Auctions and Economic Design

13 May 2011

by Simon Loertscher and Tom Wilkening, University of Melbourne

Over the last three decades, auctions have become an increasingly popular means of allocating assets. On a regular basis auctions are now being used by governments both for procuring goods and services and for selling or leasing use rights for government owned resources including radio spectrum, taxi medallions**,** or gaming machines. The internet also has provided an abundance of new auction sites, with eBay constituting one of the largest trading platforms in the world.

The economic analysis of auctions, pioneered by Vickrey (1961, 1962), has witnessed an upsurge of interest in economics that began in the late 1970s and 1980s. By now, the economics of auctions is an integral part of a new field of economics that has become known as economic design (Roth 2002), and it is celebrated as one of the great successes of modern microeconomics (Maskin, 2005).

The use of auctions as a form of allocating government owned assets to private users has not always been the standard, except perhaps for tenders used for issuing treasury bonds and in procurement. Other allocation mechanisms, such as lotteries, or processes that allocated resources to final users based on vaguely defined criterions like merit and quality, were used extensively until about twenty years ago. In this article, we introduce the basic conceptual framework within which economists analyze auctions and other allocation mechanisms, and we explain why alternative procedures such as lotteries are not satisfactory. We then review the main results from the economics literature on auction design and explain how new auction rules are being developed for complex environments where multiple interrelated objects are being sold. The article concludes with a discussion of the broader issues that are considered critical for successful auction design.

I. The Mechanism Design Framework and a Comparison of Mechanisms

Before delving into auction theory itself, it is useful to briefly lay out the conceptual framework of auction theory and mechanism design. In this framework, there is a sharp distinction between the pieces of the economic system which are under the control of the designer -- the so called *mechanism* -- and the pieces of the economic system which the design cannot directly influence, which is typically called the *environment*.

- The mechanism can be thought of as the *rules of the game* which are used to govern the interaction of individuals who come together and exchange. In the auction setting, these rules include the way in which individuals can make bids, the way that those bids assign the auctioned goods to winners, and the payments which occur at the end of the auction.

- The environment, on the other hand, is composed of the economic agents and their potential needs and wants. These are the pieces of the design problem which cannot be directly controlled by the designer. Thus, they form a set of restrictions that the designer must account for. The environment includes the number of potential bidders, their potential valuations and preferences, and the information they have about the values of the other players.

Given an economic environment, the design problem is to find, or design, mechanisms whose *outcomes* achieve certain given goals such as Pareto efficiency or maximizing seller's revenue, taking into account that individuals will *act strategically* to best achieve their individual objectives. If a government designs the mechanism, then its objective may to maximize total welfare.[1]

It is useful to consider a simple example that illustrates the mechanism design framework and explains key differences between auctions and alternative allocation mechanisms such as lotteries. As shown in Figure 1, the environment consists of a single, indivisible good that is to be allocated to one of two individuals. In principle, this good could be anything, but to fix ideas it may be useful to think of it as a license of radio spectrum.[2] In this case, the individuals would be two providers of mobile telephone services. Each individual knows privately his value for the good. Both individuals and the seller know that each individual's value lies between $500,000 and $1,000,000 and that the seller has no use for the good, so that his value is $0. The design problem is to find a way of allocating the good to the individual who values it most.

In line with the historic development for allocating radio spectrum in many countries, including the U.S. and Australia, we compare two mechanisms for the *primary market*, that is, for allocating the good initially from the seller to one of the two individuals, a lottery and an auction. In the lottery, the good is randomly allocated to one of the two individuals, say, by flipping a fair coin. The auction is an English (or open ascending) auction, which stops when no further bids are submitted and allocates the good to the bidder with the highest standing bid. Importantly, under either mechanism after the primary market closes, the two individuals are allowed to subsequently trade in a *secondary market* if they find that such trade is to their mutual benefit.

---

[1] An allocation is said to be Pareto efficient if no individual can be made better off without making another individual worse off. Pareto efficiency is a necessary condition for maximizing total welfare.

[2] A radio spectrum license is a right to broadcast over a specific frequency of the radio magnetic spectrum to the exclusion of all other companies. It is used, for example, by mobile phone companies to send and receive data and by radio and television companies to broadcast programming.

Environment

- Two Individuals
- Each with a value between \$500,000 and \$1,000,000
- Values are individuals' private information

Mechanism

- Mechanism 1: Lottery
- Mechanism 2: English Auction

Strategic Interaction

- Nash Equilibrium: Individuals do what is best for them, taking into consideration the actions of others.

Outcome

- Lottery: Inefficient Allocation in the Primary Market
- English Auction: Efficient Allocation in the Primary Market

Figure 1: The mechanism design approach to deciding between a Lottery and an English Auction for allocating a single good to one of two individuals.

Since each individual is awarded the good with a probability of ½ independently of his value, the lottery allocates the good to the low value bidder with probability of ½ if each individual is a priori equally likely to have the higher value. Consequently, the lottery is an allocation mechanism whose outcome is inefficient in the primary market with strictly positive probability.[3] In the English auction, on the other hand, each individual is best off to continue to bid as long as the price is less than or equal to her own valuation and she is not already winning the object. Therefore, the individual with the highest valuation will win the good in the English auction at the valuation of the second highest bidder. As can be seen immediately, under the English auction the allocation of the good is always efficient in the primary market. Consequently, there is no scope or need for an efficient secondary market if the English auction is used in the primary market. Therefore, the comparison between the English auction and a lottery hinges, in a sense, on the efficiency of the secondary market under the latter.[4]

[3] It is easy to see how these results generalize when there are $N \geq 2$ individuals interested in the single object being sold and that the resulting inefficiency in the primary market increases in N. Moreover, in reality the windfall profits a lottery and the subsequent secondary market offer are likely to induce many individuals to participate in the lottery who have no use for the good at all, which exacerbates the inefficiency of this allocation mechanism.

[4] In another sense, it does not depend on the secondary market at all – because the English auction achieves efficiency in the primary market, it already follows that the lottery cannot do better. We next show that it actually does worse.

Intuition suggests that the secondary market may indeed be able to correct any inefficiency resulting from the initial allocation. This idea is largely based on the model of perfect competition and the Coase Theorem, which states that absent transaction costs trade will always yield an efficient final allocation regardless of the initial allocation of property rights (Coase, 1960). In the model of perfect competition, Pareto efficiency is achieved because every agent has a dominant strategy to choose the quantity that is optimal for him. The problem of inducing efficient allocations thus reduces to finding the right prices, that is, the prices that represent the social opportunity costs of the various resources being used. If all agents act non-strategically, finding such prices is indeed possible. However, in a secondary market with only one buyer and one seller, or more generally, with only few buyers and few sellers, individuals will act strategically and take into account the effects of the information they reveal on the prices they face.[5] Foreshadowed by findings reported by Vickrey (1961) and Hurwicz (1972), Myerson and Satterthwaite (1983) have shown that there is actually no mechanism that does not require an injection of funds from an outside source and that achieves an efficient outcome in the secondary market when the buyer's and the seller's valuations are their private information.[6]

This comparison between an auction and a lottery as allocation mechanisms, in conjunction with the results on the inefficiency of secondary markets due to Myerson and Satterthwaite, is helpful in improving our understanding of the evolution of mobile telephony. As mentioned, prior to the early 1990s radio spectrum was allocated to private users by way of lotteries. This resulted in a fractionalized pattern of ownership and very limited use of mobile phone in the U.S. In 1993, President Clinton signed an act that laid the path for the first radio spectrum auction, which was run by the Federal Communications Commission (FCC) in 1994, and shortly thereafter mobile phones became so widely used that by now it is hard to think of what the world was like without them.[7]

II. Auction Theory

Most auction formats can be categorized according to two dimensions: (a) what particular price the winner pays, and (b) whether or not bidders see each others' bids. When bids are observed,

[5] Consider a situation with only one buyer and one seller who are privately informed about their valuations and costs, which are such that it is sometimes but not always efficient to trade. In this environment both privately informed agents have some market power because they can affect the probability that trade occurs. Their incentives to strategically misreport their information are therefore very similar to those of a profit maximizing monopoly who optimally charges a positive markup-up price despite the inefficiency this involves.

[6] Intuitively, because of the seller's incentives to exaggerate costs and the buyer's incentives to underreport his value, an efficient trading mechanism requires a subsidy that is injected from an outside source. Note also that the Myerson-Satterthwaite result does not say that no efficient market clearing prices exist, but that they cannot be found because of the individuals' incentives to strategically distort information.

[7] See Milgrom (2004) for a broader discussion.

the format is called open. Otherwise, it is customary to refer to it as a sealed bid format. The most widely used pricing mechanisms are first and second price auctions. In either format, the winner is the bidder with the highest submitted bid. In the first price auction, the winner pays his own bid, whereas in a second price auction he pays the second highest bid.

Single Object Auctions. The auction model that is best understood theoretically and most studied experimentally has one seller of a single indivisible good and *N* risk neutral buyers who face no budget restrictions. If every buyer's valuation for the good is an independent draw from the same distribution, which is also known as the independent private values (IPV) model, then a very remarkable result holds, which has become known as Revenue (or Payoff) Equivalence Theorem (RET). RET says that in the IPV model the seller's expected revenue is the same irrespective of the auction format he chooses, as long as the format is such that the good ends up in the hands of the buyer who values it the most.[8] Thus, any standard auction format, such as the open ascending (or English) auction that is used to sell houses in Australia, in particular in Victoria and New South Wales, by Sotheby's and Christie's to sell art, and by eBay, the open descending (or Dutch) auction that is used in wholesale flower markets in the Netherlands , the first price sealed auction that is frequently used in procurement, and the second price auction proposed by Vickrey (1961) all yield the same expected revenue. Relying on RET, Myerson (1981) has derived the selling mechanism that maximizes a seller's expected profit when buyers are *ex ante* identical and their valuations are private and independent. Importantly, this mechanism can be implemented with an English auction and an appropriately chosen reserve price. This suggests an explanation for why this format is widely used, for example, by real estate agents, Sotheby's, Christie's, and more recently by eBay.

It may seem that RET makes auction design irrelevant, just as the Coase Theorem may seem to imply that the initial allocation of property rights is irrelevant for the efficiency of the final allocation. However, exactly like the Coase Theorem provides an important theoretical benchmark but rests on the unlikely assumption that there are no transaction costs, RET provides an important theoretical benchmark for auction design but hinges critically on the underlying assumptions of independent private values, risk neutrality and no budget constraints. In a seminal paper, Milgrom and Weber (1982) showed that open auction formats will fare better if bidders' valuations are not purely private because the open format will allow bidders to infer something about the value of the good from observing each others' bids. In the presence of risk averse bidders, by contrast, the seller's expected revenue is larger in a first price auction than under a second price auction (Cox et al, 1982). Revenue equivalence also ceases to hold if some bidders face binding budget constraints (Che and Gale, 1998) or if buyers are boundedly rational (Crawford and Iriberri, 2007). Empirically, the English auction is very widely used, and esoteric, but theoretically revenue equivalent formats are never used. This strongly indicates that

[8] Due to the standard (albeit often implicit) assumption that there is no cost associated with running the auction or selling the object, the seller's revenue coincides with his profit.

differences in design, such as transparency and simplicity, are factors of critical importance in the real world.

Auctioning Many Objects. An important result in the economics of auctions is that with private values there always exists a mechanism which can efficiently solve the initial allocation (or primary market) problem, independent of the number and nature of objects to be sold. Vickrey (1961) constructed such a mechanism for both the allocation of a single indivisible good and for the allocation of many units of the same homogenous goods.

In the single object case, the Vickrey auction is a second price sealed bid auction. It gives every bidder a dominant strategy to bid one's valuation because no bidder can ever affect the price she pays. The reason for this remarkably simple bidding strategy is that prices are based on social opportunity cost, which is the second highest submitted bid. If many items of the same homogenous good are for sale and if every bidder has demand only for one of them, then the Vickrey auction induces dominant strategies to bid truthfully by pricing the good at the highest losing bid.

Importantly, the efficient mechanism that gives a dominant strategy to every bidder can be generalized considerably, as shown partly by Vickrey, and then subsequently and independently, by Clarke (1971) and Groves (1973). The resulting mechanism has become known as Vickrey-Clarke-Groves or VCG mechanism. The VCG mechanism generalizes the second price auction to any number of possibly heterogeneous goods and to any number of buyers. Like the second price auction, it endows buyers with dominant strategies to bid their values because no buyer can ever affect the price she pays, and it allocates goods efficiently because prices represent social opportunity costs.

To gain an intuition for the VCG mechanism, it is useful to first discuss the important complexity that arises when multiple objects are sold simultaneously. Unlike in the single-object case, from the point of view of different bidders objects sold in multi-object auctions may be both *complements* and *substitutes*. For example, consider a simple auction where two spectrum auction licenses are being sold, one for Sydney and one for Melbourne. For a mobile telephone carrier who wishes to serve all of Australia, the value of both licenses is higher than the value for one license. Thus, this bidder would prefer to win both licenses but may have very limited value for a stand alone license. By contrast, a different company may wish to enter a single market and may prefer to purchase one license or the other.

In order to facilitate both types of preferences, the VCG mechanism allows for individuals to express preferences for each good individually and for the package of both licenses. The goods are assigned to maximize efficiency while the payments are constructed so that a bidder can never influence their own payments.

Unfortunately, from a practical perspective the VCG mechanism has serious limitations. While in the example with two goods, the VCG mechanism requires only 2^2-1=3 bids, bid complexity grows exponentially with the number of goods being sold. If, for example, there are 10 goods for sale, every bidder would be asked to submit valuations for every out of the 2^{10}-1=1023 possible packages he could in principle be allocated. For rather obvious reasons, it will typically not be possible to ask bidders to submit so many bids, let alone expect them to evaluate each possible package so that they have any confidence of how much they should bid on it. Thus, the VCG mechanism provides a theoretical benchmark for what can possibly be achieved. However, in order to be practical, a satisfactory auction design must impose much less of a (computational) burden on the bidders. How such more practical designs have been developed, and what kinds of other obstacles they have to overcome or mitigate, is addressed next.

III. Combinatorial Auction Design

Despite the complexity which comes with allowing for complements and substitutes, the fundamental intuition of the VCG mechanism is informative for developing more practical auction formats:

1. Optimal design must limit the strategic behavior of individuals so that the submitted bids are as close as possible to the true valuations of the bidders.

2. Bidders must have the ability to reflect their true preferences in a way which protects them from undue risk.

3. The auction itself must take care of some of the complexity inherent in the assignment problem so that individual bidders themselves do not have to be experts at bidding and can instead concentrate on understanding their core business instead.

Recent designs of so called combinatorial auctions has attempted to address these three core criterions both by managing the number of potential packages and by developing mechanisms that reduce the action space and thereby limit strategic behavior.

The first generation spectrum auctions used by the FCC relied on a format known as the *simultaneous multiple round auction (SMR).* It auctions off multiple goods simultaneously in an ascending auction but does not allow for any package bidding. While this design drastically reduced the complexity of the auction and vastly outperformed the lotteries which they replaced, the SMR design had a number of weaknesses which newer designs have tried to address. First, the SMR generates an *exposure problem* for bidders requiring complementary goods, which consists of the risk a bidder faces who may only win a subset of complementary licenses, so that he loses money in the auction process. These losses led to a large number of bankruptcies in early FCC auctions and to cautious bidding by bidders who desired multiple complementary units in subsequent offerings.

In order to address the exposure problem, new designs have increasingly allowed for package bidding on a subset of potential packages. As the strategies that might be adopted in these complex environments are difficult to model theoretically, economists have often turned to experiments as a way of studying auction formats and analyzing bidding patterns. The latest FCC auction, for example, used a design that included *Hierarchical Package Bidding (HPB)* (Goeree and Holt, 2010), which has allowed bidders to submit package bids on a subset of potential license combinations. This auction format was adopted after extensive testing of a number of potential designs which varied in the number of packages (Brunner et. al, 2011).

A second problem with the SMR auction is that the bidding process allowed bidders too much freedom in their bids which unduly allowed for strategic behavior. By allowing individual bidders to control the way prices unfold in each round, the SMR has the problem that individuals can often use their bids to signal to others profitable ways to collude.[9] A major focus has been to develop auctions which reduce the actions of the bidders while still allowing for an expression of preferences. This body of research has led to the adoption of combinatorial clock auctions, which can be thought of as a generalization of the English auction to settings with multiple objects. In these auctions, prices rise over time according to an independent clock and individuals can choose whether to remain bidding on an object or dropping out. In experimental tests, these clock auctions have outperformed full combinatorial designs and designs which do not allow for packages bids (Porter et al., 1993).

IV. Final Remarks

Auctions such as the open ascending auction (or English) auction have been in use for the allocation of a small number of/single indivisible objects for millennia, dating at least as far back as Roman times. It has only been in the last thirty years, however, with the advent of computers and more sophisticated mathematics that the true potential of auctions as a way to create efficient initial allocations has been recognized.

Despite its successes, design economics is still at a nascent stage. As pointed out by Roth (2002), much like the shift from physics to engineering or biology to medicine, the move from positive economics to normative design has required a stronger understanding of how small differences in the environment can influence optimal design. Experiments have highlighted the need to tailor auctions to the specific auction environments which are being considered. In the context of spectrum auctions, the amount of available bandwidth relative to the maximal usage

[9] A striking example of this is reported in Klemperer (2002) in which Germany sold ten blocks of spectrum using an SMR auction with a requirement that each block had to exceed the previous high bid by at least 10 percent. One of the two major players bid 18.18 million deutschmarks per megahertz on blocks 1-5 and 20 million deutschmarks on blocks 6-10 thereby signaling to its rival that it should bid 20 million on blocks 1-5 and split the market. While there was no standing agreement between parties, the ability to adjust bids allowed the two firms to tacitly collude.

by potential buyers plays an important role in determining auction format. In Australia, for instance, where the amount of spectrum is large relative to demand, the auction format must take into consideration the potential of tacit collusion amongst bidders which can hurt both efficiency and revenue generation. By contrast, in the US, where the amount of spectrum is small relative to demand and the bidders are highly heterogeneous, ensuring that bidders who want only a small number of units can overcome bids by individuals who want a large number of units play a more important role in auction design.

There is a tendency in any new, exciting field, to overemphasize what is new and ignore the lessons learned from similar fields in the past. With respect to auction theory, practitioners often spend a lot of time considering small differences in the details of the mechanism, neglecting other and equally important parts of the allocation problem. For example, attracting new bidders who have interest in the goods being sold is more important than setting the reserve prices optimally (Klemperer, 2002).[10] Likewise, political constraints generated by lobbying and special interest groups are among the most common reason for inefficient outcomes. Design economics requires not only a strong understanding of new theoretical and experimental methods, but also a strong understanding and sound use of more traditional economic insights stemming from Industrial Organization and Political Economics.

References

Brunner C., Goeree, J.K., Holt, C.A. and Ledyard, J.O. 2010, 'An Experimental Test of Flexible Combinatorial Spectrum Auction Formats,' *American Economic Journal: Microeconomics*, vol. 2(1), pp. 39-57.

Che, Y.K. & Gale, I.L. 1998, 'Standard Auctions with Financially Constrained Bidders', *Review of Economic Studies*, vol. 65(1), pp. 1-21.

Coase, R.H. 1960, 'The Problem of Social Cost', *Journal of Law and Economics*, vol. 3(1), pp. 1–44

Cox, J., Roberson, B. and Smith, V. 1982, 'Theory and Behavior of Single Object Auctions', in V. Smith ed., *Research in Experimental Economics*, vol. 2, Greenwich: JAI press.

Crawford, V.P. and Iriberri, N. 2007, 'Level-k Auctions: Can Boundedly Rational Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?', *Econometrica*, 75(6), pp. 1721-1770.

[10] Entry of new bidders can be facilitated in a number of ways, including removing regulatory entry barriers and bid credits for new entrants, where bid credits $c>0$ are such that if a new bidder bids x on some object an incumbent bidder must bid $(1+c)x$ to win the object.

Goeree, J. K. and C. Holt (2010), 'Hierarchical Package Bidding: A Paper& Pencil Combinatorial Auction', *Games and Economic Behavior*, vol. 70(1), pp:146-169.

Hurwicz, L. 1972, 'On informationally decentralized systems', In McGuire, CB and Radner, R (editors): *Decision and Organization*. North-Holland.

Kagel, J., Levin, D. and Harstad, R. 1987, 'Information Impact and Allocation Rules in Auctions with Affliated Private Values: A Laboratory Study', *Econometrica*, 55(6), pp. 1275-1304.

Klemperer, P. 2002, 'What Really Matters in Auction Design', *Journal of Economic Perspectives*, 16(1), pp. 169-189.

Maskin, E. 2005, 'The Unity of Auction Theory: Milgrom's Masterclass', *Journal of Economic Literature*, 42(4), pp. 1102-1115.

Milgrom, P. and Weber, R.J. 1982, 'A Theory of Auctions and Competitive Bidding', *Econometrica*, vol. 50, pp. 1089-1122.

Milgrom, P. 2004, *Putting Auction Theory to Work*, Cambridge University Press.

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

Myerson, R.B. and Satterthwaite, M.A. 1983, 'Efficient Mechanisms for Bilateral Trading', *Journal of Economic Theory*, vol. 29(2), pp. 265-281.

Porter, D., Rassenti, S., Roopnarine, A. and Smith, V. 2003, 'Combinatorial auction design.', *Proceedings of the National Academy of Sciences of the United States of America*, vol. 100(19), pp. 11153-11157.

Roth A.E. 2002, 'The Economist as Engineer: Game Theory, Experimentation and Computation as Tools for Design Economics', *Econometrica*, vol. 70(4), pp. 1341-1378.

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

Vickrey, W. 1962, 'Auction and Bidding Games', in *Recent Advances in Game Theory* (Papers delivered at a meeting of the Princeton University Conference, October 4-6 1961), Philadelphia, PA: Ivy Curtis Press, 15-27.