Modeling an online auction through the lens of Game Theory:

What is my best bidding strategy for Sotheby's *Natively Digital: Oddly Satisfying* online auction for Lara Zankoul's *High Maintenance* NFT?

BASC0017 Interdisciplinary Game Theory

Ema Pop

March 2023

Introduction:

Sotheby's are having their annual NFT auction this March called Natively Digital. This online auction has the most renowned names in the digital space. Sotheby's broke several records in 2021, having surpassed $7.3bn in sales. The highest total in company history, due to the 95% of sales being online. (Cassady, 2021) One of this year's theme is Oddly Satisfying which presents the biggest 3D artists and animators. Renowned auction houses representing digital artists is incredibly important because they place digital art pieces within an institutionalised context of art history. Therefore, legitimising the digital creator economy (Pizzardello, 2021). Satisfying 3D loops have found themselves to be a staple of the contemporary internet culture which is why work that focuses on quantifying those aesthetics is even more important (Bak, 2022). My work aims to bring more diversity into this digital art space by breaking down the strategies to bid for an online NFT auction at Sotheby's by providing an introduction to potential collectors. Furthermore, democratising the decentralised art world and using a theoretical model to calculate the best responding bid (b_i^*), given the bidder's valuation (v_i).

Defining the terms of this auction and the assumptions:

Let's define a **valuation** (v_i) as *"each individual bidder's valuation"*, the term **NFT** as *"a non-fungible token established on a blockchain"* which in this case would be the item offered for sale for Lot 20 (Figure 1). Every bidder has a **maximum bid** that is unknown to the other bidders defined as *"the highest amount someone would like to bid for a lot that is only known to them"*. Therefore this is modeled as a game of incomplete information with a Bayesian Nash Equilibrium (b_i^*) (Blümel et al., 1986). The **highest bid** is defined as *"the current biggest bid from all placed bids"*. Let's take Lot 20 for example (Figure 1). It has an estimation of $1,000 - 2,000$ EUR and no **reserve** (Sotheby's, 2023). The estimation is done by Sotheby's according to the speculation of future value and comparison to similar art pieces. When a piece has a reserve/starting bid, it means that the Lot can only be sold if the winning bid has been met a minimum price decided beforehand. In this case, no reserve means that there is no starting bid and the NFT will still go to the winning bid, may that be below the auction house's expectation.

If I value Lara Zankoul's NFT at $v_i = 1100$ EUR, I wouldn't want to place my bid at 1100 EUR, 6 days before the auction ends. Because I could **signal** other bidders that we value the NFT higher than the current highest bid, increasing other bidder's valuations. To avoid this, I will hold my maximum bid until almost the end of the auction. Placing bids in the final moments of an auction is called **sniping** (Backus et al., 2015). I will also assume that most bidders are rational and will do this because this is an online auction. Currently, there are 4 bids placed with the highest bid being 200 EUR. To bid now, I would only bid 300 EUR, respecting the automatic bid increments shown in Figure 2. Bidding less than my valuation is called **shading** the bid. I'm also assuming that every bidder would be doing this in order to get the biggest payoff. My payoff, (U_i) is my valuation (v_i), minus the price I had to pay for the NFT if I won (p): $U_i = v_i - p$. We also assume that no bidder would be bidding more than their valuation in order to avoid the **winner's curse**. Which is when the winner pays more, than the art piece is worth.

Rules of the auction:

The auction for Lot 20 ends on the 24th of March at 3:20pm GMT. The bidder with the highest bid placed before that time wins the NFT. Because every bidder can see how many bids there are, and what the highest bid is, this could be called an open English auction. The rules of an English auction are as follows: All bidders start the auction at the lowest price. They incrementally keep bidding until they choose to drop out when the highest price is too high for them. Once they drop out, they cannot reenter the auction. The highest bid wins and pays the second highest price (O'Reagan, 2005). Furthermore, all bidders can see when another bid is placed and are in a scenario where the second highest price and highest price would have an insignificantly small price difference e.g. eBay auction (Myerson, 1981). However, in our case, in order to place a bid, Sotheby's asks the bidders to only place their **maximum bid** in their platform. It will be executed by Sotheby's at the lowest price possible, and never for more than the maximum bid amount the bidder indicates.

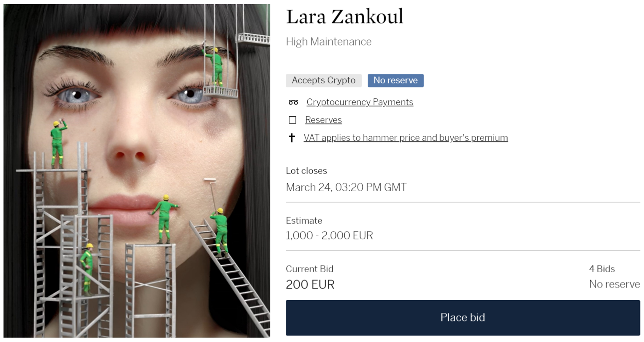

Figure 1: Lara Zankoul - High Maintenance, Lot 20, screenshot taken on March 18th at 6:45pm GMT

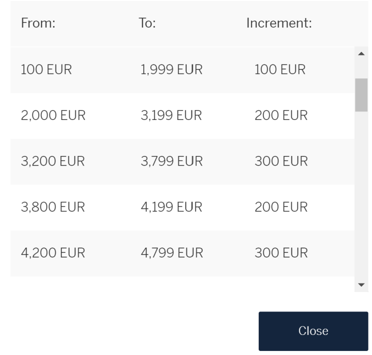

From:	To:	Increment:
100 EUR	1,999 EUR	100 EUR
2,000 EUR	3,199 EUR	200 EUR
3,200 EUR	3,799 EUR	300 EUR
3,800 EUR	4,199 EUR	200 EUR
4,200 EUR	4,799 EUR	300 EUR

Figure 2: Bid increments table

If two players enter the same maximum bid, the bid which was entered first will have priority (Sotheby's, 2022). This was an attempt to decrease sniping form bidders. Sniping usually decreases trust in the platform which is why this method of placing bids reduces the number of bids placed in the last few minutes. However, the maximum bids can be increased at any point. Therefore, leaving room for every maximum bid placed to be increased in the last essential moments, still causing some sniping (Sotheby's, 2023).

Why is this a First-Price Sealed Bid and not a Second-Price Sealed Bid or English Second-Price Open auction?

It could be argued that the bids are simultaneous if we look at the time frame right before the end of the auction when maximum bids can be placed. **Maximum bids** will be coming in the last essential moments, even seconds or milliseconds, so fast that we can assume the website will have trouble displaying all the bids in true real time (because we are assuming bidders will **snipe** their **maximum bids**). Which makes this auction more similar to a Sealed Bid auction where the bids are coming in simultaneously and it is very difficult to know what the exact highest bid is. Hence, in the last moments and last moments only, it is **not** an English open auction (Levin, 2008). Furthermore, the increments between the bids are already set and create a significant difference as the bids enter different price brackets (Figure 2). Meaning that the winning bid will pay the highest price in full. Hence we are choosing a First-Price Sealed Bid auction model and not a Second-Price Sealed Bid.

Why should assume it is common value instead of private value?

Lot 57 for example, is a haighly valued NFT and has a reserve of $60,000$ EUR. Sothebey's valuation is that someone will pay at least that price. We can assume that a bidder has a personal attachment to the art piece because they saw it live displayed at Printworks. Their personal valuation will be different compared to places a bid because it is auctioned at Sotheby's and doesn't know the artist story. For that, the equilibrium strategy is to bid exactly the same as the personal valuation, given that it would be the winning bid. This is considered a weakly dominant strategy. $b_i^* = v_i$. However, I am assuming that bidders for Lot 20 will have a common value because there is a more specific estimation, it is not a highly valued NFT and the bids are similar but not similar enough for this to be a Second Sealed Bid auction. For a this First Sealed Bid and common value auction, we can use the cumulative distribution function (CDF) to find the Bayesian Nash Equilibrium in this imperfectly and incomplete game (Peters, 2008) $\therefore F(V) = V$ valuation is independently and uniformly distributed on [0,1].

Calculating Payoffs for our First-Price Sealed Bid auction:

The utility of my bid b_i given my valuation $v_i = 1100$ EUR is: $U_i = (v_j - b_i) \times Pr\{iwin\}$. Let's say Bidder j has valuationv_j. Their utility would be $U_j = \beta(v_j) = rv_j$ where r is a fraction showing that Bidder j is also shading their bid.

Let's calculate the probability of me winning:

$Pr\{iwin\} = Pr\{b_i \geq \beta(v_j \neq i)\}$ where my bid is the highest when the auction ends and where $v_j \neq i$ is every other bidder other than me. Now, for independent events: $\prod_{j \neq i} Pr\{b_i \geq \beta(v_j)\}$ (Prof. Oyakawa, 2017) where the probability of the joint event is the product of the probabilities of individual events. For the uniformity assumption: $Pr\{b_i \geq \beta(v_j)\} = Pr\{b_i \geq \beta(v_k)\}$ where the v_s are uniform random variables. The valuations are uniformly distributed, therefore each bidding strategy is the same. $\therefore$ $b_i > v_j$ is the same as $b_i > v_k$. Therefore, $[Pr\{b_i \geq \beta(v_j)\}]^{N-1}$ because β_j and β_k are drawn from the same distribution where v_s are uniform. Where N is the nr of bidders (multiply the same probability N-1 times). Now, we write the probabilities as cumulative distribution function where $b_j = rv_j$:

$$\left[Pr\{V_j \leq \frac{b_i}{r}\}\right]^{N-1} = F\left[\frac{b_i}{r}\right]^{N-1} = \left[\frac{b_i}{r}\right]^{N-1} = Pr\{iwins\}$$

Now we plug back into the utility function:

$$EU_i = (V_i - b_i)\left[\frac{b_i}{r}\right]^{N-1}$$

$$= V_i\left[\frac{b_i}{r}\right]^{N-1} - \frac{b_i^N}{r^{N-1}}$$

Now calculate our best response by differentiating the utility function to find the maximum:

$$\frac{\mathrm{d}U_i}{\mathrm{d}b_i} = \frac{v_i}{r}\left(\frac{b_i}{r}\right)^{N-2}(N-1) - \frac{N}{r^{N-1}}(b_i)^{N-1} = 0$$

Solving for b_i^* (Prof. Oyakawa, 2017):

$$\left(\frac{1}{r^{N-1}}\right) v_i[b_i]^{N-2}(N-1) = \left(\frac{1}{r^{N-1}}\right) N[b_i]^{N-1} = 0$$

$$\therefore v_i[b_i]^{N-2}(N-1) = N[b_i]^{N-1}$$

$$\therefore \frac{v_i}{b_i}(N-1) = N$$

$$\therefore b_i^* = \frac{N-1}{N}v_i$$

Conclusion:

Therefore, the best response for my bid b_i, assuming there will be 10 bidders $(N = 10)$ is to bid $b_i^* = 990$ EUR. Respecting the increments, I would place my maximum bid at 1000 EUR according to the Bayesian Nash Equilibrium. If more people are bidding, I should shade my bid less, resulting in the best response b_i^* being closer to my value v_i. Any slight variation in the auction rules could change the model of the game completely and therefore, the bidding strategy. If this NFT had a reserve, the Bayesian Nash Equilibrium would be $b_i^* = v_i$ because it would be private value. My result also only works if, and only if, **ALL** stated assumptions are actually true. In reality, every auction is a mixture of common and private value bidder types, and mostly an English auction. Therefore, modeling it to perfection through an algorithm is challenging but provides an insight into the mechanisms behind online auctions.

References

M. Backus, T. Blake, D. V. Masterov, and S. Tadelis. Is sniping a problem for online auction markets? In *Proceedings of the 24th International Conference on World Wide Web*, pages 88–96. International World Wide Web Conferences Steering Committee, 2015. ISBN 978-1-4503-3469-3. doi: 10.1145/2736277.2741690. URL https://dl.acm.org/doi/10.1145/2736277.2741690.

G. Bak. The aesthetic measure of an NFT, 2022. URL https://www.rightclicksave.com/article/the-aesthetic-measure-of-an-nft.

W. Blümel, R. Pethig, and O. von dem Hagen. The theory of public goods: A survey of recent issues. 142 (2):241–309, 1986. URL https://www.jstor.org/stable/40750871.

D. Cassady. Sotheby's sales for 2021 surpass $7.3bn, the highest total in company history, 2021. URL https://www.theartnewspaper.com/2021/12/15/sothebys-73bn-sales-2021-highest-total-company-history.

J. Levin. Auction theory, 2008. URL https://web.stanford.edu/~jdlevin/Econ\%20286/Auctions.pdf.

R. B. Myerson. Optimal auction design. 6(1):58–73, 1981. ISSN 0364-765X, 1526-5471. doi: 10.1287/moor.6.1.58. URL http://pubsonline.informs.org/doi/10.1287/moor.6.1.58.

R. T. O'Reagan. A look at the game theory of online auctions. 2005. URL http://hdl.handle.net/2345/406.

H. Peters. *Finite Games with Incomplete Information*, pages 59–71. Springer Berlin Heidelberg, 2008. ISBN 978-3-540-69290-4 978-3-540-69291-1. URL http://link.springer.com/10.1007/978-3-540-69291-1_5.

M. Pizzardello. The new aesthetics of digital art, 2021. URL https://www.michaelpizzardello.com/.

P. G. Prof. Oyakawa. First-price sealed-bid auctions. pages 1–9, 2017. URL https://cs.brown.edu/courses/cs1951k/lectures/2020/first_price_auctions.pdf. line 62.

Sotheby's. How to bid online, 2022. URL https://www.sothebys.com/en/buy-sell#how-to-bid-online.

Sotheby's. Guide for buyers at auction, 2023. URL https://www.sothebys.com/en/docs/pdf/guide-for-buyers-at-auction-paris-en-09-02-2023.pdf.