Auction Theory and Applications

Prof. Dr. Benny Moldovanu*†

Introduction

In an auction participants (repeatedly) submit bids representing their demand or supply functions. Then, accepted trades and transaction prices are calculated by some explicit aggregation procedure. Auctions directly implement the ideas implicit in the Walrasian analysis of finding prices that are consistent with individual maximization and that equate demand and supply. In contrast to general equilibrium analysis, auction theory is based on the premise of individual strategic behavior. It offers explicit models of price formation and allocative distribution that can be applied also to small markets. Auctions have been continuously used since antiquity, and they remain simple and ubiquitous means of conducting multilateral trade. The belief that auctions yield competitive outcomes even if information is dispersed is behind the practical appeal of auctions, and behind their recent popularity.

Modern, large-scale applications include treasury auctions, spectrum auctions, private and government procurement auctions, CtoC, BtoC and BtoB internet

*Benny Moldovanu has been a full professor at the University of Bonn since 2002. Further, he was visiting professor at Yale University in 2004, at the University College London in 2005 and at Northwestern University in 2009. The winner of the 2010 Advanced Investigators Grant, awarded by the European Research Council, is one of the leading researchers in the field of mechanism design and has published in the leading international academic journals. Moldovanu is member of CEPR and of the European Economic Association and has designed several auctions for large government and private sector projects.

†The present material is based on insights found in Jehiel and Moldovanu (2006) and Jehiel and Moldovanu (2003). I am grateful to Philippe Jehiel for many years of fruitful collaboration.

auctions, real-estate auctions, commodities auctions and asset auctions following bankruptcy. In practice auctions are often used in order to achieve other important goals, besides, or instead efficiency: 1) Revenue-maximization. The seller of an item at eBay, say, does not care much about allocative efficiency, but rather aims at maximizing the price he gets. 2) Information Aggregation and Revelation. Auctions aggregate bids which are made based on the basis of privately known demand and supply functions. Hence, the resulting prices can also be seen as aggregators of private information. For example, data collected at auctions of liquidity organized by central banks is an important indicator for future monetary policy. 3) Transparency and Speed. In serious auctions the rules are precise, fixed in advance, and applied equally to all participants. Auctions provide the means for achieving quick, frictionless transactions. Speed is particularly important for perishable goods such as fish, vegetables or flowers. Accordingly, most whole-sale markets for agricultural products have long been organized as auctions.

The above goals are not independent of each other. Depending on the economic environment and on the market process, there are subtle relations among them, and sometimes even severe conflicts.

Traditionally, the focus of auction theory has been on models that view auctions as isolated events. In practice, however, auctions are often part of larger transactions: for example, in privatization exercises such as license allocation schemes auctions shape the size and composition of future markets. Thus the auction typically affects the nature of the post-auction interaction among bidders. On the other hand, anticipated scenarios about the future interaction influence bidding behavior: already at the bidding stage agents need to care about who gets what, and about the information revealed to, or possessed by others, since these features will be reflected in the equilibrium of the post-auction interaction. Thus allocative and informational externalities naturally arise in models that embed auctions in larger economic contexts.

General equilibrium analysis has identified several forms of externalities as obstacles on the road towards economic efficiency. The First Welfare Theorem fails in the presence of allocative externalities, i.e., when agents care about the physical consumption bundles of others. Akerlof's famous analysis demonstrated that the First Welfare Theorem may also fail when agents care about the information held by others. Thus, given the obstacles created by externalities in general equilibrium, it is of interest to understand what are the parallel consequences of external effects in auctions. This has been one of my main research topics (in cooperation with Philippe Jehiel and other colleagues).

An Illustration

As an illustration, consider the European process of allocating UMTS spectrum and licenses to telecom firms in $2000 - 2001$. The auctioned objects were licenses to operate a third-generation mobile telephony network in a certain country. Since per-firm industry profit in oligopoly decreases in the number of active firms, incumbents were also driven by entry preemption motives (e.g., the need to avoid further losses relative to the status quo) which translate into increased willingness to pay for licenses and capacity. To see this force at work, let us recall the German experience. Bidders did not directly submit bids for licenses but, instead, on 12 blocks of spectrum. A bidder obtained a license only if he acquired at least two blocks, and a bidder was allowed to acquire three blocks. Thus, the number of licensed firms could vary between 0 and 6. There were 7 bidders, including 4 GSM incumbents. The auction lasted for 173 rounds, and the winning firms were the 4 incumbents and two new entrants. Each licensed firm acquired 2 blocks and each license cost approximately Euro 8.4 Bn (4.2 Bn per block).

The most interesting thing occurred after one of the potential entrants left the auction in round 125, after the price reached Euro 2.5 Bn per block. Since 6 firms were left bidding for a maximum of 6 licenses, the auction could have stopped **immediately**. Instead, bidding in order to acquire more capacity and to reduce the number of competitors continued until round 173 (only intense pressure from stock markets and bond rating agencies stopped it). Compared to round 125, there was **no change in the physical allocation** but firms where, collectively, Euro 20 Bn **poorer**! Given the immense sum that was paid for licenses, the two new entrants went bankrupt (this is called a "winner's curse" in auction jargon) and did not build networks, thus leaving Germany with the four old network operators.

Single-Object Auctions

The simplest theoretical setup is one in which bidders know how much they value the good for sale, but are uncertain as to how much other bidders value the good. This is the so-called "private value" paradigm. In the ascending price (or English) auction the price gradually increases, bidders may drop out at any point in time, and the auction stops when there is only one bidder left. The last active bidder buys the good at the price where the auction stopped. In the private value setting, the ascending price auction induces an efficient outcome The reason is that it is a dominant strategy to drop out whenever the price reaches one own's valuation. In a Nobel-prize winning paper, Vickrey (1961) proposed a condensed version of this auction, now called the Vickrey auction or, in the context of one-object auctions, the sealed-bid second-price auction, in which agents secretly place bids. The bidder with the highest bid wins the object and pays the second-highest bid. Vickrey observed that in the second-price sealed-bid auction it is a dominant strategy to bid one own's valuation. Hence this format is here equivalent to the ascending price auction. It turns out that these formats, augmented by a reserve price are also **revenue maximizing** whenever bidders' signals about their values are independent of each other, agents are risk neutral, ex-ante symmetric, and the seller is bound to sell her good.

Laymen feel that more money can be extracted by requesting that the winner pays the highest bid, rather than the second highest bid. Such a format corresponds to the commonly used first-price sealed-bid auction. The argument ignores that bidders react to a change of format by adjusting their bidding strategies. In a first-price auction, a bidder who bids her own value never makes any money. Hence bidders bid less than their value, and bids will be lower than in the second-price version (but the seller receives the highest bid, not the second highest). It turns out that, in expected terms, the first-price auction generates exactly the same amount of revenue as the second-price auction as long as the agents are symmetric and risk neutral, and obtain independent signals about their respective values. This result is the celebrated Revenue Equivalence Theorem, first noticed by Vickrey. Many of the above strong conclusions heavily rely on ex-ante symmetry, risk-neutrality, the absence of budget constraints, signal independence, and the absence of informational or allocative externalities, and need to be adjusted if these features are not present!

The analysis of auctions with externalities is quite subtle even in the one-object case. One reason is that the notion of "valuation" is not well defined a-priori. Specifically, how much a bidder is ready to pay depends on his expectation about what will happen if he does not buy the object. If he expects the winner to be a tough competitor, he will be ready to pay a high price; if he expects the winner to be a soft competitor, his willingness to pay will be low.

Example 1. ***A Takeover Contest.*** *Consider three firms bidding for a fourth in a takeover contest. Consider the following expectations about future scenarios in the post-takeover industry: due to synergies, each bidding firm expects to make* *an extra profit of* π *if it wins the contest; if firm* 1 *wins, firm* 2 *expects a relative decrease in profits of* α *(and vice versa); firm* 3 *expects a decrease in profits of* $\gamma < \alpha$ *if firm* 1 *or firm* 2 *wins; finally, firms* 1 *and* 2 *expect to be unaffected if* 3 *wins. From the firms' viewpoint (e.g. abstracting from consumers' surplus), the efficient buyer is firm* 3 *(since the other firms do not expect to suffer a loss in that scenario). Consider a first price sealed-bid auction: In one equilibrium, firm* 3 *indeed wins by bidding only* π*. But in another, firms* 1 *and* 2*, who are very afraid of each other, engage in a race and one of them wins by bidding up to* $\pi + \alpha$*. In that case, both firms suffer a loss of* $-\alpha$*! If firms* 1 *and* 2 *think that the expensive race is going to happen because they cannot coordinate to let* 3 *win, they will have incentives to commit not to participate at the auction in the first place. For example, if firm* 1 *withdraws,* 3 *necessarily wins since it is willing to bid up to* $\pi + \gamma$*, while firm* 2 *is willing to bid only up to* π *(since* 1 *poses no danger anymore). This scenario is, in fact, better for firm* 1 *than participating in the race with* 2.

If, as above, the allocative externalities are due to market structure considerations, the notion of economic efficiency should not be solely based on considerations about firms' welfare. Instead, the welfare considerations should also include the consumer's surplus in each possible future scenario. But, a more thoughtful design that addresses the consumers' interest may generate low revenue.

Example 2. ***Competition over Monopoly Rents.*** *There are two licenses* A *and* B *for sale. There are two firms* $i = 1, 2$ *competing for the two licenses. Each firm* i *is allowed to buy both* A *and* B*. If firms* 1 *and* 2 *each buy one license, price-competition is assumed to drive profits down significantly (say to zero). If firm* i *buys both licenses, it earns monopoly profits* π_i*. We assume that* $\pi_1 > \pi_2$*. Consider the standard format where firms simultaneously submit two bids* b_i^A, b_i^B *for licenses* A *and* B*, respectively. Each license is allocated to the highest bidder* *on that license, who pays the bid. Since one license is worthless (as long as the other is also sold), the outcome of this auction is that firm* 1 *gets the two licenses and pays* π_2 *for it. In other words, the auction selects a monopoly structure, and the resulting market structure is not desirable from the consumers' viewpoint. If the government allows each firm to buy only one license, a duopoly may emerge (which is presumably better for consumers and total welfare) but the auction's revenue will be low.*

Multi-object auctions

Multi-object auctions raise a large number of difficulties. Even when externalities are absent, multi-unit demand, heterogeneity and complementarities induce complex demand functions, which are difficult to map in reasonably simple auction formats. Simple formats necessarily restrict bidders in some dimension, creating complex strategic effects that affect the auction's performance. We illustrate below several difficulties arising in such auctions.

With single unit demand bidders and with k homogeneous units, the $k+1^{th}$ -price auction induces an efficient allocation. When bidders have multi-unit demand, a seemingly easy generalization of this format is the uniform-price auction: bidders submit demand curves (i.e., bids for 1 up to k units), and the units are allocated to maximize the values expressed by the submitted demand curves; every allocated unit is sold at the same minimum price where aggregate demand coincides with the number of supplied units k.

Unfortunately, if bidders have multi-unit demand, the uniform price auction may lead to an inefficient allocation since bidders have an incentive to lower their demand on all units (but the first). In doing so, they affect downwards the selling price and pay a lower price on the remaining units.

Example 3. ***Demand Reduction.*** *Three identical cases of Bordeaux wine are sold through a uniform price auction. There are three potential bidders,* $i = 1, 2, 3$. *Bidders* 2 *and* 3 *are interested in one case only; their valuations are* 1 *and* 0.25*, respectively. Bidder* 1 *is potentially interested in all three cases. His valuation is* 10 *for the first unit,* 5 *for the second, and* 2 *for the third. Efficiency in the above example dictates that bidder 1 gets all three cases. Bidders 2 and 3 have a dominant strategy—to bid their values for one case. If bidder* 1 *expresses his true demand, the minimum price where demand equals supply is (slightly above)* 1*, and bidder 1's payoff is given by* $10 + 5 + 2 - 3 = 14$. *However, in the equilibrium of the uniform price auction, bidder 1 will decrease his demand to only two cases (e.g., he will bid* 15 *for either two or three units). This lowers the selling price per unit to* 0.25*, yielding for bidder 1 a payoff of* $10 + 5 - 0.5 = 14.5$. *The allocation is inefficient since bidder 2 obtains one case.*

Complementarities are particularly troublesome in auction formats where bids can be placed only on individual objects, but not on bundles (or packages). Bids cannot then fully reflect the magnitude of complementarity, creating inefficiencies. The theoretically correct way to deal with complementarities is to allow for "combinatorial" auctions where agents can place bids directly on bundles. Forbidding such bids may give rise to the so called *exposure problem:*

Example 4. ***The Exposure Problem.*** *There are two parking slots, and two bidders. Bidder 1 has a car and a trailer, and he values the two parking slots together at $100, while attaching a value of zero to each individual slot. Bidder 2 has only a car and values any slot at $75. The value of two slots is also $75 for this bidder. Efficiency dictates that bidder 1 gets both slots. But if the auction format does not allow bids on the whole package of two slots, any positive bid on an individual slot exposes bidder 1 to the danger of obtaining only that slot alone—an alternative valued at zero. The only equilibrium in a simultaneous* *ascending auction without combinatorial bids is for the first bidder to drastically reduce demand by non-participation (since otherwise he needs to bid up to $75 per slot in order to outbid the other player). Bidder 2 inefficiently wins a slot by placing a minimum bid on it. In the presence of incomplete information, bidder 1 will bid only if he attaches a sufficiently high probability to the event in which bidder 2 has a low valuation. If bidder 1 decides to bid (based on his information), and if it turns out that bidder 2 has a high value, bidder 1 will regret his decision.*

The main problem with combinatorial auctions is that they may be very complex to conduct and participate at. Besides the structural complexity arising in large auctions, combinatorial bids also induce some subtle strategic problems, such as inefficient free riding.

Example 5. ***Free Riding.*** *Two regional radio licenses are put for sale. There are three potential bidders* $0, 1, 2$. *Bidder* 0, *who needs national coverage, values only the bundle* $\{1, 2\}$ *at* v^{12}. *Bidder* i, $i = 1, 2$, *values only license* i *at* v^i. *Assume that* $v^1 + v^2 > v^{12}$. *Each bidder* i *simultaneously submits a bid* b_i *for whatever good or bundle she wishes. The goods are allocated so as to maximize the revenue generated by the bids, and each bidder pays for the goods he receives according to the bid he submitted ("pay-your-bid" auctions). Efficiency dictates that bidder* i *receives object* i, $i = 1, 2$. *But in equilibrium there will be a "war of attrition" between bidders* 1 *and* 2. *Instead of bidding up to* v^1 *on object* 1, *bidder* 1 *prefers to place a low bid on object* 1 *(say* $v^{12} - v^2$*), hoping that bidder* 2 *will make a high bid on object* 2 *(say* v^2*). Similarly, bidder* 2 *prefers to place low bid (say* $v^{12} - v^1$*), hoping that bidder* 1 *will make a high bid on object* 1 *(say* v^1*). As a consequence, there is an equilibrium in mixed strategies where bidder* 0 *gets the bundle with positive probability.*

In one-object symmetric settings, standard auctions are efficient and revenue-maximizing (at least if the seller is not allowed to retain the good). Thus, efficiency and revenue go hand in hand. This ceases to be true in multi-object settings, even if the situation is symmetric and there are no complementarities.

Example 6. ***Efficiency or Revenue?*** *Consider an auction for two objects A and B, and two bidders, 1 and 2. For both agents, the valuations for the bundle $\{A, B\}$ are given by the sum of the valuations for the individual objects, and assume these to be as follows:*

$$\begin{aligned} v_1^A &= 10; v_1^B = 7 \\ v_2^A &= 8; v_2^B = 12 \end{aligned}$$

The value maximizing auction (which puts the objects in the hand of those who value them most) is simply given by two separate second-price auctions, one for each object. Then object A goes to bidder 1 for a price of 8, while object B goes to bidder 2 for a price of 7. Total revenue is 15. But, consider now a single second-price auction for the entire bundle $\{A, B\}$. Then the bundle will be acquired by bidder B, for a price of 17 ! Hence, revenue is higher in the bundle auction, but object A is mis-allocated.

The presence of multiple objects for sale creates many new possibilities for "tacit" collusion in which firms coordinate their bids instead of competing. The main idea is that it may be preferable to share the objects at low prices instead of trying to buy more of them while pushing prices up. Ascending price formats are more vulnerable to such behavior since they offer repeated opportunities to signal intentions and future behavior. In contrast, sealed bids greatly reduce the scope of signalling, but run the risk of yielding an inefficient allocation since private information does not get properly aggregated.

Example 7. ***The German GSM Auction.*** *In October 1999 Germany auctioned 10 additional blocks of paired spectrum to the four GSM incumbents. Nine blocks were identical, each consisting of* 2×1 *MHz, while the tenth block consisted of* 2×1.4 *MHz. After the first round, the high bidder on all 10 blocks was Mannesmann (one of the two large players), which offered DM 36.360.000 for each of blocks 1-5, DM 40.000.000 for each of the blocks 6-9 (which, recall, were* ***identical*** *to blocks 1-5), and DM 56.000.000 for the larger block 10. In the second round, T-Mobil (the other big player) bid DM 40.010.000 on blocks 1-5, and the auction closed! Hence, each of the two larger firms got 5 blocks , at a price of DM 20.000.000 per MHz. Here is what one of T-Mobil's managers said: "No, there were no agreements with Mannesmann. But Mannesmann's first bid was a clear offer. Given Game Theory, it was expected that they show what they want most."*

Conclusion

Each market process creates specific strategic incentives for the participants and leads to specific distortions. Thus, the "rules of the game" do matter and deserve attention. Different auction formats lead to unavoidable trade-offs among goals. Thus, a clear determination of the auction's goals is indispensable before the choice of the auction format can be discussed. The precise structuring of traded goods and feasible bids play a main role in determining whether the auction procedure can accurately represent the agents' preferences and lead to a desired outcome.

Allocative or informational externalities, multi-unit demand, heterogeneity and complementarities induce complex demand or supply functions, which are difficult to map in reasonably simple auction formats. Such formats necessarily restrict bidders in some aspects, creating complex strategic effects that affect the auction's performance. If the auction's allocation influences some future interaction among bidders, these will take this effect into account at the bidding stage. Thus, the future interaction also influences the auction's outcome through the participants' expectations. In complex environments serious design calls for an integrated approach that combines the insights of Auction Theory with the traditional concerns of regulation and competition policy.

References

Jehiel, P., and B. Moldovanu (2003): "An Economic Perspective on Auctions," *Economic Policy*, 18(36), 271–308.

——— (2006): "Allocative and Informational Externalities in Auctions and Related Mechanisms," in *The Proceedings of the 9th World Congress of the Econometric Society*, ed. by T. P. Richard Blundell, Whitney Newey. Cambridge University Press.

Klemperer, P. (1999): "Auction Theory: A Guide to the Literature," *Journal of Economic Surveys*, 13(3), 227–286.

Krishna, V. (2002): *Auction Theory*. Academic Press.

McAfee, R. P., and J. McMillan (1987): "Auctions and Bidding," *Journal of Economic Literature*, 25(2), 699–738.

Milgrom, P. R. (1987): "Auction Theory," in *Advances in Economic Theory: Fifth World Congress*. Cambridge University Press.

——— (2004): *Putting Auction Theory to Work*. Cambridge University Press.

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

Wilson, R. (1993): "Strategic Analysis of Auctions," in *Handbook of Game Theory*, ed. by S. H. R. Aumann. Elsevier.

Wolfstetter, E. (1996): "Auctions: An Introduction," *Journal of Economic Surveys*, 10(4).