A Brief Introduction to On-Line Auctions

Rotem Shaul

September 7, 2010

Abstract

In this brief paper we illustrate the basic properties of auction theory and in addition portray specific online auction models and their properties.

1 Introduction

Auctions are a commonly used tool for selling goods in cases where a true market does not exist. In the typical case multiple buyers aim to buy some good from a single seller, and the seller wishes to sell the good for the highest possible price. Many types of auction mechanisms have been considered in the literature.

Online auctions can be categorized as a subject studied in Game Theory, and as such it is a branch of applied mathematics that is studied both in social sciences, such as Economics, and in exact sciences, such as Computer Science. Online auctions from computer scientists point of view is an online computation that must produce a sequence of decisions which are made based on past events without secure information about the future. Such computation is usually referred to as online algorithm.

2 Definitions and Terminology

2.1 Game Theory (Informal) Definitions

- *Strategy* - a player's strategy in a game is a complete plan of action for whatever situation might arise; this fully determines the player's behavior. A strategy is not to be mistaken with a single move, a strategy is a complete algorithm for playing the game, telling a player what to do for every possible situation throughout the game.
- *Utility* - utility is a measure of the satisfaction gained from a good or service.
- *Dominant Strategy* - A strategy is Dominant if it is better than any other strategy for one player, no matter how that player's opponents may play.
- *Undominated Strategy* - A strategy that is not dominated by any other strategy. A player without full information is indifferent to playing any

of the undominated strategy for it loses nothing playing those. And the player favors them over dominated strategies evidentally.

- *Nash-Equilibrium* - Nash equilibrium specifies a strategy for each player, in such a way that each players strategy yields the player at least as high a payoff as any other strategy of the player, given the strategies of the other players. In short - at Nash equilibrium no player can benefit by changing his or her strategy.

2.2 Auction Theory (Informal) Definitions

- *Auction* - any mechanism or set of trading rules for exchange.
- *First Price Auction* - a mechanism such that the winner pays his own bid.
- *Second Price Auction* - a mechanism such that the winner pays the second highest bid.
- *Sequential Auction* - a repetitive mechanism formed with rounds, each round item is sold. relationship between round prices and participants is mechanism dependant.
- *Ascending Auction* - a mechanism such that the price ascends and bidders choose to drop, last bidder that doesn't drop is declared as the winner.
- *Revenue* - total payment the auctioneer receives.
- *Social Welfare* - the utility sum of the auction winners. the social efficiency is higher as the players with the higher valuations of an item receives it.
- *Incentive Compatible Auction* - a mechanism in which the bidders best (dominant) strategy is to declare their true valuations, as a consequence such a mechanism achieves optimal social welfare
- *Vickrey Auction* - a famous second price mechanism which is incentive compatible and social efficient.

3 Sequential Auctions with Dynamic Arrivals

3.1 The Model

This model introduces dynamic arrivals and an activity rule to a sequential auction. In addition it lifts two common assumptions used in the classic format of sequential auctions.

1. Static-Set Assumption - this commonly used assumption assumes that there is a static set of players present throughout the auction from the beginning to its end. This assumption is unrealistic in today's dynamic environments such as internet auctions where players arrive dynamically over time.

2. Common-Prior Assumption - this assumes that all players have common beliefs about the future, this is evidentally an unrealistic assumption, even more so considering internet auctions where the players come from a variety of geographical and cultural locations.

Allowing players to have different and even contradicting beliefs combined with their dynamic arrivals has the immediate consequence of losing the truthfulness of the mechanism, players may be better off not saying their true valuation of an item at a given time. This immediately costs in the social welfare of the mechanism, and the main results for this model is the bounding of the social welfare loss.

Losing Truthfulness and Optimal Social Welfare by Example
Assume a seller that conducts two ascending auctions, and three participants players arrive at time one whilst each player desires one item. (They're indifferent about which item to win)
Without dynamic arrivals this is an easy scenario, there is an easy nash-equilibrium where the players with the top two valuations for the item win it, this holds because the mechanism would be truthful and there is no reason for each of them not to bid their highest value at each auction which results in the top player winning at the first round, whilst the second player winning in the second round.
With dynamic arrivals a lot of future scenarios exist, and the players strategies are now dependant on their future beliefs, which as explained earlier are not assumed to be common.
For instance, there is a positive probability for a new highest bidder arriving at time two, depending on the players beliefs at time one, the following can occur: the top player in round one believes he better off waiting for round two, this can occur if for instance he believes that no other player will join in the next auction, and that the the second top player will participate in round one. Since this is a second price auction top player of round one is better off winning only with player three bidding with him.
If his belief does not hold though, the top player of round one will never win an item, this immediately results in a social welfare loss since the optimal welfare is obtained with the top player of each round winning an item, which is not the case here.
Introduction of an Activity Rule
In order to bound the social welfare loss two things are needed, the first is the assumption that the players play undominated strategies, this is a rationality assumption. The second thing is adding an activity rule for the mechanism of the auction.
The Sequential Auction with the activity rule

- When the t'th auction ends(and there are K-t more auctions ahead), only the K-t highest bidders besides the winner are qualified to participate in the next auctions.
- Auction t+1 starting price is the price which exactly K-t+1 players are still active in auction t.

With this activity rule in place, and the assumption of players playing undominated strategies it can be shown that in a worst-case analysis, i.e with an adversary determining the number of players, their arrival times and their (undominated) strategies the lowest possible ratio between this mechanism resulting welfare and the optimum welfare is **at least one half.**

3.2 Formal Settings

The seller sells K identical and indivisible items using a sequence of K single-item ascending auctions.
The set of bidders is $I = \{1, 2, .., n\}$.
Bidders do not discount time - they are indifferent about winning the item at any auction they participate in.
A bidder's type is a pair $\theta_i = (r_i, v_i)$ whereas r_i is the players arrival time i.e it participates in auctions $r_i, .., K$ and v_i is the players value.
Each bidder desires exactly one item out of the K items. Each player's value is private, and the utility for an item is $v_i - p_i$ if she wins an item and pays p_i.
The set of possible types for bidder i is Θ_i.

Each ascending auction is formally a "Japanese" auction, where a "price clock" increments continuously and each bidder is free to drop at any price.
The last bidder that remains is the winner and the price paid the of the second highest bidder who dropped out.
Each time during an auction if a player drops, can lead to other players dropping depending on their strategy (e.g a strategy taking number of players left in consideration) so each time a player drops there is a "clock stop", that is the price stops and all players who want to drop can do so.
The order of dropping is important because it may determine the winner.
Afterwards the auction continues with the price it stopped and the remaining players.
Thus the strategy of a bidder for this auction is a function of her type, history of all previous auctions, and the history of the current auction up to the current price point.

Concrete Auction Example Assume an auction with four players with the following values $v_1 = 5, v_2 = 10, v_3 = v_4 = 8$.
Assume the bidders strategies are as following:

- Bidder one remains until her value
- Bidder two remains until her value or until there remain at most two other bidders.
- Bidder three and four remain until their value or until there remain at most one other bidder.

Given these properties the auction first round will proceed as follows - the price clock will ascend until it reaches a price of 5, then the clock stops and bidder 1 drops, his drop will lead to bidder 2 dropping (because after bidder 1 drops only two other bidders besides bidder 2 are left), and afterwards bidders 3 and 4 will also drop. Note there is a tie only between players 3 and 4, for they wish to drop together, but bidder 2 is considered to drop before them even without a tie-breaking rule. According to a tie-breaking rule either player three or four will win.

3.3 Worst Case Analysis

Proposition 3.1 *In every auction $t = r_i, ..., K$ bidder i does not drop before there remain at most $K - t + 1$ active bidders, unless the price reaches her value before that.*

This means that the bidders with $K - t + 1$ highest values among all bidders that arrive up to time t and have not won yet are qualified for auction $t + 1$. More formally, denote Λ_t the set of bidders that arrive up to time t and have not won yet.
Denote X_t as the set of bidders that participate in auction t.
Denote Q_t as the set of players at auction t that are qualified for auction $t + 1$.
We get to the following :

Corollary 3.1 *if $|X_t| \geq K - t + 1$ and all bidders play a weakly undominated strategy then at each auction t the $K - t + 1$ highest value bidders in Λ_t have the same set of values as the bidders in Q_t.*

This is a direct result of the activity rule in place. The formal proof can be found in the appendix of [1].

With the following understanding we now understand better how the algorithm will behave, and who will participate and proceed in future auctions. (notice fig.1)
We also get to the following result, if **no** dynamic arrivals occur, and all players arrive at auction one. (Or more precisely if the top K players arrive at auction one) this mechanism is social efficient. This easily follows from the fact that whilst not necessarily the top player will win at auction one, but one of the top K will. In addition, the other $K - 1$ top players are guaranteed to proceed to the next auction, and under the assumption that no new arrivals with higher valuations arrive, once again one of the top $K - 1$ will win and the rest are guaranteed to proceed. And so forth it will continue until the last auction, while none of the top K will drop.
So whilst the order of winning isn't guaranteed the fact that the top K will win is guaranteed. This is obviously not the case with dynamic arrivals, which requires a different analysis which is now described.

We now look at the analysis of the social efficiency of this mechanism, under the assumption that players may choose to play any tuple of undominated strategies. More formally:
Fix any realization of the player's type $\theta = (\theta_1, ..., \theta_n)$ where $\theta_i = (r_i, v_i)$.
We denote W as a subset of players which are valid set of winners. A set is a valid set of winners if for a given θ there exists an assignment of items to players in W such that no item is assigned to more than one player, and if item t is

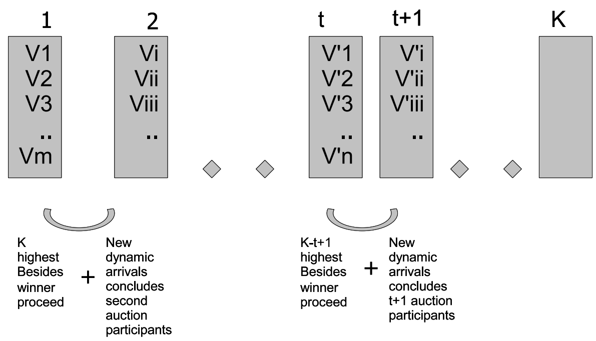

Figure 1: Mechanism with Activity rule in place.

assigned to player i then $t \geq r_i$. Let $W(\theta)$ denote the set of all valid sets of winners.
The value of some $W \in W(\theta)$ is $v(W, \theta) = \sum_{i \in W} v_i$.
$OPT \in W(\theta)$ is socially efficient if $v(OPT, \theta) = max_{W \in W(\theta)} v(W, \theta)$.
Fix any tuple of types θ, let OPT be a valid assignment with maximal efficiency and let A be an assignment that results from this mechanism.

Theorem 3.2 *$v(OPT)$ is at most twice $v(A)$ i.e $v(OPT) \leq 2v(A)$*

Denote $v_1^{OPT} \geq v_2^{OPT} \geq ... \geq v_K^{OPT}$ i.e the value of the winners of OPT ordered in a non increasing orders.
Similarly denote $v_1^A \geq v_2^A \geq ... \geq v_K^A$

Lemma 3.3 *For any index $0 \leq l \leq \lfloor K \div 2 \rfloor$ we have $v_{l+1}^A \geq v_{2l+1}^{OPT}$*

Proof:
Assume by contradiction that there are at most l winners in A with values $\geq v_{2l+1}^{OPT}$.
Denote $K - t$ as the last auction at which the winner in A has value $< v_{2l+1}^{OPT}$.
By the contradiction assumption $t \leq l$ since after auction $K - t$ there remains only t more auctions.
Denote X as the set of the $2l + 1$ highest value players in OPT.
Denote $Y = \{i \in X | r_i \leq K - t\}$. That is the $2l + 1$ highest players in OPT that arrive before time $K - t$.
Denote Z as the set of players in Y that win in **A** before auction $K - t$. Thus $Y - Z \subseteq \Lambda_{K-t}$.
Since all winners in A after auction $K - t$ have values larger than v_{2l+1}^{OPT} we get $|Y| \geq (2l + 1) - t$.
From the definiton of t after auction $K - t$ all winners in A have values at least

v_{2l+1}^{OPT}, and there are at most l such winners by assumption we get $|Z| \leq l - t$.
Thus $|Y - Z| \geq (2l + 1 - t) - (l - t) = l + 1 \geq t + 1$.
Since $Y - Z \subseteq \Lambda_{K-t}$ and $|Y| \geq t + 1$ the $t + 1$-highest value in Λ_{K-t} is larger or equal the minimal value in $Y - Z$.
According to the proposition 3.1, the winner in A at auction $K - t$ has a value at least as large as the $t + 1$-highest value in Λ_{K-t}.
Thus the winner in A at auction $K - t$ has value at least as large as the minimal value in $|Y - Z|$ which is by definition of $Y - Z$ at least v_{2l+1}^{OPT}.
This contradicts the assumption that the winner in A at auction $K - t$ has value strictly smaller than v_{2l+1}^{OPT}.□

This lemma immediately proves the theorem, it shows that $v_1^A \geq v_1^{OPT}, v_2^{OPT}$ and $v_2^A \geq v_3^{OPT}, v_4^{OPT}$ and so on.
More generally $v(OPT) \leq 2 \sum_{k=1}^{\lfloor K \div 2 \rfloor + 1} v_k^A \leq 2v(A)$.□

3.4 False Identities Caveat

As seen in the worst case analysis and as explained in the model description, the activity rule is vital for the social welfare bound to hold. Yet, there is no real mechanism to disallow players from returning in later auctions with a different, false identity. This issue breaks the auction from becoming practical. Dealing with the issue of false identities requires an enforcing tool, or a mechanism a bit different that makes it inefficient for a player to drop and return with a false identity later on. One such suggested mechanism can be a "Customer Club" that rewards players that stays longer, in which case they wouldn't want to use false identities. It won't pay off. This point is an open point for future research for most models using activity rules similar to this one.

3.5 Model's Conclusions

The model main result as already explained is the social-welfare loss bound of factor two in the worst case analysis, in addition it can be shown that for K=2 the average loss bound for the social welfare is better and is approximately thirty percent. In addition as shown this model is equivalent to a Vickrey auction, which is truthful and has optimal social welfare incase that all participants arrive at the beginning. (No dynamic arrivals) And as noted before the biggest caveat of this model, which makes it impractical is the possibility of avoiding the activity rule with false identities. A similar model with a few different properties (e.g expiring items) and assumptions can be found at [3]. It is also non truthful and the social efficiency bound shown there is one third in the worst case.
We will turn to introduce a less similar model, in order to capture the variety of mechanisms available.

4 Competitive Analysis of Incentive Compatible On-Line Auctions

In the previous model we have seen a non-incentive compatible (not truthful) auction, and analyzed its properties. Now we will examine a different mechanism and characterize incentive compatible auctions. For this model we will only describe the mechanism, the assumption, the definitions and the acquired results. The analysis and proofs for of the results can be found at [2]

4.1 The Model

This model settings is such where each different bidder arrive at different time (dynamic arrival) and the auction mechanism is required to make decision about each bid *as it is received*.
Such setting is similar to what happens in computerized settings such as bandwidth allocation for communication, and other resource allocation mechanisms, such as cache space or CPU time allocation.

In this model K identical items are sold in an auction, each bidders has a private valuation for each quantity of the items. (Note the difference from the unit demand as we have seen in the previous model)..
The bidder learns this valuation at a certain time (The moment it "needs" the resources) and must make the bid at that time.
The auction mechanism decides as the bid is received and before seeing future bids how many items to allocate to the bidder and at what price.
The main result for this model is the full characterization of what it takes for the model to be incentive compatible.

4.2 Formal Settings and Main Results

Each player has some positive benefit from receiving some quantity of the goods. This benefit is privately held. The marginal valuation of player i is $v_i(q)$. His total valuation for q goods is $\sum_{j=1}^{q} v_i(q)$.
We assume that all players have downward sloping marginal valuation function i.e $\forall_i \forall_q v_i(q+1) \leq v_i(q)$.
When player i receives q items and pays for them a total payment of P_i his utility is denoted $U_i(q, P_i) = \sum_{j=1}^{q} v_i(j) - P_i$
The bid is a non increasing function $b_i(q)$ of the form $b_i : [1, .., K] \rightarrow R$.
With these definition there is no guarantee that the bid function will equal the true valuation of the player, that is $b_i(q) \neq v_i(q)$ for some or all q and i. It is proved what it takes for this not to happen and for the auction to be incentive compatible. (for this mechanism)

Definition 4.1 *Supply Curve - an auction is called "based on supply curves" if before receiving the i'th bid it fixes a function $p_i(q)$ based on previous bids, and,*

1. *The quantity q_i sold to bidder i is the quantity q that maximizes the players utility i.e the sum $\sum_{j=1}^{q}(b_i(j) - (p_i(j))$*

2. *The price paid by agent i is $\sum_{j=1}^{q_i} p_i(j)$*

Theorem 4.2 *An auction of this model is incentive compatible if and only if it is based on supply curves.*

Proof: can be found at [2].

4.3 Competitive Analysis

This analysis defines the performance measurements (revenue and social efficiency) and exact meaning of a performance guarantee (competitiveness). For a worst case analysis, it is assumed that all marginal valuation are taken from some known interval $[\underline{p}, \bar{p}]$.

Definition 4.3 *The revenue of an auction A for a valuation sequence σ, denoted as $R_A(\sigma)$ is the total payment he received plus his valuation of the quantity he did not sell.*

$$R_A(\sigma) = \sum_i P_i + \underline{p}(k - \sum_i q_i)$$

where q_i is the quantity sold to player i and P_i is the total price paid by him.

Definition 4.4 *The social efficiency of an auction A for a valuation sequence σ, denoted as $E_A(\sigma)$ is the sum of the resulting utilities of all players including the auctioneer.*

$$E_A(\sigma) = \sum_i \sum_{j=1}^{q_i} v_i(j) + \underline{p}(k - \sum_i q_i)$$

The comparison done for the revenue and social efficiency is done regarding those of the offline Vickrey auction. Vickrey is used as the benchmark due to its popularity, due to it being the only auction with dominant strategies and due to the fact that it is always optimal in terms of social efficiency.

Definition 4.5 *An auction A is c-competitive with respect to the revenue if for every valuation sequence σ,*

$$R_A(\sigma) \geq R_{vic}(\sigma)/c$$

And similarly an auction A is c-competitive with respect to the social efficiency if for every valuation sequence σ,

$$E_A(\sigma) \geq E_{vic}(\sigma)/c$$

Theorem 4.6 *Every incentive compatible auction has a competitive ratio of at least $\Theta(log(\bar{p} \div \underline{p})$ with respect to either the revenue or the social efficiency.*

4.4 Model's Conclusions

This model was given to show the variety of models, and the variety of properties each may hold and may be analyzed. This model has further extensions and more results and can be found at [2] for full description.

References

[1] E. Segev and R. Lavi, *Efficiency Levels in Sequential Auctions with Dynamic Arrivals*

[2] N. Nisan and R. Lavi, *Competitive Analysis of Incentive Compatible On-Line Auctions*, Theoretical Computer Science 310(1), pp. 159-180, 2004.

[3] N. Nisan and R. Lavi, *Online Ascending Auctions for Gradually Expiring Items*, 17th ACM-SIAM Symposium on Discrete Algorithms (SODA'05), 2005.