the common practice of also using the term "bid" to refer to the bid price alone. A bidder's *valuation* is the maximum amount he is willing to pay; there are a number of equivalent terms for valuation, including *value*, *willingness to pay*, and *reservation price*, but not *reserve price*, which is the minimum price at which the auctioneer is willing to sell. Two bidders are said to be of the same *type* if they have identical information, beliefs, and preferences relevant to their decision-making. A *reverse auction*, also known as a *procurement auction*, is one in which the auctioneer is a buyer and the bidders are sellers. The theory for reverse auctions is fundamentally the same as for *forward auctions*, the latter term being used when necessary to distinguish the more familiar type of auctions from reverse auctions.

27.1.1 Efficiency vs. Optimality

Let us begin by immediately clearing up a common point of misunderstanding about auctions when discussed by non-economists, viz., the meaning of the word "efficient." This important term in the auction lexicon does not mean revenue maximizing or cost minimizing. An auction is said to be *efficient* if it allocates the items to those who value them the most *ex post*. In contrast, an auction is said to be *optimal* if it results in the most profitable allocation for the seller, e.g., if it maximizes expected revenue or minimizes expected cost. Efficient auctions are not necessarily optimal. Optimal auctions are not necessarily efficient.

Which is the proper goal of an auction, efficiency or optimality? Optimality seems more natural in many circumstances, but efficiency might be the appropriate objective, for example, in the case of auctions used to facilitate the transfer of public resources to the private sector. Thus, when Vice President Al Gore opened the first series of U.S. spectrum license auctions on December 5th, 1994, he proclaimed their purpose as being "to put licenses into the hands of those who value them the most" (White House 1994). The fact that these nine auctions ultimately raised $20 billion (Cramton 1998) was presumably just a convenient truth.

27.1.2 Terminology

As mentioned above, the primary reason items are priced via an auction rather than other methods is that the seller is sufficiently unsure of the valuations of the potential buyers that he has an incentive to elicit additional information from them. This is an example of *information asymmetry*, that is, there is information available to one or more parties that is not available to others. In fact, William Vickrey shared his 1996 Nobel Prize in Economics for "fundamental contributions to the economic theory of incentives under asymmetric information" (Royal Swedish Academy of Sciences 1996).[2]

27.1.2.1 Private values

A specific case of information asymmetry is an auction in which each bidder knows his own value of the item at the time of the bidding; this is called a *private values* setting. In this scenario no bidder knows with certainty the valuations of the other bidders but, even if he were to learn what they are, this would not alter his valuation. An *independent private values* setting is a private values setting in which the bidders' values are independent random variables.

[2] Four other auction theorists have been awarded the Nobel Prize in Economics: Vernon L. Smith, "for having established laboratory experiments as a tool in empirical economic analysis, especially in the study of alternative market mechanisms" [awarded 2002]; and Leonid Hurwicz, Eric S. Maskin, and Roger B. Myerson, "for having laid the foundations of mechanism design theory" [2007].

Independent private values is both the most unrealistic auction setting and the one most studied by economists. Milgrom (2004, p. 157) provides an honest assessment of the situation:

> Relaxing the private-values and independence assumptions raises a host of new issues... Bidders' ignorance of their values leads us to study what information bidders are likely to acquire, whether they will share this information or keep it secret, and whether the auctioneer can improve the outcome by gathering and disseminating information on its own. The independence assumption is an essential premise of... the revenue equivalence theorems. Relaxing this assumption forces us to reevaluate the most basic results of auction theory.

That being said, it is nevertheless common to add two further assumptions to the independent private values model. First, the bidders are *symmetric*, that is, the bidder types are not only independent but identically distributed according to some continuous density. Second, the bidders are *risk neutral*, that is, each seeks to maximize his expected profit, also called his *bidder surplus*, the difference between his valuation and the price he pays. All these assumptions are so standard, beginning with the seminal paper of Vickrey (1961), that the independent private values model with symmetric, risk-neutral bidders is often referred to as the *benchmark model* (see, e.g., Milgrom 2004). One hastens to add that the benchmark model does lend considerable insight into more general and applicable auction settings.

27.1.2.2 Common value and the winner's curse

Let us leave aside, for now, the benchmark model. It is possible, and in fact not unusual, that a bidder's precise valuation for an item is unknown to himself at the time of the bidding. However, he may have some private information that is correlated with the item's true value to him; this information is called his *signal*. An example of a signal would be an expert opinion or a test result. From an information standpoint, the opposite of the private values model is the *common value model*, the scenario in which all bidders have the same valuation for the item, but where this valuation is unknown to the bidders at the time of the bidding. In a common value setting, different bidders in general have different signals.

What is arguably the most influential paper on common value auctions is not a theoretical analysis by an auction theorist ensconced in academia, but rather is a thought piece by three petroleum engineers employed by the Atlantic Richfield Company. In their paper "Competitive Bidding in High-Risk Situations," E.C. Capen, R.V. Clapp, and W.M. Campbell (1971) described what has become known as the *winner's curse*.[3] As might be expected, the value of an oil exploration lease depends on how much oil is under the particular tract covered by the lease. Capen and his colleagues ask the reader to play a little game they call a "think sale." Think of yourself as a manager whose task is to set bids on parcels in an impending sale. On any of your parcels, you will have a consensus property value put together by your experts. Allow that *on average*, your value estimates are correct. However, your opponents *on this particular tract* will have better or worse information—that is, information that is either more accurate or less accurate than yours. Thus, there will be quite a divergence of opinions as to the tract's value, where some bidders have overestimated the true value of the parcel, and others have underestimated it. Capen et al. (1971, p. 643) ask:

> Can we not then conclude that he who thinks he sees the most reserves, will *tend* to win the parcel in competitive bidding? This conclusion leads straightway to another: *In competitive bidding, the winner tends to be the bidder who most overestimates true tract value*. (Capen et al. 1971: 643)

In other words, winning against a number of rivals in a common value auction implies that your value estimate is in fact an overestimate, conditional on the event of winning. In fact, this happens not only in a "pure" common value setting. As long as there is *any* common value component to the bidder valuations—i.e., the general model of interdependent values discussed immediately below—then you are at risk of falling victim to the winner's curse. This raises the question: In formulating a bid, how do you avoid being cursed by this "adverse selection" effect? Capen, Clapp and Campbell suggest placing "a lower bid than one might come up with otherwise," i.e., to engage in *bid shading*, a strategy that is in fact backed up by theory.

Competitive bidding becomes especially interesting when one of the parties knows the value of the item with certainty and the others do not. Consider the following refinement of the above scenario. You are bidding against another company that already owns oil rights to a contiguous parcel and therefore, by drilling "offset control wells" on the boundary, the other company has obtained nearly perfect information on the value of the rights on the parcel in which you are both interested. Your company is not so fortunate in that it has no contiguous parcel, and therefore only has imperfect seismic and other information on which to act. How should you bid? Actually, this question applies more generally to the case where your opponent does not have perfect information,

[3] Despite a great many citations, the phrase "winner's curse" does not appear anywhere in the paper of Capen, Clapp, and Campbell (1971). The term appears on page 1078 of Oren and Williams (1975), who attribute it to a conference presentation of Capen, Clapp, and Campbell.

just better information than you. This question, motivated by a case study of two major oil companies bidding via a sealed tender for rights to an offshore parcel, was considered by Wilson (1967, 1969). His analysis of it, together with the highly influential (though unpublished) PhD dissertation of Armando Ortega-Reichert (1968), opened up the entire field of common value auctions.

27.1.2.3 Interdependent and affiliated values

A general model lying between the two extreme settings of private values and common value allows each bidder's value to be a function of his own signal as well as the signals of all the other bidders, which are typically unknown to him. This setting is called *interdependent values*. This setting models the familiar situation in which a bidder's precise valuation is unknown at the time of the auction but would be affected by information available to other bidders.

An important special case, introduced by Milgrom and Weber (1982), is that of *affiliated values* in which the winning bidder's payoff depends upon his personal preferences, the preferences of others, and the intrinsic qualities of the item being sold. This is a strong form of positive correlation that allows for both private values uncertainty and common value uncertainty. Affiliation between two given bidders implies that if one bidder has a high estimate of value, it is more likely that the other bidder's estimate of the value is high. Klemperer (2004a, Chapter 1) provides the following intuition: "[Y]our value for a painting may depend mostly on your own private information—how much you like it—but also somewhat on others' private information—how much they like it—because this affects the resale value and/or the prestige of owning it."

27.2 THE FOUR BASIC AUCTIONS

The four basic auctions are the English auction, the Dutch auction, the Vickrey auction, and the first-price sealed-bid auction. Each of these, in its simplest form, is an auction for a single item. Although economists tend to be more interested in efficiency than revenue, one of the first results in auction theory, which is due to Vickrey (1961), is that under the benchmark model the expected revenues of the four basic auctions are identical (see Section 27.3.3).

The English and Dutch auctions are both *dynamic auctions*, i.e., they provide multiple opportunities for bidders to bid, where some information about the bidding is generally revealed to the bidders during the course of the auction. Thus, bidders have the opportunity to adjust their subsequent bids accordingly; this process is called *price discovery*.

27.2.1 The English Auction

The *English auction* is an ascending-bid auction. Bidders submit successively higher bids for the item until no bidder is willing to bid higher. The last bidder to bid wins the item and pays the amount of his last bid.

The English auction is probably the format most people think of when they think of auctions, especially for auctions of fine art, and indeed, almost all fine art is auctioned via the English auction. The major auction houses are the English firms Christie's and Sotheby's. Beggs and Graddy (2009) explain how art auctions work. If someone wishes to sell a piece of art, an expert at the auction house will provide advice on the likely valuation of the piece, and together the seller and the expert will agree on a secret reserve price. Prior to the auction, information on the piece is published in a presale catalogue, which includes a low- and a high-price estimate for the work; the low estimate is invariably set at or above the secret reserve price. Bidding starts low, and the auctioneer subsequently calls out higher and higher prices. When the bidding stops, the item is said to be "knocked down." However, not all items that are knocked down are in fact sold. If the bidding does not reach the level of the secret reserve price of the item, it will go unsold, in which case it will be put up for sale at a later auction, sold elsewhere, or taken off the market. Tunca and Wu (2009) report that for industrial procurement, buyers most often use a reverse English auction.

It is probably impossible to determine when the English auction was first used, although Herodotus in his *Histories*, circa 440 B.C., describes an English auction that even allowed for negative prices, i.e., subsidies.[4] This was at least a thousand years before the word "English" existed in any form. In fact, the term "English auction" is surprisingly recent, as the earliest known reference (Charlton 2008) appeared in *The Times* (London) in 1891: "All fish shall be sold by English auction" (*Times* 1891). One may hypothesize that the term arose to distinguish the auction format that had been common in England for many years from the relatively recent arrival of the Dutch auction (see Section 27.2.2).

Three interesting variations on the English auction are the *candle auction*, the *Japanese auction*, and the *silent auction*.

[4] This auction described by Herodotus, the earliest auction known, served an important societal function. For more details, see Herodotus (1914, paragraph 196). This is also discussed in Cassady (1967, Chapter 3).

27.2.1.1 The candle auction

In a *candle auction*, a short burning candle is used to determining the stopping time of the auction, at which time the highest bidder wins the item at his bid price. In his famous diary, Samuel Pepys (1926) reports observing a candle auction on two occasions for the sale of ships by the British Admiralty. The first time was on the 6th of November, 1660:

> From thence Mr. Creed and I to Wilkinson's, and dined together, and in great haste thence to our office, where we met all, for the sale of two ships by an inch of candle (the first time that ever I saw any of this kind), where I observed how they do invite one another, and at last how they all do cry, and we have much to do to tell who did cry last.

By "cry," Pepys means "cry out"; the English auction is also called an *open outcry* auction. The second entry was on the 3rd of September, 1662:

> After dinner by water to the office, and there we met and sold the Weymouth, Successe, and Fellowship hulkes, where pleasant to see how backward men are at first to bid; and yet when the candle is going out, how they bawl and dispute afterwards who bid the most first.
>
> And here I observed one many cunninger than the rest, that was sure to bid the last man and to carry it; and enquiring the reason, he told me that just as the flame goes out the smoke descends, which is a thing I never observed before, and by that he doth know the instant when to bid last—which is very pretty.

As Klemperer (2004b) points out, this is the first recorded instance of *sniping*, i.e., withholding one's true bid until shortly before the closing of the auction in the hope that it will be too late for other bidders to respond.

In England, the transition from candle auction to English auction began in the late 18th century. In 1778, two advertisements appeared in *The Times* side by side. The first was headed "Sales by Candle" and lists "The following GOODS, viz., 124 Chests of East India Indigo, 1 Barrel of French ditto, 2650 Pounds of Nutmegs, 239 ditto Mace..." The second advertisement was headed "Sale by Auction" for "All the Genuine HOUSEHOLD FURNITURE, China, Books, an Eight Day Clock, and other Effects of Mr. White, deceased" (Times 1778). Here "by Auction" refers specifically to the English auction. The candle auction was in use in England as late as the early 19th century, as can be seen from a notice of a sale by candle in *The Times* in April 1825 (Times 1825).

27.2.1.2 The Japanese auction

The *Japanese auction* is a variant on the English auction in which the bid price is raised continuously by the auctioneer, where each bidder must meet the current price to stay in the auction, and where all bidders are aware at all times as to the number of active bidders. Milgrom and Weber (1982, Section 5) analyze the English auction using the Japanese variant.

An important feature that distinguishes it from the standard English auction is that the Japanese variant precludes *jump bids*. Jump bids are bids that exceed the minimum required increment, often significantly so. Avery (1998) points out that a jump bid can serve to discourage competition, for two related reasons. First, it suggests that the jump bidder values the item more than anyone else. Second, if the jump bidder drops out in favour of your later bid, his action is strong evidence that you have overbid.

In an implementation of the Japanese auction known as the *button auction*, bidders hold down a button to show that they are active as the auctioneer steadily increases the price; when the price exceeds a bidder's valuation, the bidder releases the button, which locks to prevent him from pressing it again. Another variation is where the bidders are assembled in a room and are asked to stand at the beginning of the auction, and to sit as the price is raised above their valuation. The last bidder to remain standing pays the current price.[5]

27.2.1.3 The silent auction

One generalization of the English auction for multiple items is the *silent auction*. In a silent auction, the multiple items are auctioned simultaneously, where all items are on display to all bidders—which might in fact be a description or photograph of the item. All bidders are permitted to bid on any number of items. As the name implies, there is no "oral outcry." Next to each item is a piece of paper, called a bid sheet, which specifies the minimum price, that is, the auctioneer's reserve price, and the minimum amount that the existing high bid must be raised for the new bid to be valid, that is the *bid increment*. There may also be an estimate of the value of the item. An interested bidder can decide whether he wishes to increase the *standing high bid* by writing his name and the new bid on the item's bid sheet. Bidding on all items closes simultaneously according to a common clock.

Silent auctions are typically used when there is a charitable or public goods component to the seller's revenues. In a charity auction, this component is often the donation or reduced cost to the

[5] Auctions are popular in Japan, notably for the selling of fish, but the auctions used in Japan are not Japanese auctions. The Tsukiji fish market, the largest wholesale fish market in the world, uses the standard English auction. Nevertheless, "Japanese auction" is now a useful term to distinguish this variant of the English auction.

charity of the items for sale at the auction. Charity auctions are commonly employed in church sales. In cases where the winner reneges, the auctioneer will typically allocate the item to the second-highest bidder at his bid price and then attempt to shame the high bidder into donating the difference.[6]

At silent auctions, it is not uncommon to observe the following behavior (Milgrom 2004, Section 7.2). Very close to the announced ending time, there is someone who approaches a table, lifts the pencil and slowly writes his name and bid as the bell rings announcing the end of the auction. Often, this is the only bid the bidder ever makes for the item. The bidder's intent is to keep the price as low as possible and then place a bid only when no one has any time left to respond. In other words, this is another instance of sniping, although that term is rarely, if ever, used to describe this behavior in a charity auction context.

Sniping is not very harmful at charity auctions, Milgrom points out, since bidders might be happy to pay a higher price to acquire what they want, having contributed more to a worthy cause. However in online auctions, where people may not be feeling as charitable as at a church sale, this tactic can be a more serious issue. For more on sniping, see Ockenfels, Reiley, and Sadrieh (2006, Section 4).

27.2.2 The Dutch auction

The *Dutch auction* is a descending-bid auction, where the auctioneer starts at a high price and announces successively lower prices. The first bidder to bid wins the item and pays the price current at the time of his bid. The Dutch auction is often referred to as the *descending clock auction* since, although it can be conducted orally, it often employs a clock, with the clock hand starting at a high price determined by the auctioneer, and dropping until a buyer stops the clock by pushing a button to bid for a lot at the price determined by the clock hand.

The Dutch auction is not only of conceptual interest; in terms of quantity of items sold and magnitude of revenues generated, it is one of the most significant auctions in the world today. Cassady (1967) reported Dutch auctions being used throughout Europe, in certain Middle Eastern countries, and in particular for the sale of fish in both Hull, England and in Israel. It has been used in Australia at the Sydney Fish Market since 1989. But the most famous application of the Dutch auction is, of course, the selling of flowers in the Netherlands. On January 1, 2008, the two great Dutch flower auction cooperatives, the Aalsmeer Flower Auction and FloraHolland, merged to

[6] Popkowski Leszczyc and Rothkopf (2010) studied motives and bidding in charity auctions by conducting field experiments consisting of simultaneous charity and non-charity auctions for identical products on an internet auction site. Among their results, they found that auctions with 25% of revenue donated to charity had higher *net* revenue—i.e., after the charitable donation was subtracted—than non-charity auctions.

become the new FloraHolland. According to a press release (FloraHolland 2008), during 2007 exporters, wholesalers, florists, and major retailers bought 11 billion cut flowers and 1.2 billion house and garden plants from the Aalsmeer Flower Auction and FloraHolland. The total 2007 turnover of flowers and plants for the soon-to-be-merged cooperatives was in excess of 4 billion euros ($6.2 billion).

The day-to-day working of the Dutch flower auctions, which are auctions for multiple units, is explained by van den Berg, van Ours, and Pradhan (2001). The wall in front of the auctioning room contains a large board with a "clock" and an electronic display of the properties of the product to be auctioned, as well as the minimum price. The clock is actually a circle of small lamps, each corresponding to a given monetary value. Once the clock is set in motion, consecutive lamps light sequentially clockwise, corresponding to a decrease in the value. Just prior to the start of the auction, the flowers or plants are transported through the room and a few items are shown to the buyers. The auctioneer decides on a starting position for the clock that corresponds to an "unreasonably high" price for the product, and sets the clock in motion. The value drops continuously until some buyer stops the clock by pushing a button, where the value indicated by the clock at that moment is the price he is to pay for a single item. The buyer then announces how many units he wants to purchase, where a unit is a fixed amount of items for the product, e.g., it might be 120 flowers. The identity of the buyer is shown on the electronic display. If he does not purchase the entire lot, the clock is re-set to a high value and the process starts again for the remaining units; this continues until the entire lot is sold. If the clock passes the minimum price, the remainder of the lot is destroyed. Van den Berg et al. report that the average duration of a single auction transaction is a couple of seconds.

27.2.2.1 An historical note

There exists a general consensus that the term "Dutch auction" derives from its use in the Dutch flower auctions and that these were originally devised by farmers in Holland during the 1870s (see, e.g., Learmount 1985, p. 74); however, this consensus is incorrect. There is no doubt that the term goes back decades earlier. In March 1830, the *Virginia Literary Museum & Journal of Belles-Lettres, Arts, Sciences &c.* (an ephemeral weekly in publication for a year at the University of Virginia) published a short story entitled "The Country Belle" (K 1830).[7] The story tells of a Miss Patty Starkie who, upon leaving boarding school, had prescribed a number of qualifications she required in a husband. Unfortunately, "it always happened that some one of these essentials was wanting, or was not possessed in sufficient quantity." Consequently:

[7] "K" is probably George Tucker, the journal's editor.

> After a few years Miss Patty began to relax from the strictness of her conditions, and to follow the example of a Dutch auctioneer, whose practice is to set up his wares at the highest price, and thence bid downwards til he meets with a purchaser; but, unfortunately, however she fell in the price, it was always above the rate of the market, and the goods remained on hand; or to speak without a figure, she continued Miss Patty Starkie still, until the time of which we speak, when she had reached the sober and discreet age of thirty five.

We should point out that the term "Dutch auction" is sometimes used by practitioners to refer to some variation of the "uniform-price auction" (discussed in Section 27.3.2), such as the "Dutch auction" tender offer for common stock, and the eBay "Dutch auction" for auctioning multiple units on the internet. In this chapter, the term *Dutch auction* will always refer to the descending-bid auction described above.

27.2.2.2 The slow Dutch auction and markdown management

Just how important is speed in determining the outcome of the Dutch auction? We will see in Section 3 that one of Vickrey's key results is the revenue equivalence between the Dutch auction and the "first-price sealed-bid auction" (discussed in Section 27.2.3). However, Lucking-Reiley (1999) reports results of field experiments of various auction formats on the internet for collectable cards ("Magic: The Gathering"), and found that the Dutch auction earned approximately 30 percent more revenue than the first-price sealed-bid auction. Lucking-Reiley suggests that the longer timescales in these online auctions may be a factor, viz., days and hours, versus minutes and seconds. Carare and Rothkopf (2005) present both decision-theoretic and game-theoretic models that support this assertion. They find that bidders in such Dutch auctions, when faced with a positive cost of returning to the auction site, prefer to purchase the object sooner at a higher price so as to economize on the cost of return. Therefore, when transaction costs are accounted for, these Dutch auctions yield, on average, higher revenue than first-price sealed-bid auctions.

Thus, a distinction can be made between two types of Dutch auction according to their speed: the *descending clock auction* described earlier, and the *slow Dutch auction* considered by Lucking-Reiley and Carare and Rothkopf. The slow Dutch auction is similar to the retail pricing policy in which a merchant holding a fixed inventory of a perishable good or a fashion good finds it advantageous to *mark down*—i.e., discount—his stock rather than allowing it to depreciate or go out of style, and thus subjects it to increasingly deeper discounts until either the inventory of the product is completely sold or a final sell-by date is reached. This policy was made famous by

Filene's Basement, a chain of department stores in the Boston area. In Filene's Basement, the price tag on each item is marked with the date it arrives on the selling floor. The longer an item remains unsold, the more the price is automatically reduced; first 25%, then 50% and finally 75% of the original price; what is not sold is donated to charity.[8] Determining an optimal markdown policy is known as the *markdown management problem*. We refer the interested reader to Chapter 25 by Ramakrishnan.

The close connection between markdown pricing and the Dutch auction was first pointed out by Lazear (1986). Specifically, he observes that when markdown pricing is employed such that the reduction in price is small each time, and such that the number of potential buyers is sufficiently small that they may behave strategically with respect to the waiting time, then this procedure is essentially the Dutch auction. Lazear shows that "bargain basement" behavior can be predicted and that the price cutting rule can be specified, and he addresses the question of when a rigid rule of this sort is an optimal pricing policy. Under the assumption that all bidders have the same valuation, Lazear finds that both the initial price and the speed of the fall in price increase as the number of customers per unit of time increases, and as prior uncertainty about the value of the good increases.

Gallego, Phillips, and Sahin (2008) observe that in fashion retailing, stores often schedule the markdown of products to occur at the same times each year, hence shrewd customers quickly learn that they are likely to save if they are willing to wait. In effect, such sellers are training customers to wait. Gallego et al. investigate how this training effect might change the seller's optimal markdown policy and, in particular, whether a seller would be better off allowing merchandise to perish rather than selling it at a discount. Like Lazear, they abstract to a two-period model, where inventory perishes at the end of the second period. However, unlike Lazear, who considered the case when all customers have the same valuation for an item but the seller is *ex ante* uncertain whether that valuation is low or high, Gallego et al. assume that the customers follow a distribution of willingness to pay. They find that the seller's optimal equilibrium policy is to set a single price for both periods and to not restrict second-period sales, even if customer willingness to pay is lower in the second period than the first.

27.2.3 The first-price sealed-bid auction

In the first-price sealed-bid auction, bidders simultaneously submit sealed bids for the item; the highest bidder wins the item, and pays the amount of his bid. Klemperer (2004a, Chapter 5) mentions that an advantage of the first-price sealed-bid auction is that it does not disincentivize

[8] Ironically Filene's Basement, which had been operating under Chapter 11 bankruptcy court protection, was itself sold at auction in June 2009.

anyone from entering the bidding, in contrast to the English auction, in which a weaker (potential) bidder knows that a stronger bidder can always re-bid so as to top his bid. But, as Klemperer explains, this advantage also highlights a disadvantage of the first-price sealed-bid auction, viz., it allows bidders with lower valuations to occasionally beat opponents with higher valuations, and thus is more likely to lead to inefficient outcomes.

Applications of first-price sealed-bid auctions come in two forms: the "buy" format, where bidders are purchasers and the high bid wins, and the reverse (i.e., procurement) format, where bidders are suppliers and the low bid wins. Vickrey (1961) provides as examples the sale of property and the underwriting of securities in the former case, and bidding for construction contracts in the latter.

McAfee and McMillan (1987) state that first-price sealed-bid auctions are sometimes used in the sale of artwork and real estate, but that they are of greater quantitative significance for government contracts, including mineral rights in government-owned land, as in the paper of Capen, Clapp, and Campbell (1971) discussed in Section 27.1. McAfee and McMillan wrote: "For many government contracts, firms submit sealed bids; the contract is required by law to be awarded to the lowest qualified bidder." This situation still holds widely in the United States. Zullo (2006), for example, examined public-private contracting in Wisconsin, and found that the state's Department of Transportation (DOT) has a well-developed system for competitive bidding on road and bridge construction. Before the Wisconsin DOT issues a request for proposals, DOT engineers estimate the cost of a project and prepare detailed specifications based on both experience and industry standards. Then bidders respond with proposals, and the lowest bid is selected. However, the DOT has, in effect, a reserve price: If proposals deviate significantly from the DOT estimate, the project is either re-evaluated or cancelled.[9]

27.2.4 The Vickrey auction

The Vickrey auction is a second-price sealed-bid auction, where bidders simultaneously submit sealed bids for the item. The highest bidder wins the item and pays the amount bid by the second-highest bidder. The auction is named after William Vickrey who developed the theory behind it in his celebrated 1961 paper in the *Journal of Finance*.

[9] Although "lowest qualified bidder" is the *de facto* standard for state construction auctions in the United States, other countries take a more heterodox view. For example, in Italy, Portugal, Peru, Korea, Denmark, and Taiwan, the standard practice is to *exclude* the lowest bid, as well as the highest, and award the contract to the bid coming closest to the average of the remaining bids. See Lambropoulos (2007).

It is easily seen that in the Vickrey auction it is a dominant strategy for a bidder to bid his true value.[10] A bidder has no incentive to bid lower than his value, since such a bid will not affect the price he pays if he wins, and could lose him the item that he might have otherwise won at an acceptable price; the bidder also has no incentive to bid higher than his value, since this will win him the item he would have otherwise lost only in the case where he will be required to pay more than his value. This truth-telling property is called *incentive compatibility*. Thus, the Vickrey auction always awards the item to the bidder who values it the most, that is, it is efficient.

An important way of viewing the Vickrey auction is that *the winning bidder pays the opportunity cost of winning the item*. That is, he pays the incremental value that would be derived by assigning the item according to the next-best use among the bidders. In this way, a winning bidder achieves a profit equal to his incremental contribution to total value. Of course, in an auction with a single item, this next-best use is to assign it to the second-highest bidder, and the incremental value is the second-highest bid. However, viewing the Vickrey auction in terms of opportunity cost allows for the possibility that the auction could be generalized for multiple items. We will return to this idea in Section 27.5.5.

With regard to theory, Vickrey was a pioneer; with regard to practice, there were forerunners. In 1797 Goethe devised a second-price auction with a reserve price in order to sell a manuscript to a single bidder (Moldovanu and Tietzel 1998), while stamp auctioneers have been using second-price sealed-bid auctions since 1893 (Lucking-Reiley 2000).

The Vickrey auction was employed by the government of New Zealand in 1989-90 in the first spectrum auction ever held (see Section 27.4). Unfortunately, the New Zealand auction was not ideally designed for achieving efficiency, and evidence suggests that the outcome was indeed inefficient (see Milgrom 2004, Section 1.2.2). To make matters worse, the auction was dreadful in terms of generating revenue (see Mueller 1993). Fifteen years after the auction, in what some might describe as quintessential political understatement, the New Zealand Ministry of Economic Development (2005) issued a report stating: "The results of the second price tenders held in New Zealand attracted some criticism as the return to the Crown was below what some people considered to be the true value of the spectrum." The Ministry attributed the disappointing revenue to three factors: (1) very thin markets, (2) lack of information concerning the value of the spectrum, and (3) "instances where only nominal bids were placed and no reserves were set, meaning that some licences were essentially given away." We will discuss weaknesses of the Vickrey auction in Section 27.5.5.

---

[10] A *dominant strategy* for a bidder is one that is at least as good as any other strategy for that bidder, no matter how the other bidders may bid.

Vickrey's paper is significant for both the depth and breadth of the contributions it makes, and is required reading for anyone interested in auction pricing. We thus discuss it in some detail in the next section.

27.3 ANALYZING AUCTIONS

…………………………………………………………………………………………………...

"Counterspeculation, Auctions, and Competitive Sealed Tenders" (Vickrey 1961) is the foundation paper of auction theory. In this remarkable work, Vickrey analyzed the four basic auction types under the benchmark model.[11] In the following subsection, we present a concise summary of his analysis.

27.3.1 Analyzing the four basic auctions (benchmark model)

- *The English auction and the second-price sealed-bid (Vickrey) auction.* In an English auction, a bidder has a dominant strategy to continue bidding up until the point at which the price reaches his valuation. The bidder with the highest valuation will win the item and pay a price just above the second-highest bidder valuation. A sealed-bid procedure that is strategically equivalent to the English auction is where the winner pays the price of the second-highest bid, since it will be a dominant strategy for a bidder to bid his valuation. The intuition is as follows. Bidding in an English auction will stop at a level approximately equal to the second highest value among the values that the bidders place on the item, since at that point there will be only one interested bidder left; the object will then be purchased at that price by the bidder to whom it has the highest value.
- *The Dutch auction and the first-price sealed-bid auction.* In the Dutch auction, the first and only bid is the one that concludes the auction. Each bidder, in attempting to determine at what point he should be prepared to make a bid, will need to take into account whatever information he possesses concerning the probable bids by others. Under the benchmark model for the Dutch and first-price sealed-bid auctions, the unique equilibrium strategy determines how players should shade their bids, depending on their beliefs about the distribution of other bidders' bids. Thus, a bidder has no dominant strategy. The first-price sealed-bid auction is strategically

[11] It should be said that, until the appearance of his paper, "sealed-bid tenders" were not considered to be auctions. Vickrey himself described them in his paper as "not... auctions as such" (p. 20).

equivalent to the Dutch auction. Therefore, a bidder in the Dutch auction should wait until the price falls to what he should have bid if he were participating in a sealed-bid auction.

- *Revenue comparisons among the four auctions.* The English and Dutch auctions can be shown to produce the same average expected price and thus the same expected profit to both the seller and to the bidders, respectively. Thus, all four auctions produce the same expected profit to the seller and to the bidders, respectively. However, in order to maximize his expectation of profit, a bidder in the first-price auction, as in the Dutch auction, must concern himself not only with his own valuation, but also with his estimates of the valuations of other bidders as well as the bidding strategies that he thinks they will follow, which can involve a considerable amount of expenditure of bidder resources. In contrast, the second-price method makes any such general market appraisal—as in the English auction—entirely superfluous.

The four basic auctions under the benchmark model are summarized in Table 27.1.

Table 27.1 The four basic auctions

	Bidding procedure	
Dominant strategy exists?	Dynamic	Sealed-bid
Yes	↑ English	≡ Vickrey
No	↓ Dutch	≡ first-price sealed-bid

27.3.2 Other contributions of Vickrey

Probably due to the considerable significance of his main results, the other contributions of Vickrey's paper are often overlooked. Five other topics were covered by Vickrey in this paper: shill bids, third-price auctions, the second-price Dutch auction, multiple-unit ascending auctions, and multi-unit sealed-bid auctions.

- *Shill bids*.[12] Vickrey was aware that the second-price method may not be "automatically self-policing to quite the same extent as the first price method," and that a shill could be employed by the auctioneer to jack up the price by putting in a bid just under the top bid.

- *Third-price auctions*. What would be the equilibrium strategy in a *third*-price sealed-bid auction? Bid somewhat *higher* than one's valuation, since the danger of the payment exceeding the valuation is offset by the increased probability of gains in cases where the second-highest bid exceeds this value but the third-highest bid falls below it. Beginning with Vickrey, anyone who has ever discussed third-price auctions has felt an understandable obligation to mention that they are a theoretical construct of no known practical significance.

- *The Second-price Dutch auction*. An interesting variation on the Dutch auction would be where the winner pays the second-highest bid price rather than the first, so as to make it equivalent to the second-price sealed-bid auction. The clock mechanism could be altered so that the first button pushed would pre-select the signal to be flashed, but there would be no indication until the second button is pushed, stopping the clock, indicating both the price determined by the second button and the winner determined by the first.

This idea was followed up forty-eight years later by Mishra and Parkes (2009) who, under the assumption of private values, introduced a "Vickrey-Dutch" auction for two different environments: (i) multiple items with buyers each having demand for a single item; and (ii) multiple units with buyers each having non-increasing marginal values for additional units. Both variants of the Mishra and Parkes auction retain the advantages of the Dutch auction of speed and elicitation (which the authors show via simulation), and inherit the Vickrey property of supporting truthful bidding that terminates in an efficient allocation (which they show analytically).

Multiple-unit ascending auctions. In the multiple-unit generalization of the English auction, each bidder is interested in at most one unit. There are two variations. In the first, m units are offered simultaneously, and each bidder is permitted to raise his bid, even when this does not make it the high bid. When no bidder wishes to raise further, one unit is awarded to each of the m highest

[12] For readers outside of North America: The term *shill* refers to a seller's accomplice posing as a customer.

bidders, who pay the same price equal to the $m+1^{st}$ highest bid. The result is *Pareto optimal*, i.e., no one could be made better off without someone else being made worse off. This is a *uniform-price auction.* In the second variation, the assumption is that there may be minor quality differences among the items, in which case they can be auctioned off successively. In this variation, bidders will need to bid strategically.

Multiple-unit sealed-bid auctions. Vickrey describes bonds as often being sold in multiple-unit sealed-bid auctions, with bids accepted from highest to lowest until all units are allocated. He discusses the two variants, in which bidders are asked either to: (1) pay their bid, or (2) pay a uniform price set to the price of the last accepted bid. In the uniform-price case, Vickrey advises that it be set to the first bid rejected rather than the last bid accepted. There is still a debate on as to which of the two variants is best for selling government securities, *pay-your-bid* or *uniform-price*.

The U.S. Treasury currently uses multiple-unit sealed-bid auctions to sell Treasury bills, notes, bonds, and TIPS (Treasury Inflation-Protected Securities) in order to determine their rate or yield. Each year, the Treasury conducts approximately 200 public auctions and issues more than $4.2 trillion in securities. All Treasury auctions operate as follows. Each bidder specifies an amount (up to 35% of the issue amount) and the yield that is acceptable to him. At the close of an auction, the Treasury accepts competitive bids in ascending order in terms of their yields, until the quantity of accepted bids reaches the offering amount. All successful bidders receive the same yield at the highest accepted bid, i.e., the last bid accepted.

27.3.3 Optimal auctions

Which auction should a seller use so as to maximize his expected revenue? This is the *Optimal Auction Problem*. As discussed earlier, Vickrey showed in his classic 1961 paper that the four basic auctions under the benchmark model yield the same expected revenue to the seller. What could be described as the "Fundamental Theorem of Auctions" is a significant generalization of Vickrey's result, and is called the *Revenue Equivalence Theorem*. The result is due to Myerson and, independently, to Riley and Samuelson:[13]

[13] There are various formulations of the revenue equivalence theorem, which differ primarily in their degree of generality, and which can be described collectively as the "revenue equivalence theorems."

Revenue Equivalence Theorem (Myerson 1981, Riley and Samuelson 1981)
Consider auctions in which the item goes to the bidder with the highest signal and in which any bidder with the lowest-possible signal expects zero profit. Assume that bidders are risk-neutral and that each has a private signal independently drawn from a common, strictly increasing continuous distribution. Then every auction yields the same expected revenue, and results in each bidder making the same expected payment as a function of his signal.

Note that the theorem only specifies that the potential buyers have private *signals*, not private *values*. So, this result applies far more generally than to the private values setting considered by Vickrey (1961).

Myerson also showed that any of the four standard auction formats, together with an appropriately chosen reserve price, is an optimal auction. That is, no other auction format can result in a higher expected revenue for the seller. The necessity of a reserve price was dramatically demonstrated by the New Zealand auction (discussed in Section 27.2.4).

Virtually all subsequent work on optimal auctions has built on Myerson's paper. Independently, Riley and Samuelson (1981) established the Revenue Equivalence Theorem under somewhat less general conditions. However, unlike Myerson, they also considered the case of risk-averse bidders.

27.3.3.1 The linkage principle

The Revenue Equivalence Theorem requires the strong assumption that each bidder's private information is independent of his competitors' private information. This leads to the question of what can be said under the more realistic assumption that the bidders have "affiliated values" (discussed above in Section 27.1.2). Milgrom and Weber (1982, Theorem 21) show that for bidders with affiliated values, the English auction generates higher average prices than the Vickrey auction.

Klemperer (2004a, Chapter 1) provides the following intuition. The winning bidder's surplus is due to his private information. The more that the price paid depends on others' information, the more closely the price is related to the winner's information, since information is affiliated. Now in an English auction with common value elements, i.e., the interdependent values model, the price depends on all other bidders' information, while in a Vickrey auction, the price depends on only one other bidder's information, so it follows that the English auction generates higher average prices than the Vickrey auction. Similarly, when bidders have affiliated values and are risk-neutral, the Vickrey auction generates higher average prices than the Dutch and first-price sealed-bid auctions, which ignore all information except for the value of the winning bid. Klemperer points out that this general principle will imply the following practical rule of thumb: If the seller has access to any private source of information, it would be in his interest to reveal it honestly, as it is likely to result in higher revenue to him.

This general principle, that expected revenue is raised by linking the winner's payment to information that is affiliated with the winner's information, is sufficiently important to have a name:

The Linkage Principle (Milgrom and Weber 1982)
On average, a seller in an auction will enhance his revenue by providing the bidders with as much information as possible about the value of the item.

One consequence of the Linkage Principle, pointed out by Klemperer (2004a, Chapter 1) , is that if the winner's value can be observed (even imperfectly) after the conclusion of the auction, then the seller can earn more revenue by making the winner's payment depend on this observation.

The Linkage Principle, useful as it is, does come with one important caveat: It is only guaranteed to hold in the case of a single item; there are counterexamples for the case of multiple units. See Perry and Reny (1999) for a full discussion.

27.3.3.2 Optimality and information structure

Most of the literature on mechanism[14] design maintains the assumption of Myerson (1981) that the information held by market participants is exogenous. However, there are many auctions in which the precision of the information available to the buyers is at least partially controlled by the seller. For example, in U.S. offshore "wildcat" oil tract auctions[15] the firms involved in bidding are permitted to gather information about the lease value and their drilling costs prior to the sale using seismic information, but no on-site drilling is allowed. In contrast, in US offshore "drainage" oil leases, some bidders are intentionally given access to superior information by allowing them prior drilling in the area. It would be helpful to know what information structure maximize the seller's revenues.

This question was taken up by Bergemann and Pesendorfer (2007), who consider the joint decision problem of a seller who wishes to sell an object to one of multiple bidders with private valuations, where the seller can decide (1) the accuracy by which bidders learn their valuation, as well as (2) to whom to sell at what price. Bergemann and Pesendorfer show the existence of an optimal information structure, characterize properties of optimal information structures, and illustrate that the case of Myerson (1981) emerges as a special case when the seller informs the

[14] Auctions comprise a subset of a class of objects called *mechanisms*, a topic that is beyond the scope of this chapter. Presentations of mechanisms and mechanism design can be found in Krishna (2010) and Milgrom (2004).
[15] These auctions were also discussed in the paper of Capen et al. (1971) considered in Section 27.1, above.

bidders perfectly. Bergemann and Pesendorfer provide other examples in which the seller can control the precision of the information available to the buyers: wholesale used car auctions, licensing for motion pictures, and competition of brokers for the trade of a large portfolio on behalf of an institutional asset manager.

27.4 THE SIMULTANEOUS ASCENDING AUCTION

A new form of mechanism for auctioning *multiple items* appeared on the scene in 1994, the *simultaneous ascending auction* (SAA). This auction was designed for use by the U.S. Federal Communications Commission to allocate licenses for the right to use bands of the electromagnetic spectrum. The simultaneous ascending auction is unique among auction mechanisms in that it was conceived for a very specific application and developed over a short period of time. Paul Milgrom and Robert Wilson, in consultation with John McMillan and Preston McAfee, developed and proposed this auction procedure, quite literally, over a period of a few weeks (Milgrom 2004, p. xi). This is all the more extraordinary in that the SAA has proved to be the most influential new auction design of the past century.

Cramton (2006) points out that auctions have become the preferred method of assigning spectrum—not only in the United States, but also in Europe and around the world. Since 1994 there have been over 80 *spectrum auctions* in the U.S. alone, mainly for Personal Communications Services (PCS) and Advanced Wireless Services (AWS), most of which have made use of the simultaneous ascending auction. Klemperer (2004a, Chapter 5) reports that the early auctions in Europe for third generation (3G) mobile wireless licenses (Austria, Germany, Italy, the Netherlands, Switzerland, and the UK in 2000; Belgium, Denmark, and Greece in 2001) raised in total almost $100 billion (or over 1.5 percent of GDP). Hazlett (2008) reports that the U.S. auctions alone—this is prior to the 700 MHz auction in mid-2008—raised in excess of $25 billion. The 700 MHz auction raised over $19 billion. Cramton (2006) observes that the SAA has been refined and extended to the sale of divisible goods in electricity, gas, and environmental markets.[16]

[16] See Wilson (1979) for the classic paper on auctions of divisible goods.

27.4.1 How the SAA works

The SAA proceeds by discrete rounds, where multiple items are auctioned simultaneously, each with its own price. Each bidder is free to bid in a round on any number of items, subject to: (i) a minimum bid increment; (ii) an *activity rule* that restricts the pace of the bidding; and (iii) an allowance for *bid withdrawal*, whereby players can withdraw their bids subject to a payment equal to the difference between the withdrawn bid and the bid that replaces it. The auction terminates when a round completes in which no new bids have been submitted on any item. The winner on each item is the high bidder on that item, who is required to pay his bid.

The minimum bid increment is typically specified as a percentage of the current price, and is subject to change throughout the auction. Minimum bid increments assure that the auction concludes in a reasonable amount of time.

The activity rule determines the bidder's current *eligibility*, which is the maximum quantity of items on which he may bid. Thus, the activity rule requires bidders to maintain a minimum level of bidding activity during the course of the auction where, as the auction progresses, the level of activity increases. The activity rule forces a bidder desiring a large quantity of items to bid for a relatively large quantity earlier in the auction, thus preventing against the "snake in the grass" strategy, in which a bidder maintains a low level of activity early in the auction and then greatly expands his demand late in the auction. The activity rule also promotes price discovery.

The bid withdrawal rule[17] facilitates the realization of bidder *synergies*, the situation in which a bidder has a higher valuation for a particular set of items than the sum of their values to him individually, i.e., the items in the set are *complements*. Thus, the bid withdrawal rule mitigates what is known as the *exposure problem*, in which a bidder is exposed to a possible loss by bidding on a set of items where his bid price accounts for synergistic gains that he might not achieve.

The simultaneous ascending auction can be seen to be a generalization of the silent auction, albeit with an electronic bid sheet.

27.4.2 Development of the spectrum auctions

In 1985, a report entitled "Using Auctions to Select FCC Licensees" was issued by the Federal Communications Commission's Office of Plans and Policy (Kwerel and Felker 1985). Although the basic idea had been around for some time—in 1959 Ronald Coase famously proposed a general regime of spectrum property rights (Coase 1959)— the Kwerel-Felker paper can be credited as being the key document that successfully advocated the utilization of auctions for the allocation of

[17] The bid withdrawal rule was proposed by Preston McAfee.

spectrum licenses. Prior to this date, it had been generally accepted that there were only two ways to assign radio and television licenses.

The first method was *comparative hearings*. Depending on your point of view, comparative hearings meant that licenses were: (i) assigned by the statutory standard of "public interest, convenience or necessity" (Communications Act of 1934), or (ii) "simply handed to politically preferred parties" (Hazlett 2008). In any event, comparative hearings suffered from being very slow, as well as being wasteful of resources and lacking transparency. The method was cumbersome and eventually led to a large backlog of unassigned licenses.

In the early 1980s, these drawbacks led to the replacement of comparative hearings by a system that could work quickly: *lotteries*. As Milgrom (2004) explains, the lotteries did speed up the license approval process, but the fact that lottery winners were permitted to resell their licenses meant that a huge number of new applicants—many of whom were speculators with no interest in the telephone business—were randomly rewarded with windfalls in the form of licenses worth millions of dollars. Huge amounts of resources were again wasted, and the resulting chaos delayed the introduction of nationwide mobile telephone services in the United States. A new method to allocate licenses was needed desperately.

Finally, eight years after the appearance of the Kwerel-Felker paper, Congress passed the 1993 Omnibus Budget Reconciliation Act, which gave the FCC authority to use competitive bidding to choose from among two or more mutually exclusive applications for an initial license. Kwerel (2004) reports in the foreword to Milgrom's book that one of the first auction design issues that the FCC considered was whether to use an ascending- or sealed-bid mechanism. If precedent were the guide, a sealed-bid design would have certainly won the day, since the Federal government already made use of simple sealed-bid auctions, especially for offshore oil and gas leases. Kwerel explains that the FCC chose the ascending bid mechanism because they believed that providing bidders with more information would likely increase efficiency. Complete details of the development of the auction are provided in Kwerel's foreword.

27.4.2.1 An historical note

Although it is generally accepted that the concept of auctions for radio and television licenses was first thought of by a law student named Leo Herzel who proposed it in the *University of Chicago Law Review* in 1951 (see, e.g., Hazlett 2008), the idea had in fact appeared the previous year in the *American Journal of Economics and Sociology* in a short note by Will Lissner, the journal's editor (Lissner 1950). Lissner discusses the decision by the government of South Vietnam to license gaming (i.e., gambling) houses: "As a matter of social policy the licenses had to be limited. Who was to get the privilege? Persons favored by the politicians? No, an auction was held on Nov. 27, 1948 at which the privilege of opening gaming houses in the Saigon-Cholon region went to the highest bidders." Lissner ends with an admonishment regarding a missed opportunity for another type of government auction closer to home, describing the status quo as follows: "Radio and television channel licensees in the United States enjoy as their private income the unearned income produced by their privilege."

27.4.3 The clock auction

The *clock auction* is an ingeniously simple variant on the SAA. Each item has its own associated upward-ticking clock, which is controlled by the auctioneer and indicates the current price for the item, where the clock price for that item applies to all the available units. The clocks are started at low prices, that is, the reserve prices. Each round, bidders are given a fixed amount of time to submit their bids for items they would like to purchase at the current clock prices. Where there is excess demand for an item, the price for that item ticks upward. A new round is started, and the auctioneer requests new bids. When demand equals supply on all items, the auction ends. A significant positive feature of the clock auction is that it obviously precludes the possibility of jump bids, the occurrence of which can forestall or signal competition.

There are two main variants of clock auction. The first, called the "combinatorial clock auction," was developed during 1999 by David Porter, Stephen Rassenti, Anil Roopnarine, and Vernon L. Smith (2003). At the time, they were testing versions of the Federal Communication Commission's "simultaneous multiple round (SMR) auction"—the original term for the SAA—on hundreds of auction participants. Porter and his colleagues had four objectives in mind: efficiency in achieving all gains from exchange, task simplicity for the bidders, efficacy in handling complexity in the allocation problem, and computational feasibility. An important component of the combinatorial clock auction is that a final phase is required when, after a particular clock price increases, e.g., from $90 to $100, the demand for that item at the higher price, i.e., $100, becomes less than is available. In such cases where there is excess supply for at least one item, and demand exactly equals supply for all the other items, the auctioneer must solve an integer programming problem to find the allocation of items that maximizes his revenue. This integer programming problem is called the "winner determination problem" and is discussed in more detail in Section 27.5.2 below.

The second variant of clock auction, called the "dynamic clock auction," is due to Ausubel and Cramton (2004). The key difference here is the allowance of *intraround bids*, in which bidders express their demands in each round at all possible combinations of prices between the start-of-round price to the end-of-round price. Thus, in the combinatorial clock auction of Porter et al. (2003) where the price increased from $90 to $100 in a round, a bidder was only able to express the quantity he desired at $90 and at $100. However, in the dynamic clock auction of Ausubel and Cramton, a bidder expresses his desired quantity at *all* prices between $90 and $100. This feature eliminates the need to solve an integer programming problem at the end of the auction.

Although Porter, Rassenti, Roopnarine, and Smith slightly anticipated Ausubel and Cramton in their development of a clock auction, the latter authors admit that the idea goes all the way back to Walras (1874, Lesson 12), who introduced a theoretical process to study price adjustments in a market involving a number of different commodities. A fictitious auctioneer announces a set of prices for each of the commodities, bidders respond by reporting the quantity of each item they wish to purchase at these prices, and the auctioneer increases or decreases the price on each item according to whether the excess demand is positive or negative. This iterative process, called *tâtonnement* (literally, "groping") continues until a set of prices is reached at which excess demand is zero, and trade occurs only at the final set of prices. Of course, the clock auction is not identical with Walrasian tâtonnement. In both types of clock auction, the prices can only rise; in Walrasian tâtonnement, the prices can both rise and fall, and in general there can be no guarantee of convergence, as the procedure may cycle indefinitely.

When should the clock auction be employed rather than the SAA? Milgrom (2004, Section 7.2) explains that when there are a few homogeneous classes of items, each with many goods, the clock can run much faster than the standard SAA; further, it leads to the same near-competitive outcomes with *straightforward bidding*, i.e., the strategy in which a bidder bids the minimum amount in each round so as to maximize his surplus under the current prices. The earliest practical use of the clock auction is the Electricité de France (EDF) power plant auctions, which employed a dynamic clock auction (Ausubel and Cramton 2004). EDF's use of clock auctions began in 2001, and they have been successfully used by the company in 42 quarterly auctions, selling in total some 10 billion euro of electricity contracts as of December 2011 (Ausubel 2012).

27.4.4 Further reading

The spectrum auctions are ongoing, and the definitive history of these auctions is yet to be written. However, the reader wishing to acquire a good overview with a minimum of technical detail is directed to the following program of readings: McMillan (1994), McMillan (1995), McAfee and McMillan (1996), Cramton (1997), Cramton (2006), and Hazlett (2008). Klemperer (2004a, Chapters 5 and 6) discusses the "third generation'' (3G) mobile telecommunication (UMTS) auctions held in 2000 and 2001, where Chapter 5 focuses on the European auctions overall, and

Chapter 6 presents a first-hand account of designing the British 3G auction. For a technical presentation of the simultaneous ascending auction in the context of the spectrum auctions, see the book by Milgrom (2004).

27.5 COMBINATORIAL AUCTIONS

…………………………………………………………………………………………………...

A *combinatorial auction* is an auction in which bidders can place bids on combinations of items, called *packages*, rather than just on individual items. Combinatorial auctions can also include the possibility of *Boolean bids*, package bids joined up by the Boolean connectives: AND, OR, and NOT. Note that the combinatorial clock auction, despite its name, is not a true combinatorial auction, as it does not allow package bids.

Although combinatorial auctions have been discussed in the literature for almost thirty years, most expositions of auctions by economists have little or no mention of them. The reason for this is clear: Combinatorial auctions are cross-disciplinary, requiring an understanding not only of economics but of combinatorial optimization as well. In addition, economists tend to seek equilibrium results, and finding an equilibrium in a combinatorial auction is a daunting task.[18]

Combinatorial auctions provide fertile ground for future research in, and applications of, auction pricing. One important caveat from Milgrom (2004, p. xiii): "Unlike auctions for a single object, in which efficiency and revenue objectives are usually at least roughly aligned, multi-item auctions can involve radical trade-offs between these two objectives."

27.5.1 Airport time slots

Just as the work of William Vickrey (1961) is widely accepted as the foundation paper in auctions, a strong claim can be made for the work of Stephen Rassenti, Vernon L. Smith, and Robert L. Bulfin (1982), "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," as being the foundation paper in combinatorial auctions. Smith and his co-authors considered a topic that was, and is, real and urgent: the allocation of airport runway slots at congested airports.

It is certainly obvious that, when an airline requires a take-off slot for a flight at the originating airport, it also requires a landing slot at the terminating airport and, in cases where there are interconnecting legs, there will also be demands for take-off and landing slots at the intermediate airports. Smith et al. addressed the problem of designing a combinatorial sealed-bid auction to

[18] I am grateful to Itai Sher for this additional point.

serve as the primary market for allocating airport slots for which individual airlines would submit package bids.

Smith et al. first formulated the auctioneer's problem as an integer programming problem, which they recognized as a variant of the *set-packing problem*. They next provided an algorithm that yields an allocation to packages that maximizes efficiency and determines individual slot resource prices, which are then used to price packages to winning bidders not exceeding the amount they bid. Finally, the authors conducted a series of experiments where students were paid according to how well they did in the auction.

As pointed out in the introductory chapter of Cramton, Shoham, and Steinberg (2006), the paper of Rassenti, Smith, and Bulfin is significant not only for being the first on combinatorial auctions, but also for introducing many of the key ideas in the field. These include the mathematical programming formulation of the winner determination problem, the connection between the winner determination problem and the set packing problem, and the related issue of computational complexity. The paper described Boolean bids. It made use of techniques from experimental economics for testing combinatorial auctions. It raised the issue of incentive compatibility in combinatorial auctions. Even the very term "combinatorial auction" was introduced in this paper, as was "smart" exchange or market, a now-standard term. For the intriguing story of the origins of combinatorial auctions arising from airline deregulation in the U.S., see Smith (2006). For more on auctions for airspace system resources, see Ball, Donohue, and Hoffman (2006).

27.5.2 The winner determination problem

The idea of a combinatorial auction is simple enough, viz., to allow package bids in addition to individual bids. However, one rather significant problem arises, the notorious *Winner Determination Problem* (*WDP*). This is the auctioneer's problem of labeling bids as either winning or losing so as to maximize the sum of accepted bids, under the constraint that each item can be allocated to at most one bidder. This is a computationally intractable problem, since the WDP is equivalent to the *weighted set packing problem* in combinatorial optimization, and thus is *NP-hard*. What this means in practice is that, for realistically-sized problems, the computational burden can be—astonishingly enough—beyond the capability of any existing computer.[19] For more details on the Winner Determination Problem, see Lehmann, Müller, and Sandholm (2006).

---

[19] In some restricted cases the problem becomes tractable, but the restrictions required are invariably draconian. See Rothkopf, Pekec, and Harstad (1998).

Another difficulty that arises in combinatorial auctions is known as the *threshold problem*. This is where the allowance for package bidding can favour bidders seeking larger packages, since bidders on smaller packages may not have the resources individually to overtake a large package bid, or may not have the ability to coordinate with each other in order to do so.

27.5.3 Combinatorial auctions in practice

Combinatorial auctions have been proposed for assigning universal service support for competing telephone companies (Kelly and Steinberg 2000). They have been used for truckload transportation in the United States (Caplice and Sheffi 2006), bus routes in London (Cantillon and Pesendorfer 2006), school milk programs in Chile (Epstein et al. 2002), and industrial procurement worldwide (Bichler et al. 2006). They were used in the allocation of spectrum in the United States in 2008 and in Britain in 2009.

Most of this probably would not have happened if it were not for the paper of Rassenti, Smith, and Bulfin (1982). But again, practice preceded theory. Twenty-seven years earlier, auction firm executive Louis McLean (1955) in his article, "Auction Anecdotes," retailed a number of stories about auctions involving lawyers, where three of his stories describe bankruptcy auctions allowing for *entirety bids*. An entirety bid is a package bid on all the items in the auction, where the highest entirety bid wins only if it exceeds the sum of the bids on the individual items. Auctions incorporating entirety bidding are still in common use today, especially for bankruptcy and real estate sales.

Below, we discuss three practical combinatorial auction designs.

27.5.4 Practical combinatorial auction designs

27.5.4.1 The ascending proxy auction

In an *ascending proxy auction*, each bidder reports his preferences for packages or contracts to an electronic *proxy agent* that subsequently bids on the bidder's behalf. Preferences for packages or contracts can be much more than reservation prices. A contract could specify, for example, price, quality, and closing date.

The ascending proxy auction works as follows. In each round, if a given bidder is not among the provisional winners, the proxy agent submits the bid that the bidder most prefers according to his reported preferences. The auctioneer then considers all bids from the current and past rounds and selects his most preferred feasible collection of bids according to his objective, under the restriction that accepted bids can include at most one bid from each bidder. The auctioneer's selected bids become the new provisional allocation, the associated bidders are designated provisional winners, and the process repeats until no new bids are submitted. For more on ascending proxy auctions, see the paper of Ausubel and Milgrom (2006b), which introduced it.

27.5.4.2 The clock–proxy auction

The *clock-proxy auction* is, as the name implies, a hybrid between the clock auction and the ascending proxy auction; it was proposed by Ausubel, Cramton, and Milgrom (2006). The auction operates in two phases, where the first phase is a clock auction, in which the bidders directly submit bids, and the second phase is a proxy round, where bidders have a single opportunity to input proxy values, which is then run as a proxy auction. The conclusion of the proxy phase concludes the auction.

In the clock-proxy auction, bids are kept active throughout the auction, i.e., no bid withdrawals are permitted. Specifically, bids from the clock phase are also treated as package bids in the proxy phase. All bids are treated as mutually exclusive, i.e., as XOR bids. There are activity rules within the clock phase, and between the clock and proxy phases.

What are the advantages of the clock-proxy auction? The clock phase is simple for bidders and provides for price discovery. The proxy phase facilitates efficient allocations and competitive revenues, as well as reducing opportunities for collusion. The clock-proxy auction design has been tested successfully in the field; see Ausubel and Cramton (2004) for more details. It has been further developed and adopted for spectrum auctions in the United Kingdom.

27.5.4.3 PAUSE: Progressive Adaptive User Selection Environment

Is it possible to design a combinatorial auction mechanism that permits all package bids, yet is computationally tractable for the auctioneer? As demonstrated by Kelly and Steinberg (2000), the answer is yes. By transferring the computational burden of evaluating a package bid to the bidder submitting the bid, the auctioneer no longer faces the Winner Determination Problem. Although in theory a bidder might face a computationally intractable problem, in practice the bidder may have, for some of his bids, a relatively easy problem, and a basic principle in auctions has been that the task of finding an appropriate bid is the responsibility of the bidder. The Kelly-Steinberg design is called PAUSE (Progressive Adaptive User Selection Environment).

PAUSE proceeds in stages. In stage 1, a simultaneous ascending auction is held for all the items, thus facilitating price discovery. After stage 1, bidders can realize their synergies via package bidding. However, the bids on packages cannot be submitted in isolation: each bidder is required to submit them as part of a *composite bid*, which is a set of non-overlapping package bids (including possibly individual bids) that cover *all the items in the auction*. Of course, a bidder will generally be interested in bidding only on a subset of the items in the auction—and in any given round, perhaps only a subset of these. A composite bid consists of the bidder's own bids, together with previously-submitted bid—including composite bids—by any of the bidders.

The following example should make this clear (see Figure 1). There are six items in the auction: α, β, γ, α′, β′, γ′ (Figure 1a). Stage 1 ended with a bid of 5 on each item, respectively, by bidders A, B, C, D, E, F, and consequently a revenue to the auctioneer from these six bids totalling 30 (Figure 1b).

In the current round, there are standing bids of 11 by bidder A on the package αα′, 20 by bidder B on package ββ′, and 14 by bidder C on package γγ′, with revenue to the auctioneer from these three bids totalling 45 (Figure 1c).

Now, two composite bids are submitted simultaneously from Bidders A and B. Bidder A has a high valuation for the package αβγ. His composite bid consists of a bid from himself of 35 on the package αβγ, together with the earlier bids of 5 each on α′, β′, and γ′, respectively, from bidders D, E, and F, with revenue to the auctioneer of 50 (Figure 1d). Bidder B's composite bid (not shown) consists of a bid from himself of 35 on the package αα′ββ′, together with the earlier package bid of 14 on γγ′ from bidder C, with revenue to the auctioneer of 49. Thus, the auctioneer chooses bidder A's composite bid. The auction progresses from there, with bidders submitting composite bids of increasingly higher revenue until the auction terminates.

α	β	γ
α'	β'	γ'

(A) The six items in the auction.

5_A	5_B	5_C
5_D	5_E	5_F

(b) End of Stage 1. Auctioneer revenue: 30.

11_A	20_B	14_C

(c) Start of current round. Auctioneer revenue: 45.

35_A		
5_D	5_E	5_F

(d) Bidder A's composite bid. Auctioneer revenue: 50.

Figure 1. Illustration of composite bidding

Composite bidding has three important consequences: (1) the auctioneer is relieved of the computational burden of the winner determination problem (since he needs only choose the highest valid composite bid); (2) each losing bidder can compare his bids with the winning composite bid to see why he lost; and (3) at the conclusion of the auction, no bidder—winning or losing—would prefer to exchange his allocation with that of another bidder. These features are called, respectively: (1) *computational tractability*, (2) *transparency*, and (3) *envy-freeness*.

Kelly and Steinberg introduced the PAUSE mechanism for a specific application in the U.S. arising from the *Telecommunications Act of 1996*, one of whose requirements was that regulators consider ways to reform the method of providing *universal service* subsidies for high-cost areas. This refers to the situation that had been in place in the United States for many years, in which telephone companies were granted a monopoly to provide telephone service within their operating region, but had a concomitant responsibility to provide basic telephone service to everyone, no matter how costly. This was mitigated by a provision in which the telephone companies would receive subsidies for designated "high cost areas." Around the time of the passage of the Act, several parties advocated that "competitive bidding"—auctions—be used to determine universal service subsidies.

Kelly and Steinberg's paper was written in response to this suggestion. Their universal service support auction was designed as a reverse auction, using the PAUSE mechanism, where firms would bid for subsides on specified areas, and the winning firm on an area would be the one that bid the lowest subsidy. Since a firm might find it less costly to serve an area if it were to serve it together with other areas, a combinatorial auction was required. Kelly and Steinberg's design also allowed for competition within areas, i.e., "multiple winners," a strong interest of the FCC at the time. Kelly and Steinberg 's paper focused primarily on the auctioneer. Land, Powell, and Steinberg (2006) take the next step by examining bidder behavior under the mechanism.

27.5.5 The VCG mechanism

The most famous combinatorial auction is the *Vickrey-Clarke-Groves* (*VCG*) *mechanism*, which is a natural generalization of the Vickrey auction to multiple items. The VCG, like the Vickrey auction itself, requires winners to pay the opportunity cost of their participation. It works as follows. Bidders report their valuations for every possible package of items to the auctioneer, who then determines which items are to be allocated to which bidders by solving the problem of maximizing total payments. However, each bidder pays not his bid price but rather the incremental value that would be derived by assigning the items allocated to him according to the items' next best use among the other bidders. Under the VCG mechanism, as in the Vickrey auction, it is a dominant strategy for the bidder to truthfully report his values. For a lucid presentation of the VCG formula, see Ausubel and Milgrom (2006a).

As the name indicates, the Vickrey-Clarke-Groves mechanism evolved from three sources: William Vickrey's (1961) famous paper on auctions, Edward H. Clarke's (1971) work on the pricing of public goods, and Theodore Groves's (1973) contribution to the theory of teams. These three papers were written independently—there are no cross-citations among them—and it is clear that neither Clarke nor Groves had an auction *per se* in mind. Clarke's idea was to propose a solution to the "revealed preference problem" for public goods, a situation in which individuals are induced to hide or understate their true preferences in order to improve their individual welfare while foregoing jointly available potential gains. Groves studied the problem of inducing members of an organization to behave as if they formed a team, where the team head's incentive problem is to choose a set of employee compensation rules that will induce his sub-unit managers to communicate accurate information and take optimal decisions. These two methods merged to become the "Clarke-Groves demand-revealing mechanism" for public goods, which later became the “Vickrey-Clarke-Groves mechanism” or, more simply, the “VCG mechanism.” Often, the VCG mechanism is itself referred to as the Vickrey auction.

27.5.5.1 Weaknesses of the VCG mechanism

The VCG mechanism has some impressive theoretical strengths. However, its list of weaknesses is distressingly long. In his afterward to his survey of auction theory, Klemperer (2004a) emphasizes that the Vickrey auction is usually impractical even in those private-value contexts in which it is (in theory) efficient. Here is Klemperer's litany of VCG woes:

> Policy makers usually find a Vickrey auction very hard to understand and operate; it often results in bidders with high values paying less for objects than bidders who win identical objects but have lower values for them (which seems strange and unfair to many people); it offers unusual opportunities for collusive behavior which are also hard to guard against; and it sometimes yields low revenues. Furthermore, it is not efficient (and may perform very badly) if bidders are risk-averse or have budget constraints or have common-value elements to their valuations.

That the VCG sometimes yields low revenues is a problem of sufficient concern that it has a name: *revenue deficiency*. These low revenues can, in fact, be as low as zero!

As Ausubel and Milgrom (2006a) explain in their felicitously-titled work, “The Lovely but Lonely Vickrey Auction,” its weakness go a long way to explaining why the VCG mechanism, “so lovely in theory,” is “so lonely in practice.” On their own list they include the auction's vulnerability to the use of multiple identities by a single bidder. This problem can be described as “shill bids by a bidder,” in contrast to “shill bids by the auctioneer” (discussed above in Section 27.3.2). More formally, this ruse is known as *false name bidding* or *pseudonymous bidding* (Yokoo 2006). Yokoo states that, while many auction researchers have discussed problems arising from collusion, compared with collusion a pseudonymous bid is easier to execute on the internet, since getting another identifier such as an email address is cheap. But Day and Milgrom (2008) point out that the problem is broader than just anonymous internet auctions. In the U.S. spectrum auctions several of the largest corporate bidders—including AT&T, Cingular, T-Mobile, Sprint, and Leap Wireless—had at times either contracts with, or financial interests in, multiple bidders bidding in the same auction. As Day and Milgrom explain, this allowed for strategies that would not be possible for a single, unified bidder.

Rothkopf (2007) includes among his "Thirteen Reasons Why the Vickrey-Clarke-Groves Process is Not Practical" two additional weakness not mentioned above: the possibility of alternative equilibria, and problems associated with the disclosure of valuable confidential information.

We next show that, by compromising the loveliness of the Vickrey auction, it became lonely no more.

27.5.5.2 The generalized second-price auction

The world's most frequently employed auction is unknown to the overwhelming majority of people who open the auction each day. This probably includes you. Each time you enter a search term into Google, you initiate an auction among a subset of Google's advertisers. The underlying auction mechanism used by Google—and by Yahoo and by most other search engines—is what Edelman, Ostrovsky, Schwarz (2007) refer to as the *generalized second-price* (*GSP*) *auction*. According to Google 's June 2009 quarterly report filed with the U.S. Securities and Exchange Commission, over 97 percent of the company's $11.03 billion revenue for the 6-month period ending June 2009 came from advertising, that is, the GSP auction.

Your Google search term generates a page of links most relevant to the search term. No auction is involved here. However, in addition, the right side of your screen displays a list of paid advertisements, called *sponsored links*, i.e., web links to advertisers who wish to target ads to you as a consequence of you entering that key word. There are a limited number of positions for sponsored links on the web page. Now, if you are sufficiently interested in a sponsored link that you decide to click on it, this will have two immediate effects: (i) you will be sent to the advertiser's web page, and (ii) the advertiser will be charged a fee by Google, viz., the advertiser's individual *price-per-click*. The assumption is that, the higher up on the list a sponsored link appears, the more likely you are to click on it. Some type of *sponsored search auction* is employed to determine the allocation of the ad positions to advertisers and their individualized price-per-click bid price, conditional on your key word. As mentioned above, the mechanism used most often is the generalized second-price auction.

The GSP auction works essentially as follows. For a specific keyword, advertisers submit bids stating their maximum willingness to pay each time a user clicks on their sponsored link. The advertiser who bids highest is allocated the first sponsored link position, but his price-per-click is set at the bid price of the second-highest bidder. The advertiser who bid the second-highest price-per-click is allocated the second sponsored link position and pays the third-highest bid price as his price-per-click, and so forth.[20]

The GSP auction obviously generalizes the second-price auction in the sense that if there were only one sponsored link position, then the GSP auction would coincide with the original Vickrey auction. Further, the GSP auction has the feature that a bidder's payment does not directly depend on his own bid. However, as shown by Edelman et al., the GSP auction is (not surprisingly) not equivalent to the VCG mechanism. In particular, unlike the VCG, the GSP generally does not have an equilibrium in dominant strategies, and truth-telling is not an equilibrium of GSP. Varian (2007) provides empirical evidence, based on a random sample of 2425 auctions involving at least five ads each on a particular day, that the equilibria of the GSP auction describe reasonably accurately the properties observed in Google's sponsored search auction.

27.6 SUMMARY AND FUTURE DIRECTIONS

The starting point of auction theory is Vickrey's seminal work, the significance of which has less to do with proposing the Vickrey auction than with founding a field that has now matured to the point where it can explain the limitations of the Vickrey auction. In summary, what basic lessons does that field provide to someone who wishes to engage in auction pricing?

First, know your objective. If it is efficiency, then go ahead and use the Vickrey auction, but be well aware of its weaknesses, too. If your objective is optimality, keep in mind that a slow Dutch auction will likely earn you more revenue than a first-price sealed-bid auction. Whatever your objective, always set a reserve price. If you choose to use a sealed-bid auction, don't even consider a 3rd-price (or 4th- or 5th- etc. price) auction. If you are selling a single item and you want to maximize revenue, reveal any private information you may have, as this will likely yield you more. And if the winner's value can be observed to any extent after the auction's close, than you will probably do better still by making the payment dependent on this observation. In selling multiple items, you may wish to use the simultaneous ascending auction, but when there are only a few kinds of items, each with many units available, an ascending clock auction would probably run faster. If you suspect that some bidders might have significant synergies for at least some of the items, then a combinatorial auction would likely be best of all.

[20] This description is in fact a simplification. Google, as well as most other search engines, currently multiply bids by "quality scores," which are often closely related to how good/clickable, the ad is. More detail is provided by Edelman et al. (2007, p. 257) and Varian (2007).

What are the important future directions in auctions? I suggest three:

- *Combinatorial auctions.* The future of auctions is combinatorial auctions. This is due to a convergence of two factors. First, as auctions have increasingly been put into practice, the limits of the usefulness of standard (i.e., non-combinatorial) auctions has been reached for most purposes, since the existence of bidder synergies is so common in many cases. Second, the internet has made it no longer necessary for the bidders to be assembled in a single location, and at the same time makes the logistics of bid submission relatively easy, even with a large number of items and many bidders. Such large auctions are more likely to involve synergies. However, it is the appropriate choice of combinatorial auction that is the key question here. Testing and evaluation of combinatorial auction procedures is what is now needed.
- *Tie bids.* There is a dirty little word in the world of auctions, and that word is *tie*. Vickrey casually dismissed the issue: "In the case of tie bids we can assume the tie to be broken by a random drawing giving each tied player the same probability of winning." Alas, this obvious procedure, now called *the standard tie-breaking rule*, is not always sufficient to ensure an equilibrium. Under interdependent values, *special tie-breaking rules* might be required. For example, tied bidders might be asked to bid in a second-price auction. In practice, tie-breaking has often been effected via *time stamps*, i.e., where preference is given to a bid submitted earlier. Time stamps had been used in the U.S. spectrum auctions, but the FCC reverted to the standard tie-breaking rule after observing the alarming practice of bidders rushing to submit their bids in an effort to win ties. Other tie-breaking rules can be more surprising. At the Tsukiji Fish Market in Tokyo, ties are often broken with a quick round of rock-paper-scissors. In Florida, current statutes require that, in any state procurement auction, ties are broken by giving preference to a business that certifies that it has implemented a drug-free workplace program. Well worthwhile would be further study regarding appropriate tie-breaking rules in theory and in practice.
- *Institutional aspects of auctions.* As we have seen, auctions currently have enormous impact in the pricing of fine art (English), spectrum (simultaneous ascending, combinatorial), internet advertising (generalized second-price), Treasury securities (multiple-unit sealed-bid), and perishable products such as flowers and fish (Dutch). They are a key tool in industrial procurement (English, combinatorial) and government contracting (first-price sealed bid, other sealed-bid). Of course, auctions play a highly significant role in many other areas of commerce. It is this author's hope that some reader of this chapter will embark on an updated study of the institutional aspects of auction procedures and processes, as it would be of considerable value to

both the academic and business communities to have a 21st century *tour d'horizon* of what Cassady called “this fascinating method of selling and price making.”

For further reading on auctions, three books comprise the gold standard: Krishna (2010), Klemperer (2004a), and Milgrom (2004). All three books provide the rigorous theory; however, the Klemperer and Milgrom volumes are ultimately aimed at applications, especially spectrum allocation, with Milgrom focusing on the American and Klemperer on the European auctions. Milgrom’s book also discusses combinatorial auctions, a topic covered thoroughly in the integrated multi-authored book edited by Cramton, Shoham, and Steinberg (2006).[21]

[21] The choice from among four major publishers for the book *Combinatorial Auctions* was determined by the book’s editors via auction, where bids were publisher offers to set the retail price of the book, and the publisher offering the lowest price would be selected the winner (MIT Press).

REFERENCES

Athey, S., Haile, P.A. (2006). "Empirical Models of Auctions," Chapter 1 (pp. 1-45) in: R. Blundell, W.K. Newey, and T. Persson (eds.), *Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress*, *Volume II*, Cambridge, UK: Cambridge University Press.

Ausubel, L.M. (2012). Private communication, February 6.

Ausubel, L.M., Cramton, P. (2004). "Auctioning Many Divisible Goods," *Journal of the European Economic Association* 2 (2-3), 480-493.

Ausubel, L.M., Cramton, P., Milgrom, P. (2006). "The Clock-Proxy Auction: A Practical Combinatorial Auction Design," Chapter 5 (pp. 115-138) in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Ausubel, L.M., Milgrom, P. (2006a). "The Lovely but Lonely Vickrey Auction," Chapter 1 (pp. 17-40) in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Ausubel, L.M., Milgrom, P. (2006b). "Ascending Proxy Auctions," Chapter 3 (pp. 79-98) in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Avery, C. (1998). "Strategic Jump Bidding in English Auctions," *Review of Economic Studies* 65 (2), 185-210.

Ball, M.O., Donohue, G.L., Hoffman, K. (2006). "Auctions for the Safe, Efficient, and Equitable Allocation of Airspace System Resources," Chapter 20 (pp. 507-538) in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Beggs, A., Graddy, K. (2009). "Anchoring Effects: Evidence from Art Auctions," *American Economic Review* 99 (3), 1027-1039.

Bergemann, D., Pesendorfer, M. (2007). "Information Structures in Optimal Auctions," *Journal of Economic Theory* 137 (1), 580-609.

Bichler, M., Davenport, A., Hohner, G., Kalagnanam, J. (2006). "Industrial Procurement Auctions," Chapter 23 (pp. 593-612) in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Bichler, M., Steinberg, R. (eds.) (2007). Special Issue on E-Auctions for Procurement Operations, *Production and Operations Management* 16 (4).

Cantillon, E., Pesendorfer, M. (2006). "Auctioning Bus Routes: The London Experience," Chapter 22 (pp. 573-592) in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Capen, E.C., Clapp, R.V., Campbell, W.M. (1971). "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology* 23, 641-653.

Caplice, C. , Sheffi ,Y. (2006). "Combinatorial Auctions for Truckload Transportation," Chapter 21 (pp. 539-572) in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Carare, O., Rothkopf, M. (2005). "Slow Dutch Auctions," *Management Science* 51 (3), 365-373.

Cassady, R., Jr. (1967). *Auctions and Auctioneers*, Berkeley and Los Angeles: University of California Press.

Charlton, M. (2008). *Oxford University Press*, private communication, July 15 and 21, 2008.

Clarke, E.H. (1971). "Multipart Pricing of Public Goods," *Public Choice* 11, 17-33.

Coase, R.H. (1959). "The Federal Communications Commission," *Journal of Law and Economics* 2, 1-40.

Cramton, P. (1997). "The FCC Spectrum Auctions: An Early Assessment," *Journal of Economics and Management Strategy* 6 (3), 431-495.

Cramton, P. (1998). "The Efficiency of the FCC Spectrum Auctions," *Journal of Law and Economics* 41 (2, Supplement), 727-736.

Cramton, P. (2006). "Simultaneous Ascending Auctions," Chapter 4 (pp. 99-114) in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Cramton, P., Shoham, Y., Steinberg, R. (eds.) (2006). *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Day, R., Milgrom, P. (2008). "Core-Selecting Package Auctions," *International Journal of Game Theory* 36 (3,4), 393-407.

Edelman, B., Ostrovsky, M., Schwarz, M. (2007). "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review* 97 (1), 242-259.

Epstein, R., Henriquez, L, Catalán, J., Weintraub, G.Y., Martinez, C. (2002). "A Combinatorial Auction Improves School Meals in Chile," *Interfaces* 32 (6), 1-14.

FloraHolland (2008). "Smooth integration after merger: FloraHolland satisfied with annual figures for 2007," Press release, Aalsmeer, The Netherlands, May.

Gallego, G., Phillips, R., Sahin, O. (2008). "Strategic Management of Distressed Inventory," *Production and Operations Management* 17 (4), 402-415.

Groves, T. (1973). "Incentives in Teams," *Econometrica* 41 (4), 617-631.

Hazlett, T.W. (2008). "Optimal Abolition of FCC Spectrum Allocation," *Journal of Economic Perspectives* 22 (1), 103-128.

Hendricks, K., Porter, R. (2007). "An Empirical Perspective on Auctions," Chapter 32 (pp. 2073-2143) in: M. Armstrong and R. Porter (eds.), *Handbook of Industrial Organization*, Volume 3, Amsterdam: Elsevier.

Herodotus (1914). *Histories*, *Book I* (*Clio*). (G.C Macaulay, translator.) London: Macmillan.

K (1830). "The Country Belle," *Virginia Literary Museum & Journal of Belles-Lettres, Arts, Sciences &c.* 1 (40) (March 17), 632-640.

Kagel, J.H., Levin, D. (2013). "Auctions: A Survey of Experimental Research, 1995-2008, in J. Kagel and A. Roth (eds), *Handbook of Experimental Economics, Vol. 2*. Princeton, NJ: Princeton University Press.

Kelly, F., Steinberg, R. (2000). "A Combinatorial Auction with Multiple Winners for Universal Service," *Management Science* 46 (4), 586-596.

Klemperer, P. (2004a). *Auctions: Theory and Practice*, Princeton: Princeton University Press.

Klemperer, P. (2004b). "An Early Example of 'Sniping' in an Auction," Lagniappe, *Journal of Political Economy* 112 (3).

Krishna, V. (2010). *Auction Theory*, Second Edition, San Diego, CA: Academic Press.

Kwerel, E. (2004). Foreword (pp. xv–xxii) to: Milgrom, P., *Putting Auction Theory to Work*, Cambridge: Cambridge University Press.

Kwerel, E., Felker, A. (1985). "Using Auctions to Select FCC Licensees," Office of Plans and Policy Working Paper No. 16, Federal Communications Commission, Washington, D.C., May.

Lambropoulos, S. (2007). "The Use of Time and Cost Utility for Construction Contract Award under European Union Legislation," *Building and Environment* 42 (1), 452-463.

Land, A. Powell, S., Steinberg, R. (2006). "PAUSE: A Computationally Tractable Combinatorial Auction," Chapter 6 (pp. 139-157) in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Lazear, E.P. (1986). "Retail Pricing and Clearance Sales," *American Economic Review* 76 (1), 14-32.

Learmount, B. (1985). *A History of the Auction*. Iver, UK: Barnard & Learmount.

Lehmann, D., Müller, R., Sandholm, T. (2006). "The Winner Determination Problem," Chapter 12 (pp. 297-317) in: P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Lissner, W. (1950). "Taxing Privilege in South Viet Nam," *American Journal of Economics and Sociology* 9 (4), p. 444.

Lucking-Reiley, D. (1999). "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet," *American Economic Review* 89 (5), 1063-1080.

Lucking-Reiley, D. (2000). "Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce," *Journal of Economic Perspectives* 14 (3), 183-192.

McAfee, R.P., McMillan, J. (1987). "Auctions and Bidding," *Journal of Economic Literature* 25 (2), 669-738.

McAfee, R.P., McMillan, J. (1996). "Analyzing the Airways Auction," *Journal of Economic Perspectives* 10 (1), 159-175.

McLean, L. (1955). "Auction Anecdotes," *The Shingle* (Philadelphia Bar Association) 18 (3), 65-70.

McMillan, J. (1994). "Selling Spectrum Rights," *Journal of Economic Perspectives* 8 (3), 145-162.

McMillan, J. (1995). "Why Auction the Spectrum?" *Telecommunications Policy* 19 (3), 191-199.

Milgrom, P.R. (2004). *Putting Auction Theory to Work*, Cambridge: Cambridge University Press.

Milgrom, P.R., Weber, R.J. (1982). "A Theory of Auctions and Competitive Bidding," *Econometrica* 50 (5), 1089-1122.

Mishra, D. Parkes, D.C. (2009). "Multi-Item Vickrey-Dutch auctions," *Games and Economic Behavior* 66 (1), 326-347.

Moldovanu, B., Tietzel, M. (1998). "Goethe's Second-Price Auction," *Journal of Political Economy* 106 (4), 854-859.

Mueller, M. (1993). "New Zealand's Revolution in Spectrum Management," *Information Economics and Policy* 5 (2), 159-177.

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

New Zealand Ministry of Economic Development (2005). "Spectrum Auction Design in New Zealand," Radio Spectrum Policy and Planning, Resources and Networks Branch, November.

Ockenfels, A., Reiley, D.H., Sadrieh, A. (2006). "Online Auctions," Chapter 12 (pp. 571-628) in: T.J. Hendershott (ed.), *Handbook of Information Systems and Economics,* vol. 1, Amsterdam: Elsevier Science.

Oren, M.E., Williams, A.C. (1975). "On Competitive Bidding," *Operations Research* 23 (6), 1072-1079.

Ortega-Reichert, A. (1968). "Models for Competitive Bidding under Uncertainty," PhD dissertation, Department of Operations Research, Stanford University, Stanford, CA.

Pepys, S. (1926). *The Diary of Samuel Pepys*. Edited by M. Bright and H.B. Wheatley. London: George Bell & Sons.

Perry, M., Reny, P.J. 1999. "On the Failure of the Linkage Principle," *Econometrica* 67 (4), 895-900.

Popkowski Leszczyc, P.T.L., Rothkopf, M.H. (2010). "Charitable Motives and Bidding in Charity Auctions," *Management Science* 56 (3), 399-413.

Porter, D., Rassenti, S., Roopnarine, A., Smith, V. (2003). "Combinatorial Auction Design," *Proceedings of the National Academy of Sciences* 100 (19), 11153-11157.

Rassenti, S.J., Smith, V.L., Bulfin, R.L. (1982). "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," *Bell Journal of Economics* 13 (2), 402-417.

Riley, J.G., Samuelson, W.F. (1981). "Optimal Auctions," *American Economic Review* 71 (3), 381-392.

Rothkopf, M.H. (2007). "Thirteen Reasons Why the Vickrey-Clarke-Groves Process is Not Practical," *Operations Research* 55 (2), 191-197.

Rothkopf, M.H., Pekec, A., Harstad, R.M. (1998). "Computationally Manageable Combinational Auctions," *Management Science* 44 (8), 1131-1147.

Royal Swedish Academy of Sciences (1996). Press release, Stockholm, Sweden, October 8.

Smith, V.L. (2006). Foreword (pp. xi–xv) to: P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

The Times (London) (1788). February 5, p. 4.

The Times (London) (1825). April 29, p. 4.

The Times (London) (1891). September 14, p. 12.

Tunca, T.I., Wu, Q. (2009). "Multiple Sourcing and Procurement Process Selection with Bidding Events," *Management Science* 55 (5), 763-780.

van den Berg, G.J., van Ours, J.C., Pradhan, M.P. (2001). "The Declining Price Anomaly in Dutch Dutch Rose Auctions," *American Economic Review* 91 (4) 1055-1062.

Varian, H.R. (2007). "Position Auctions," *International Journal of Industrial Organization* 25 (6), 1163-1178.

Vickrey, W. (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance* 16 (1), 8-37.

Walras, Leon (1874). *Eléments d'economie politique pure; ou, théorie de la richesse sociale.* (*Elements of Pure Economics; or, The Theory of Social Wealth*). Lausanne: Corbaz.

White House (1994). "Vice President Gore Opens First FCC Broadband Spectrum Auction," Office of the Vice President , press release, Washington, D.C., December 5.

Wilson, R.B. (1967). "Competitive Bidding with Asymmetric Information," *Management Science* 13 (11), 816-820.

Wilson, R. B. (1969). "Competitive Bidding with Disparate Information," *Management Science* 15 (7), 446-448.

Wilson, R.[B.] (1979). "Auctions of Shares," *Quarterly Journal of Economics* 93 (4), 675-689.

Yokoo, M. (2006). "Pseudonymous Bidding in Combinatorial Auctions," Chapter 7 (pp. 161-187) in P. Cramton, Y. Shoham, and R. Steinberg (eds.), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Zullo, R. (2006). "Public-Private Contracting and Political Reciprocity," *Political Research Quarterly* 59 (2), 273-281.