Optimal bidding strategies for an entry deterrence model of online auctions

Jean-Noël Rivasseau
Department of Computer Science
University of British Columbia
Vancouver, Canada
jnriva@cs.ubc.ca

Abstract

In this project I explore sequential auctions with potential entry in later rounds. By considering a simple model, it can be shown that the main particularity of the equilibria taking place is that they try to *conceal* precious information to the new entrant. This project's main result consists in the proof of a somewhat stronger impossibility result for the equilibria of the considered model. A review of the various strategies available to bidders, depending on the assumptions made, is then provided.

Introduction

Rules for traditionnal auctions usually only had to deal with a fixed set of potential bidders, since these auctions were set up in one common place where all buyers and sellers were gathered. Some of these auctionning events offered multiple occurences of the same (or very similar) good, and thus, some satisfied bidders on early rounds would not attend the remaining rounds, or some late potential buyers could have missed the first few initial ones and still participate in the last. However, even if in these traditionnal auctions the number of bidders could vary from round to round, the set of bidders informed of all the bids that took place was constant or could only decrease in size. Usually, a new (late) arrivant didn't have any way to know what happened when he wasn't there yet. The history and records of all the auctions and bids submitted were completely unknown to him as he was considering his strategy for participating in the remaining auctions.

In today's online auctions environments, things have changed dramatically. The ever increasing popularity of eBay or other online auctionning sites attract thousands of new potential buyers every day. And these buyers can instantly get some information about previously held auctions, simply by looking at the currently running auctions and observing the prices. The sheer number of auctions for identical or very similar items allows newcomers to first have a look at the system, observe common prices (bids) submitted, and then, if they like what they see, they can decide to participate in future auctions of the goods they are interested in. On the other hand, if they are somehow disappointed in their initial review of the system, they may decide not to participate in auctions at that site.

It makes sense to model this kind of behavior, which seems extremely common. This project will thus examine, from a very theoretical point of view, sequential auctions with entry deterrence. This means we will consider several rounds of auctions of the same good, but some of the bidders are allowed to participate only after a certain number of initial rounds (but they get to observe the results of the bids submitted in the initial rounds). This is a reasonable choice to model the behavior of "newcomers" into the system, who may elect to first have a look at how things are going, then participate in later auctions. In this project, we will focus our attention on the strategies available to the initial bidders: they can be seen as more experienced users (or professional bidders) that have to define a best strategy in order to maximize their utilities against this perpetual flow of potential newcomers. It turns out that in order to achieve optimality, they have to resort to apparently strange (at first glance) bidding strategies designed to hide their true valuations; but this behavior is conditionnal on an assumption of perfect rationality for the potential entrants, which doesn't seem very realistic in practice.

This project report is structured as follows. The next section will introduce the formalism of the model we will be considering throughout the paper. We will discuss how realistic this model is, and how some of the assumptions we are making can be plausibly explained by some observable behavior on auction sites like eBay. The model we use is in fact the one considered by Andrei Bremzen in his recent paper on sequential auctions (Brezmen 2003). His main results will thus be recalled briefly, and a section will be devoted to an extension of his results: we will prove the equilibrium Bremzen constructed is the only possible one, given certain conditions. This is the most important result in this project. The last section will consist in a discussion of online bidding strategies, specifically related to the theoretical setting we presented. There are in the literature several papers discussing that topic (see for exemple (Porter & Shoham 2003), or (Easley & Tenorio 2001)). However, none of them specifically considers the possibility of modelling online bidding behavior with entry deterrence, and this project report hopes to give a presentation of what are the best strategies (and consequent equilibria), if we are making these assumptions of entry deterrence.

Setting and Model

Type of auction

We will be considering an ascending price auction in which there are two rounds, resolved sequentially. An ascending price auction is conceptually equivalent to a second price auction with sealed bids; the price paid by the winner is determined by the price at which the last other remaining bidder dropped out. Thus in effect, the drop out price selected by each bidder is (almost) equivalent to a pre-submitted sealed bid; and the winner ends up paying the second highest bid, like in a second price auction. See (Vickrey 1962) for more details about this "ratchet" strategy, and (Klemperer 1999) for a general discussion of the various common types of auctions. The advantage of sealed bids is that it allows an immediate (or very fast) resolution of the auction, because it doesn't require any communication between the seller and the bidders except for the initial sealed bids. In practice, sealed-bid auctions are used at eBay and other online marketplaces. However, because in eBay the auction lasts for a certain amount of time, bidders can revise their initial bids, and eventually resubmit a higher bid. Normally, there shouldn't be any incentive to revise a bid, since the first bid should only be deducted from the valuation a particular buyer has of a certain good, and this shouldn't change over time. But in fact, there may be advantages in doing so, because it can induce some signaling. This issue is tackled in (Easley & Tenorio 2001): bidders using jump bids as signaling bids can achieve higher expected profits. We will not be considering this issue here, and will always treat our auction as being strictly equivalent to a sealed-bid second price auction (without any opportunity to change initially submitted bids within the same round of auction).

On spying

Note that in this case of second price auction, if we were to consider a single auction, cheating would be possible for the seller (if he somehow dishonestly spied on the bids submitted by the various bidders), but not for the bidders: in (Porter & Shoham 2003), a paper dealing with the issue of cheating as a consequence to spying on others, the authors acknowledge this fact and thus examine the cheating opportunities for bidders in a first price sealed auction (which is not incentive compatible). However, with several sequential auctions it is interesting for a bidder to spy on the other bidders, even with a second price auction. This is essentially due to the fact that in our sequential auction setting, a second price auction is no longer an incentive compatible mechanism. Effects of cheating and spying could consequently be an interesting thing to examine in the future; for this project we will assume our marketplace is secure and doesn't allow anyone to spy on other bids before submitting its own. All information is disclosed through the normal auction mechanism, when the auction resolves.

Number of bidders

Let's get back to the description of our model. 3 bidders will be allowed to participate in the auction. This small number is enough to give a theoretical insight on the issues at stake and how they affect each bidder's strategy. The two first bidders are allowed to participate in both rounds; however the third one is a "newcomer" to the system, and can only participate in the second round. He gets to observe the results of the first round: so he observes what was the sealed-bid of the loser of the first round (on eBay, that's what he would also observe: the winning bid is the amount of the reserve price for the loser). All three participants have their valuation for the good drawn from an uniform distribution between 0 and 1 (this allows for easier computations as the valuations are already normalized like probabilities!), and are risk neutral. In each of the two rounds, one item of the good is sold. Bidders are only interested in at most one unit of the good, so the winner of the first round will quit and won't participate in the second round. The loser of the first round and Player 3 will thus bid against each other in the last round.

Tie-breaking rule and entry cost

As we will see, ties will play a crucial role in our equilibrium, so we need a tie-breaking rule. We choose the simplest one: if a tie occurs between two bids, the good will be sold randomly (with probability $1/2$) to one of the players. Finally, Player 3 incurs an entry cost for entering the auction in the second round. This cost $c > 0$ will in fact never be greater than $1/2$, or else Player 3 would never enter (see (Brezmen 2003) for details). Sometimes this entry cost will prohibit Player 3 from entering; in this case the loser of the first round can get the item for free (pays 0). In a more realistic setting, the loser would actually pay the reserve price of the seller for the item (and bids wouldn't start at 0 either but at this reserve price). But from a mathematical perspective, the two situations are exactly equivalent, and it is easier to consider that the "best possible result" for the buyer is to pay 0, so we will adopt that notation.

Empirical justification of the model

It is interesting to discuss why, in a realistic online marketplace, player 3 would incur an entering cost and not the two first bidders. After all, participation in eBay is free for every potential bidder. However, if player 3 is indeed new to the system, this assumption makes sense. There are in fact actual costs for participating in online auctions: you would have to buy a computer or get access to one, pay the connection fees, then register for an account at eBay, learn how the system works, etc... All of this takes time and/or money and thus can be considered as a cost for the potential entrant (plus, the fear of paying money to a dishonest seller and never getting back anything in return prevents many people from entering online auctions - this can also be modelled in the cost). This is especially true for newcomers; experienced users already have made the investment of learning how the sytem works, and thus incur very little costs at each new auction in which they participate. But for a person moderately interested in online auctions, that is just having a look at an eBay web page describing an auction, to see "if all this online auctions frenzy can be truly worth it", the entry cost can be very real. The decision to take the time and effort to actually enter "a next auction" will very often depend on the initial impression an user gets of the system, and that initial impression is essentially conditionned on the prices observed, like in our model. This entry cost is all the more realistic as the relative price for the object is small, because the entry cost in our setting actually represents a fraction of the item cost. Thus it is unrealistic to consider an entry cost of 0.4 for a $\$100,000$ car (this would translate to a cost of $\$40,000$); but for a $\$10$ or even a $\$100$ item, our model is quite plausible.

A step function equilibrium

In this section I will recall briefly the main result presented by Bremzen. For the details of all the proofs, refer to (Brezmen 2003). He constructed an equilibrium for the three bidders in the previously described setting. Such an equilibrium consists of:

- a strategy that the two first players will follow in the first round,
- an entry decision for the third player (conditionned, of course, on the observed result of the first round),
- and finally, a strategy followed by the two bidders on the last round, if player 3 enters.

It should be noted that bidding strategies for the two first players on the first round have to be symmetric (thus we only have to define one strategy), due to the inherent symmetry of our setting. Similarly, the bidding strategies for the second round, if player 3 enters, are symmetric, and in fact are trivial to define, because we are dealing with a simple second price auction (with two players). It is then well known (see (Shoham 2003)) that the optimal strategy is to bid their respective true valuations.

The equilibrium described by Bremzen results, for the initial player strategy, in a submitted bid that is a step function of the valuation. More precisely, up to a certain valuation v_{lim}, the optimal choice is to bid 0; if the valuation is greater than v_{lim}, then the player should bid a fixed amount b^*. This strategy is presented in Figure 1.

This, combined with a decision entry for the third player, defines a Nash-Equilibrium of our auction, given that $c > c_{lim} \approx 0.1467$ and $c < \frac{1}{2}$. The entry strategy for Player 3 with valuation v, that results in an equilibrium, is the following:

- if the observed price is b^*, enter if $v > v_{lim} + \sqrt{2c(1 - v_{lim})}$;
- if the observed price is 0, enter if $v > \frac{v_{lim}}{2} + c$;
- otherwise, enter if $v > c$.

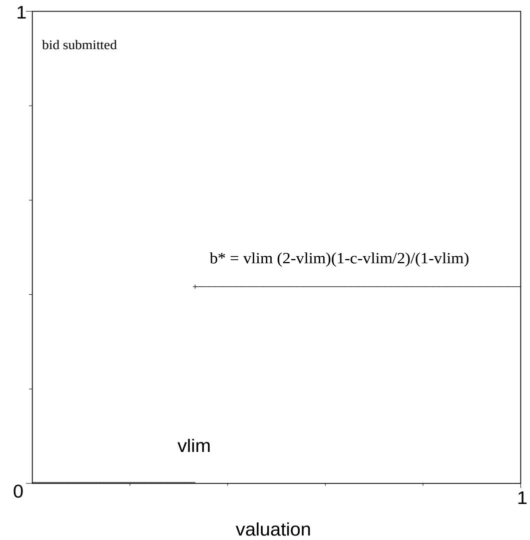

Figure 1: An equilibrium strategy as a 2 step function

It should be noted, as pointed out by Bremzen, that there is a degree of freedom in this equilibrium, in the choice of v_{lim}. Not all values are of course possible for v_{lim}, but there is an interval of possible values. b^*, on the contrary, is binded to v_{lim} and c by the following formula:

$$b^* = \frac{v_{lim} \cdot (2 - v_{lim}) \cdot (1 - c - v_{lim}/2)}{1 - v_{lim}}$$

So there is no degree of freedom here; actually, b^* is chosen so that Player 1 or 2 is indifferent between bidding 0 or b^*, if he has value v_{lim}. One important remark is that if Player 3 observes price b^*, he may correctly deduce the value chosen for v_{lim}. However, if he only observes price 0, he can't get the value for v_{lim}, which makes his entry decision (enter if $v > \frac{v_{lim}}{2} + c$) problematic. Similarly, the two first players have to choose the same exact v_{lim} in order to achieve an equilibrium. So, even if this set of strategies do define a Nash-Equilibrium, it seems rather hard to achieve in practice. We will discuss these problems in the last section.

Intuitively, this step function equilibrium occurs as a consequence of the willingness from the first bidders to hide their correct valuations to the new entrant. If the new entrant is able to determine exactly the valuation of the loser, he will enter as soon as his own valuation is greater than the loser's valuation added to his entry cost. If there is uncertainty about that valuation, he will enter if his own valuation is greater than the *expectation* of the loser's valuation and his entry cost. It is thus in the initial bidders interest to keep the entrant ignorant of their exact valuations.

Equilibria for other values of c can also be found. If $c = 0$, it is easy to understand that it is in Player 3 best interest to enter all the time, and this leads to a drop out strategy for the first round of $b(v) = v - \frac{v^2}{2}$. If $c > 1/2$, we already said earlier that Player 3 will never enter, thus the equilibrium is for both initial players to always bid 0. If $0 < c < c_{min}$, then a step function equilibrium is also possible. In this case of small entry costs, there are no longer only two steps, but k ones. k will go up when c goes to 0, but for a given c, it can be proven (again, refer to (Brezmen 2003) for details) that a step function with a finite number of steps will constitute a Nash-Equilibrium for our problem.

The only possible equilibrium strategies are step functions

This section presents the main result of the project: we will show that step functions, as described in the previous paragraphs, are the only equilibria existing for our problem. This was a question remaining to be answered, according to Bremzen's conclusions. He showed that for any $c > 0$, there existed no equilibrium bidding function $b(v)$, for the initial round, that was strictly monotone in valuation. I will go one step further and claim that:

Lemma: For a sufficiently high entry cost $c > c^*$, there exists no equilibrium bidding function $b(v)$ that is strictly monotone in valuation on any interval.

Proof: Let's assume the contrary, and let $[x_1, x_2]$ be an interval in which $b(v)$ is strictly monotone. Then it is clear that on that interval, player 3 will be able to correctly deduce the valuation of the loser (because there is a bijection on that interval between the bid and the valuation). He will enter if his own valuation is greater than $v+c$.

When player 1 decides the amount he is going to place on his sealed bid, he assumes (as part of the equilibrium) than player 2 is following $b(v)$. If his own valuation is on the interval $[x_1, x_2]$, he may consider deviating by a small amount, and claim he is of type v_{false} which is still in the interval $[x_1, x_2]$. His expected payoff will then be

$$Payoff(v, v_{false}) = W + \int_{x-1}^{v_{false}} (v - b(x)) \cdot dx + v(1 - v_{false})(v_{false} + c)$$

Here, W represents a constant payoff (not depending on v_{false}) that player 1 gets if player 2 is going to bid lower than the interval and consequently, player 1 is going to win the item on the first round. The integral represents the payoff if player 2 is going to bid on the $[x_1, x_2]$ interval but lower than player 1 (thus lower than v_{false}). Finally the last term represents the payoff if player 2 bids higher than player 1. Then player 1 goes on to the next round (with probability $(1 - v_{false})$), with player 3 believing he has valuation v_{false}. If we make v_{false} sufficiently close to v, so that $|v - v_{false}| < c$, then player 1 always lose if player 3 enters, and thus he receives something by going to the second round only if player 3 does not enter (probability $(v_{false} + c)$), and in this case his payoff is v.

$Payoff(v, v_{false})$ must reach its maximum at $v_{false} = v$ since it is an equilibrium function (and thus small deviations won't bring in more profits). Thus by differentiating $Payoff(v, v_{false})$ in respect to v_{false} we obtain the condition:

$$0 = v - b(v) - v(v + c) + v(1 - v)$$

$$b(v) = 2v - vc - 2v^2$$

To be monotone on the interval $[x_1, x_2]$, we again differentiate one time and obtain the condition:

$$v < \frac{2 - c}{4}$$

Thus x_2 has to be lower than $\frac{2-c}{4}$. We proved than $b(v)$ can't be strictly monotone on the interval [$\frac{2-c}{4}$,1]. Still, there could be an arbitrary number of intervals contained in $[0, \frac{2-c}{4}]$ where $b(v)$ could be strictly monotone and of the form (this is a requirement) $b(v) = 2v - vc - 2v^2$. We now have to prove that this is impossible.

Let's pick the interval, where $b(v)$ is monotone, *closest* to $\frac{2-c}{4}$. This interval is $[x_1, x_2]$ with $x_2 < \frac{2-c}{4}$. By construction, we know that after x_2, $b(v)$ is a step function (because it can be strictly monotone nowhere). Let's assume it has only a single step after x_2, with a bidding price of b^*; given that $c > c^*$, we can always assume that. Since the number of possible steps is linked to the value of c, for a sufficiently high c, we can be sure that we have at most two steps (see (Brezmen 2003) for a proof of that). The single step assumption after x_2 is not crucial here and is only done to simplify derivations. I am confident that a more complete proof could examine the case of small entry costs c and conclude to the same result (thus we would prove that for any c, there exists no equilibrium bidding function $b(v)$ that is strictly monotone in valuation on any interval).

In order to construct an equilibrium, the initial bidder (player 1) with valuation x_2 must be indifferent between bidding $2 \cdot x_2 - x_2 \cdot c - 2 \cdot x_2^2$ (monotone part) and bidding b^* (step part). So:

$$Payoff_{monotone}(x_2) = Payoff_{step}(x_2)$$

$$W + v(1 - x_2)(x_2 + c) = W + \frac{(x_2 - b^*)(1 - x_2)}{2} + \frac{(1 - x_2) \cdot \alpha}{2}$$

W is a constant and represents (in both cases) the payoff if we win the first round because player 2 bids lower than player 1. ($\frac{(x_2-b^*)(1-x_2)}{2}$) is the case when both initial players are tied at the step, and ($\frac{(1-x_2)\alpha}{2}$) represents the payoff if player 1 goes to the second round (α is the probability of player 3 entering: if player 3 enters, player 1 automatically loses since he is at the beginning of the step with a valuation of x_2). After derivations we obtain:

$$b^* = (1+\alpha) \cdot x_2 - 2x_2 \cdot c - 2 \cdot x_2^2$$

This b^* has to be more than the price $b(x_2)$ corresponding to the monotone part. Thus:

$$(1+\alpha) \cdot x_2 - 2x_2 \cdot c - 2 \cdot x_2^2 \geq 2x_2 - x_2 \cdot c - 2x_2^2$$

This gives us:

$$c \leq \alpha - 1$$

which is impossible, given that $c > 0$ and $\alpha \leq 1$. Thus over the interval $[x_1, x_2]$ we can't have a monotone $b(v)$, since it is impossible for such an equilibrium $b(v)$ to "link" to the next step. The rest of the proof is straightforward: since we have proved that the monotone interval which was closest to the steps defining the end of our equilibrium strategy can't exist, obviously we can consider the second next closest interval, which becomes the closest since we have just proven the previous one cannot exist. By recursion, we conclude that no interval where $b(v)$ is strictly monotone can exist. Thus our lemma.

Our lemma proves than equilibria of the type described in Figure 2 can't exist (which was something remaining to be answered, according to Bremzen). Since an equilibrium bidding strategy can't be strictly monotone anywhere, obviously the only possible equilibria are step functions, that have a finite number of discontinuities and are constant over intervals.

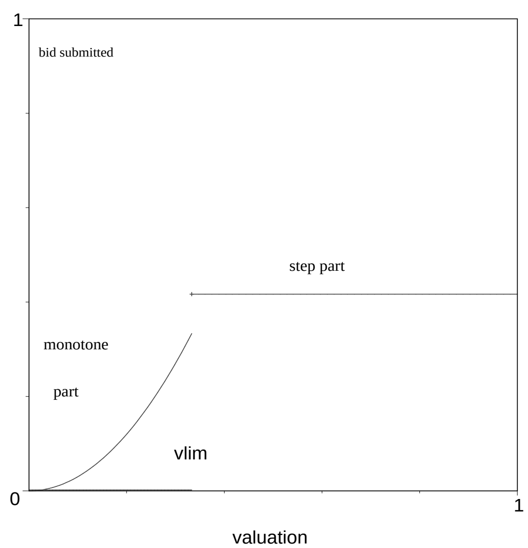

Figure 2: An impossible equilibrium strategy

Discussion of online bidding strategies for the entry deterrence model

It is clear that following an optimal strategy can significantly affect the profits of the initial bidders, as shows the following example.

Example. *Let's consider an entry cost of $c = 1/3$. We are interested in computing the average payoff of an initial bidder. By following the simple "honest" strategy of bidding your valuation for the first round, you get an average payoff of 0.3 (if your valuation is drawn uniformly from 0 to 1). By following the equilibrium strategy, with a $v_{lim} = 1/3$, the average payoff is 0.384, more than 25% more.*

On rationality

This may suggest than experienced buyers at online marketplaces should try to follow this equilibrium strategy, actually "teaming up" with other established experienced users in order to limitate the entrance of newcomers. But there is a catch in this analysis. The equilibrium discussed is a Nash-Equilibrium, which means that if all bidders are perfectly rational and aware than the others are playing this strategy, then it is in their best interest to play in this way. However, newcomers to the system cannot reasonably be considered as rational: they are very unlikely to perform this analysis and play in consequence. For exemple, they are likely, if they observe very low prices in an online auction, to try to enter subsequent auctions, because they think it will be profitable. But in fact, they may not realize that low prices are due to the fact that initial bidders want to hide their correct valuations; and these newcomers may be facing a tougher opposition than they expect on the auctions they enter, after having paid the entry cost. Like we also pointed out earlier, there is also a degree of freedom in the choice of v_{lim}, which may result in a lack of coordination for the initial bidders, and additional uncertainty for the third player. This will most probably lead to a loss of profit, although it would have to be investigated.

Thus, it would be interesting not to focus on the Nash-Equilibria of the entry deterrence model for auctions, but to find what is the best strategy for initial bidders, given that the newcomers will act "irrationally". After all, we have just seen that computing the equilibrium is no small task; realistically, almost nobody can be expected to perform this kind of pushed analysis. A newcomer can plausibly be assumed to play always naively, that is, enter if his valuation is more than the sum of the observed price and the entry cost.

Making that assumption dramatically changes the equilibrium for the two first players. We are back to a model where information hidding is not important and even undesirable. This is because since the newcomer doesn't think information hidding is occuring, any "steps" we may have in our equilibrium function now have an adverse effect for the initial bidders. The newcomer will assume initial bidders have the lowest possible valuation of the step. In these conditions, if an equilibrium exists (we can't know for sure that an equilibrium exists, because we are not dealing with a finite normal form game), it is probably very different from the ones we presented earlier.

On cooperation

One other interest thing to note is that if initial bidders cooperate (but don't follow an equilibrium strategy), they can achieve impressive profits at the expense of the third player. A very simple cooperating strategy is to always bid 0. This gives, with $c = 1/3$ and against a third player bidding naively, an average payoff of 0.358, which is already very good. Against a "rational" third player aware of the strategy, this gives a payoff of 0.458 ! This illustrates two facts: a player behaving rationally is actually better for the initial bidders, because in that case, information hiding can very easily forbid him from entering and thus leads to huge profits for the initial bidder. The second fact is that cooperation between initial players gives better rewards than following the equilibrium strategies, especially in the case of rational newcomers. But, this is of course very risky: initial players will be tempted to cheat on their partner and bid a very small positive amount, obtaining the maximum possible average payoff of 0.5, regardless of whether the newcomer is rational or not ...

What is worth noting is that in such a case, where initial bidders would agree to cooperate to keep their profit high against new comers, their best strategy is not to bid very high prices in order to discourage potential entrants, but instead to agree on very low prices for the auctions in which they are the only ones to participate. This is because our model is fairly simple and considers only two rounds. If more rounds were available, it would probably be realistic to assume that, upon observation of the price of initial auctions, a new entrant could decide either to never enter, or to enter and occur the entry cost only once, then be able to participate at every remaining auction. In that setting, intuition suggests that initial bidders willing to cooperate would then focus on an agressive bidding policy, in order to permanently eliminate potential rivals, since every new entrant would stay forever in the system without incurring additional costs.

Conclusion

In this project report we have examined a plausible model for online auctions, based on entry deterrence and described extensively in (Brezmen 2003). The crucial assumption for this model is that whenever several auctions of the same good are available, sequentially, there will be new potential buyers for the later auctions, but these newcomers incur an entry cost, reflecting the effort needed to register at the online marketplace.

We have extended Bremzen's impossibility result, and proved that the only possible equilibria (if all bidders are rational) in this setting are for the submitted bids to be step functions of the valuations. However, and although the initial bidders would prefer the potential newcomers to behave rationally in their entry decision, this would probably not occur in practice. Thus actual optimal bidding strategies may not be the ones that theoretical Nash equilibria would suggest. Also, if initial bidders can agree to a form of cooperation, they can achieve much higher profits, as is often the case: in the classic Prisonner's Dilemna game, cooperation can lead to better payoffs for both players than the Nash-Equilibrium strategy.

This project opens the way for a lot of future work. First, the main proof in this report could probably be extended to prove a stronger and more general result (withdrawing the assumption we made on the entry cost c by mean of precaution). Another very interesting result would consist in computing an equilibrium for the initial bidders strategies, given that the new entrant will behave naively. If such an equilibrium does exist, it is probably on this kind of equilibrium that experimented or professional bidders on eBay should rely in order to maximize their profits against newcomers. It would be interesting to empirically test this.

Finally, the model could be extended in order to better capture long-term behavior in online marketplaces. Considering more than 3 players is interesting and already tackled in (Brezmen 2003). However, a much more complex change would be to examine the possibility of more than 2 auctionning rounds. More generally, sellers and auction-designers strategies, in sequential auctions with potential entry in later rounds, could also be considered.

References

Brezmen, A. 2003. Sequential auctions with entry deterrence.

Easley, R. F., and Tenorio, R. 2001. Jump Bidding Strategies in Internet Auctions.

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

Porter, R., and Shoham, Y. 2003. On cheating in sealed-bid auctions. In *Proceedings of the 4th ACM conference on Electronic commerce*, 76–84. ACM Press.

Shoham, Y. 2003. Multi Agent Systems (draft version).

Vickrey, W. 1962. Auctions and bidding games. In *Recent Advances in Game Theory*. Princeton University.