An Introduction to Auction Theory for Undergraduate Students

Félix Muñoz-García*
School of Economic Sciences
Washington State University
Pullman, WA 99164

September 25, 2012

Abstract

This paper provides a non-technical introduction to auction theory. Despite the rapidly expanding literature using auction theory, and the few textbooks introducing it to upper-level Ph.D. students, most undergraduate textbooks do not cover the topic, or present short verbal descriptions about it. This paper offers an introduction to auctions, emphasizing their common ingredients, analyzes optimal bidding behavior in first- and second-price auctions, and finally examines bidding strategies in common-value auctions and the winner's curse. Unlike graduate textbooks on auction theory, the paper only assumes a basic knowledge of algebra and calculus, and uses worked-out examples and figures, thus making the explanation attractive and understandable for most economics and business majors.

Keywords: Auction theory; First-price auction; Second-price auction; Common-value auctions; Bidding strategies.

JEL classification: D44, D8, C7.

*Address: 103G Hulbert Hall, Washington State University. Pullman, WA 99164. E-mail: fmunoz@wsu.edu. Phone: (509) 335 8402. Fax: (509) 335 1173.

1 Introduction

Auctions have always been a large part of the economic landscape, with some auctions reported as early as in Babylon in about 500 B.C. and during the Roman Empire, in 193 A.D.[1] Auctions with precise set of rules emerged in 1595, where the Oxford English Dictionary first included the entry; and auctions houses like Sotheby's and Christie's were founded as early as 1744 and 1766, respectively. Commonly used auctions nowadays, however, are often online, with popular websites such as eBay, with US$11 billion in total revenue and more than 27,000 employees worldwide, which attracted the entry of several competitors into the online auction industry, such as QuiBids recently.

Auctions have also been used by governments throughout history. In addition to auctioning off treasury bonds, in the last decade governments started to sell air waves (3G technology). For instance, the British 3G telecom licenses generated Euro 36 billion in what British economists called "the biggest auction ever," and where several game theorists played an important role in designing and testing the auction format before its final implementation. In fact, the specific design of 3G auctions created a great controversy in most European countries during the 1990s since, as the following figure from McKinsey (2002) shows, countries with similar population collected enormously different revenues from the sale, thus suggesting that some countries (such as Germany and the UK) better understood bidders' strategic incentives when participating in these auctions, while others essentially overlooked these issues, e.g., Netherlands or Italy.

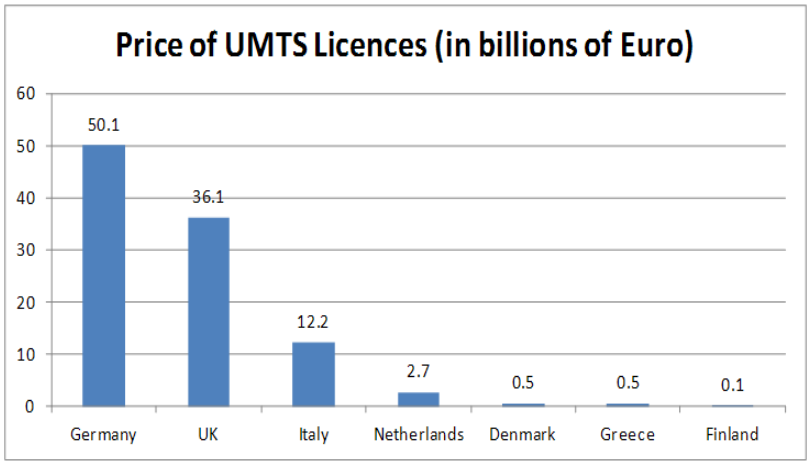

Fig 1. Prices of 3G licences.

Despite the rapidly expanding literature using auction theory, only a few graduate-level textbooks about this topic have been published; such as Krishna (2002), Milgrom (2004), Menezes and

[1] In particular, the Praetorian Guard, after killing Pertinax, the emperor, announced that the highest bidder could claim the Empire. Didius Julianus was the winner, becoming the emperor for two short months, after which he was beheaded.

Monteiro (2004) and Klemperer (2004). These textbooks, however, introduce auction theory to upper-level (second year) Ph.D. students, using advanced mathematical statistics and, hence, are not accessible for undergraduate students. In addition, most undergraduate textbooks do not cover the topic, or present short verbal descriptions about it; see, for instance, Pindyck and Rubinfeld (2012) pp. 516-23, Perloff (2011) pp. 462-66, or Besanko and Braeutigam (2011) pp. 633-42.[2] In order to provide an attractive introduction to auction theory to undergraduate students, this paper only assumes a basic knowledge of algebra and calculus, and uses worked-out examples and figures. As a consequence, the explanations are appropriate for intermediate microeconomics and game theory courses, both for economics and business majors. In particular, the paper emphasizes the common ingredients in most auction formats (understanding them as allocation mechanism). Then, it analyzes optimal bidding behavior in first-price auctions (section three) and in second-price auctions (section four). Finally, section five examines bidding strategies in common-value auctions and the winner's curse.

2 Auctions as allocation mechanisms

Consider N bidders who seek to acquire a certain object, where each bidder i has a valuation v_i for the object, and assume that there is one seller. Note that we can design many different rules for the auction, following the same auction formats we commonly observe in real life settings. For instance, we could use:

1. First-price auction (FPA), whereby the winner is the bidder submitting the highest bid, and he/she must pay the highest bid (which in this case is his/hers).
2. Second-price auction (SPA), where the winner is the bidder submitting the highest bid, but in this case he/she must pay the second highest bid.
3. Third-price auction, where the winner is still the bidder submitting the highest bid, but now he/she must pay the third highest bid.
4. All-pay auction, where the winner is still the bidder submitting the highest bid, but in this case every single bidder must pay the price he/she submitted.

Importantly, several features are common in the above auction formats, implying that all auctions can be interpreted as allocation mechanisms with two main ingredients:

a) An *allocation rule*, specifying who gets the object. The allocation rule for most auctions determines that the object is allocated to the bidder submitting the highest bid. This was, in fact, the allocation rule for all four auction formats considered above. However, we could assign the object by using a lottery, where the probability of winning the object is a ratio of

[2] Varian's (2010) textbook provides a more complete introduction to auctions and mechanism design but, unlike this paper, it does not focus on equilibrium bidding strategies.

my bid relative to the sum of all bidders' bids, i.e., $prob(win) = \frac{b_1}{b_1+b_2+...+b_N}$, an allocation rule often used in certain Chinese auctions.

b) A *payment rule*, which describes how much every bidder must pay. For instance, the payment rule in the FPA determines that the individual submitting the highest bid pays his own bid, while everybody else pays zero. In contrast, the payment rule in the SPA specifies that the individual submitting the highest bid (the winner) pays the second-highest bid, while everybody else pays zero. Finally, the payment rule in the all-pay auction determines that every individual must pay the bid that he/she submitted.[3]

2.1 Privately observed valuations

Before analyzing equilibrium bidding strategies in different auction formats, note that auctions are strategic settings where players must choose their strategies (i.e., how much to bid) in an incomplete information context.[4] In particular, every bidder knows his/her own valuation for the object, v_i, but does not observe other bidder j's valuation, $j \neq i$. That is, bidder i is "in the dark" about his opponent's valuation.

Despite not observing j's valuation, bidder i knows the probability distribution behind bidder j's valuation. For instance, v_j can be relatively high, e.g., $v_j = 10$, with probability 0.4, or low, $v_j = 5$, otherwise (with probability 0.6). More generally, bidder j's valuation, v_j, is distributed according to a cumulative distribution function $F(v) = prob(v_j < v)$, intuitively representing that the probability that v_j lies below a certain cutoff v is exactly $F(v)$. For simplicity, we normally assume that every bidder's valuation for the object is drawn from a uniform distribution function between 0 and 1, i.e., $v_j \sim U[0,1]$.[5] The following figure illustrates this uniform distribution where the horizontal axis depicts v_j and the vertical axis measures the cumulated probability $F(v)$. For instance, if bidder i's valuation is v, then all points to the left-hand side of v in the horizontal axis represent that $v_j < v$, entailing that bidder j's valuation is lower than bidder i's. The mapping of these points into the vertical axis gives us the probability $prob(v_j < v) = F(v)$ which, in the case of a uniform distribution entails $F(v) = v$. Similarly, the valuations to the right-hand side of v describe points where $v_j > v$ and, thus, bidder j's valuation is higher than that of bidder i. Mapping these points into the vertical axis we obtain the probability $prob(v_j > v) = 1 - F(v)$ which, under a uniform distribution, implies $1 - F(v) = 1 - v$.

[3] This auction format is used by the internet seller QuiBids.com. For instance, if you participate in the sale of a new iPad, and you submit a low bid of \$25 but some other bidder wins by submitting a higher bid, you will still see your \$25 bid withdrawn from your QuiBids account.

[4] Auctions are, hence, regarded as an example of Bayesian game.

[5] Note that this assumption does not imply that bidder j does not assign a valuation v_j larger than one to the object but, instead, that his range of valuations, e.g., from 0 to $\overline{v}$, can be normalized to the interval $[0,1]$.

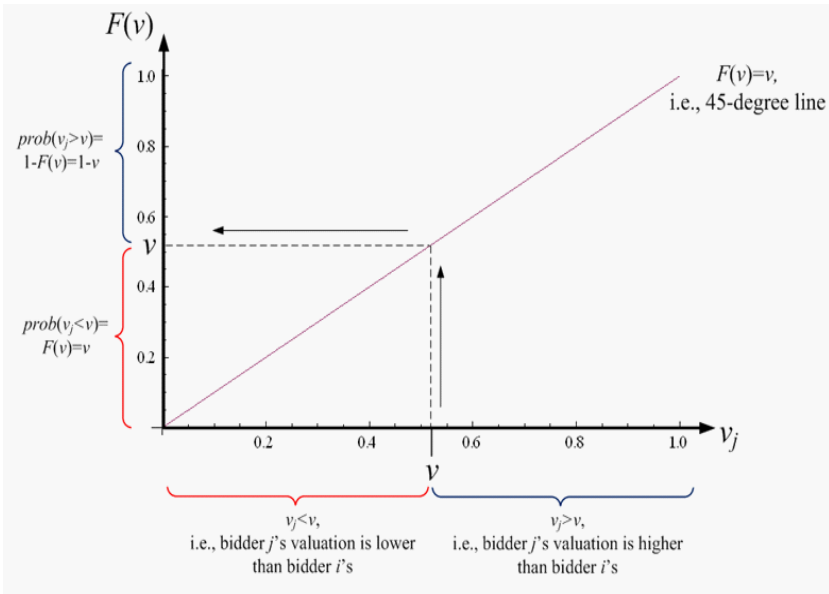

Fig 2. Uniform probability distribution.

Importantly, since all bidders are ex-ante symmetric, they will all be using the same bidding function, $b_i : [0,1] \to \mathbb{R}_+$, which maps bidder i's valuation, $v_i \in [0,1]$, into a precise bid, $b_i \in \mathbb{R}_+$. However, the fact that bidders use a symmetric function does not imply that all of them submit the same bid. Indeed, depending on his privately observed valuation for the object, bidding function $b_i(v_i)$ prescribes that bidders can submit different bids. As an example, consider a symmetric bidding function $b_i(v_i) = \frac{v_i}{2}$. Hence, a bidder with valuation $v_i = 0.4$ will submit a bid of $b_i(0.4) = 0.4/2 = \$0.2$, while a different bidder whose valuation is $v_i = 0.9$ would submit a bid of $b_i(0.9) = 0.9/2 = \$0.45$. In other words, even if bidders are *symmetric* in the bidding function they use, they can be *asymmetric* in the actual bid they submit.

3 First-price auctions

Let's start analyzing equilibrium bidding behavior in the first-price auction (FPA). First, note that submitting a bid above one's valuation, $b_i > v_i$, is a dominated strategy. In particular, the bidder would obtain a negative payoff if winning, since his expected utility from participating in the auction

$$EU_i(b_i|v_i) = prob(win) \cdot (v_i - b_i) + prob(lose) \cdot 0$$

would be negative, since $v_i < b_i$, regardless of his probability of winning. Note that in the above expected utility, we specify that, upon winning, bidder i receives a net payoff of $v_i - x$, i.e., the difference between his true valuation for the object and the bid he submits (which ultimately

constitutes the price he pays for the good if he were to win).[6] Similarly, submitting a bid b_i that exactly coincides with one's valuation, $b_i = v_i$, also constitutes a dominated strategy since, even if the bidder happens to win, his expected utility would be zero, i.e., $EU_i(b_i|v_i) = prob(win) \cdot (v_i - b_i)$, given that $b_i = v_i$. Therefore, the equilibrium bidding strategy in a FPA must imply a bid below one's valuation, $b_i < v_i$. That is, bidders must practice what is usually referred to as "bid shading." In particular, if bidder i's valuation is v_i, his bid must be a share of his true valuation, i.e., $b_i(v_i) = a \cdot v_i$, where $a \in (0, 1)$. The following figure illustrates bid shading in the FPA, since bidding strategies must lie below the 45-degree line.

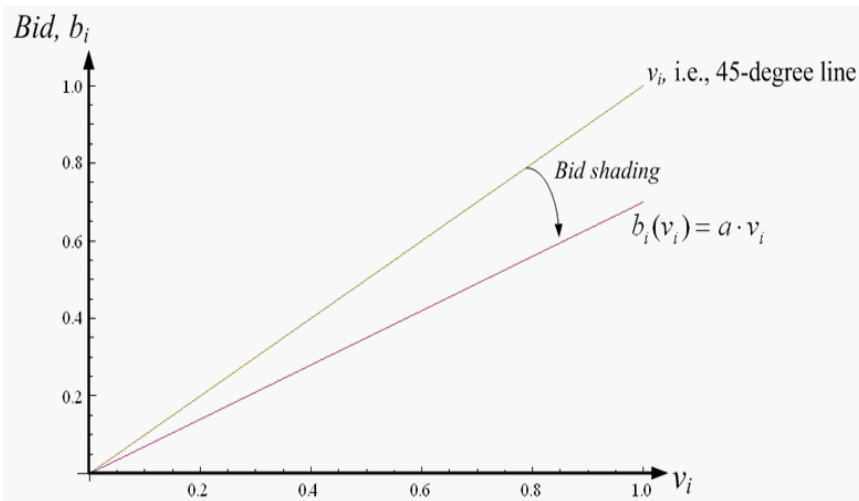

Fig 3. "Bid shading" in the FPA.

A natural question at this point is: How intense bid shading must be in the FPA? Or, alternatively, what is the precise value of the bid shading parameter a? In order to answer such question, we must first describe bidder i's expected utility from submitting a given bid x, when his valuation for the object is v_i,

$$EU_i(x|v_i) = prob(win) \cdot (v_i - x) + prob(lose) \cdot 0$$

Before continuing our analysis, we still must precisely characterize the probability of winning in the above expression, i.e., $prob(win)$. Specifically, upon submitting a bid $b_i = x$, bidder j can anticipate that bidder i's valuation is $\frac{x}{a}$, by just inverting the bidding function $b_i(v_i) = x = a \cdot v_i$, i.e., solving for v_i in $x = a \cdot v_i$ yields $v_i = \frac{x}{a}$. This inference is illustrated in the figure below where bid x in the vertical axis is mapped into the bidding function $a \cdot v_i$, which corresponds to a valuation of $\frac{x}{a}$ in the horizontal axis. Intuitively, for a bid x, bidder j can use the symmetric bidding function $a \cdot v_i$ to "recover" bidder i's valuation, $\frac{x}{a}$.

[6] Upon loosing, bidders do not obtain any object and, in this auction, do not have to pay any monetary amount, thus implying a zero payoff.

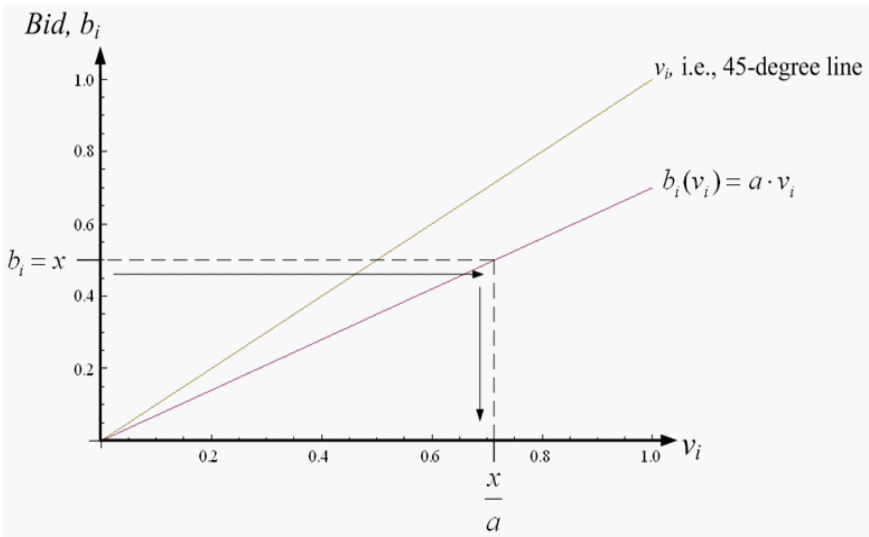

Fig 4. "Recovering" bidder i's valuation.

Hence, the probability of winning is given by $prob(b_i \geq b_j)$ and, according to the vertical axis in the previous figure, $prob(b_i > b_j) = prob(x > b_j)$. If, rather than describing probability $prob(x > b_j)$ from the point of view of bids (see shaded portion of the vertical axis in figure 5 below), we characterize it from the point of view of valuations (in the shaded segment of the horizontal axis), we obtain that $prob(b_i > b_j) = prob(\frac{x}{a} > v_j)$.

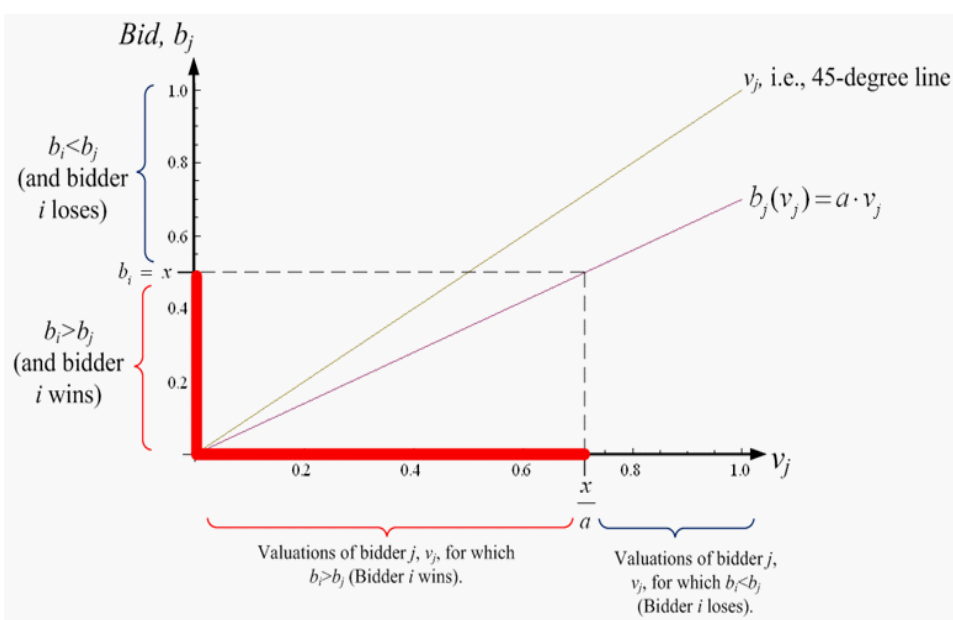

Fig. 5. Probability of winning in the FPA.

Indeed, the shaded set of valuations in the horizontal axis illustrates valuations of bidder j, v_j, for which his bid lies below player i's bid x. In contrast, valuations v_j satisfying $v_j > \frac{x}{a}$ entail that player j's bids would exceed x, implying that bidder j wins the auction. Hence, if the probability that bidder i wins the object is given by $prob(\frac{x}{a} > v_j)$, and valuations are uniformly distributed, we have that $prob(\frac{x}{a} > v_j) = \frac{x}{a}$.[7] We can now plug this probability of winning into bidder i's expected utility from submitting a bid of x in the FPA, as follows

$$EU_i(x|v_i) = \frac{x}{a}(v_i - x) = \frac{v_i x - x^2}{a}$$

Taking first-order conditions with respect to bidder i's bid, x, we obtain $\frac{v_i - 2x}{a} = 0$ which, solving for x yields bidder i's optimal bidding function $x(v_i) = \frac{1}{2}v_i$. Intuitively, this bidding function informs bidder i how much to bid, as a function of his privately observed valuation for the object, v_i. For instance, when $v_i = 0.75$, his optimal bid is $\frac{1}{2}0.75 = 0.375$. This bidding function implies that, when competing against another bidder j, and only $N = 2$ players participate in the FPA, bidder i shades his bid in half, as the following figure illustrates.

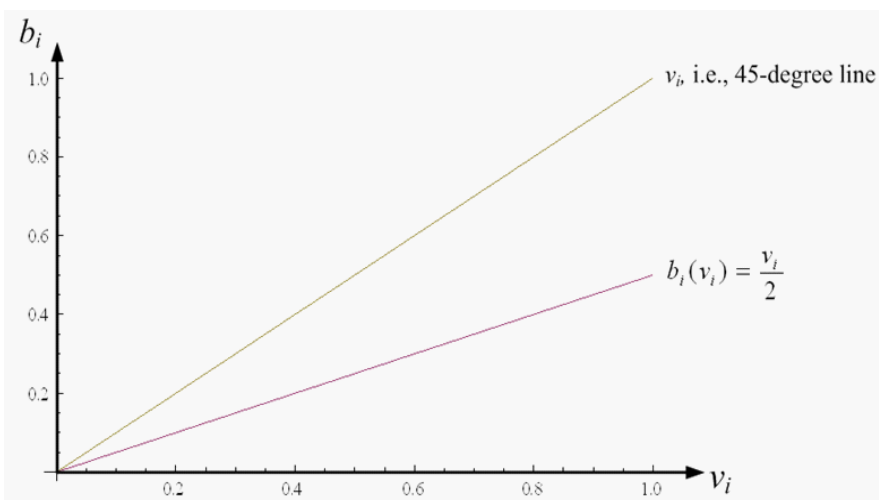

Fig 6. Optimal bidding function with $N = 2$ bidders.

3.1 Extending the first-price auction to N bidders

Our results are easily extensible to FPA with N bidders. In particular, the probability of bidder i winning when submitting a bid of $\$x$ is

$$\begin{aligned} prob(win) &= prob\left(\frac{x}{a} > v_1\right) \cdot \ldots \cdot prob\left(\frac{x}{a} > v_{i-1}\right) \cdot prob\left(\frac{x}{a} > v_{i+1}\right) \cdot \ldots \cdot prob\left(\frac{x}{a} > v_N\right) \\ &= \frac{x}{a} \cdot \ldots \cdot \frac{x}{a} \cdot \frac{x}{a} \cdot \ldots \cdot \frac{x}{a} = \left(\frac{x}{a}\right)^{N-1} \end{aligned}$$

[7] Recall that, if a given random variable y is distributed according to a uniform distribution function $U[0,1]$, the probability that the value of y lies below a certain cutoff c, is exactly c, i.e., $prob(y < c) = F(c) = c$.

where we evaluate the probability that the valuation of all other $N-1$ bidders, v_1, v_2,..., v_{i-1}, v_{i+1},..., v_N (expect for bidder i) lies above the valuation $v_i = \frac{x}{a}$ that generates a bid of exactly $\$x$. Hence, bidder i's expected utility from submitting x becomes

$$EU_i(x|v_i) = \left(\frac{x}{a}\right)^{N-1}(v_i - x) + \left[1 - \left(\frac{x}{a}\right)^{N-1}\right]0$$

Taking first-order conditions with respect to his bid, x, we obtain

$$-\left(\frac{x}{a}\right)^{N-1} + \left(\frac{x}{a}\right)^{N-2}\left(\frac{1}{a}\right)(v_i - x) = 0$$

Rearranging, $\left(\frac{x}{a}\right)^N \frac{a}{x^2}\left[(N-1)v_i - nx\right] = 0$, and solving for x, we find bidder i's optimal bidding function, $x(v_i) = \frac{N-1}{N}v_i$. The following figure depicts the bidding function for the case of $N = 2$, $N = 3$, and $N = 4$ bidders, showing that bid shading is ameliorated when more bidders participate in the auction, i.e., bidding functions approach the 45-degree line. Indeed, for $N = 2$ the optimal bidding function is $\frac{1}{2}v_i$, but it increases to $\frac{2}{3}v_i$ when $N = 3$ bidders compete for the object, to $\frac{3}{4}v_i$ when $N = 4$ players participate in the auction, etc. For a extremely large number of bidders, e.g., $N = 2,000$, bidder i's optimal bidding function becomes $b_i(v_i) = \frac{1,999}{2,000}v_i \simeq v_i$ and, hence, bidder i's bid almost coincides with his valuation for the object, describing a bidding function that approaches the 45-degree line.

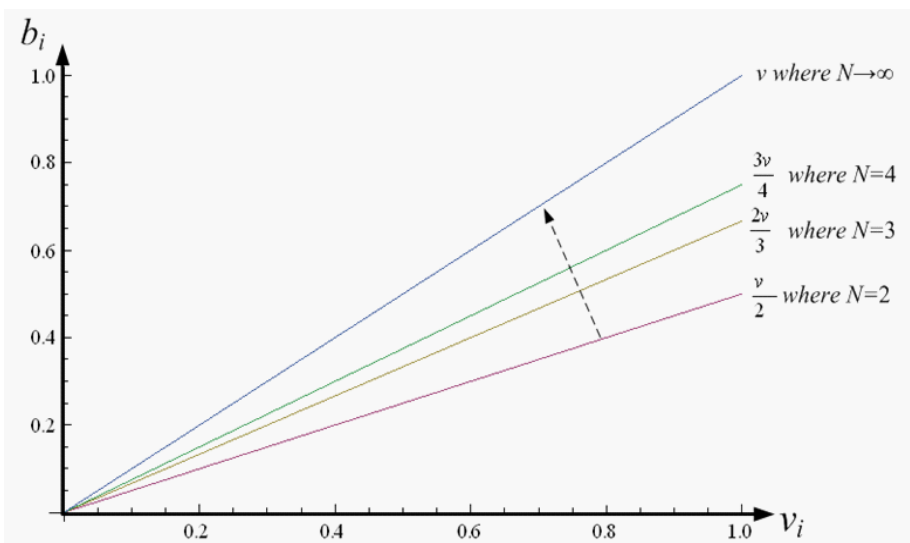

Fig 7. Optimal bidding function increases in N.

Intuitively, if bidder i seeks to win the object, he can shade his bid when only another bidder competes for the good, since the probability of him assigning a large valuation to the object is relatively low. However, when several players compete in the auction, the probability that some of them has a high valuation for the object (and, thus submits a high bid) increases. That is,

competition gets "tougher" as more bidders participate and, as a consequence, every bidder must increase his bid, ultimately ameliorating his incentives to practice bid shading.

3.2 First-price auctions with risk-averse bidders

Let us next analyze how our equilibrium results would be affected if bidders are risk averse, i.e., their utility function is concave in income , x, e.g., $u(x) = x^{\alpha}$, where $0 < \alpha \leq 1$ denotes bidder i's risk-aversion parameter. In particular, when $\alpha = 1$ he is risk neutral, while when α decreases, he becomes risk averse.[8] First, note that the probability of winning is unaffected, since, for a symmetric bidding function $b_i(v_i) = a \cdot v_i$ for every bidder i, where $a \in (0, 1)$, the probability that bidder i wins the auction against another bidder j is

$$prob(b_i > b_j) = prob(x > b_j) = prob\left(\frac{x}{a} > v_j\right) = \frac{x}{a}$$

Therefore, bidder i's expected utility from participating in this auction is

$$EU_i(x|v_i) = \frac{x}{a} \times (v_i - x)^{\alpha} + \left(1 - \frac{x}{a}\right) \times 0$$

where, relative to the case of risk-neutral bidders analyzed above, the only difference arises in the evaluation of the net payoff from winning, $v_i - x$, which it is evaluated as $(v_i - x)^{\alpha}$. Taking first-order conditions with respect to his bid, x, we have

$$\frac{1}{a}(v_i - x)^{\alpha} - \frac{x}{a}\alpha(v_i - x)^{\alpha-1} = 0,$$

and solving for x, we find the optimal bidding function, $x(v_i) = \frac{v_i}{1+\alpha}$. Importantly, this case embodies that of risk-neutral bidders analyzed above as a special case. Specifically, when $\alpha = 1$, bidder i's optimal bidding function becomes $x(v_i) = \frac{v_i}{2}$. However, when his risk aversion increases, i.e., α decreases, bidder i's optimal bidding function increases. Specifically, $\frac{\partial x(v_i)}{\partial \alpha} = -\frac{v_i}{(1-\alpha)^2}$, which is negative for all parameter values. In the extreme case in which α decreases to $\alpha \to 0$, the optimal bidding function becomes $x(v_i) = v_i$, and players do not practice bid shading. The following figure illustrates the increasing pattern in players' bidding function, starting from $\frac{v_i}{2}$ when bidders are risk neutral, $\alpha = 1$, and approaching the 45-degree line (no bid shading) as players become more risk averse.

[8] A tipical example you have probably encountered in intermediate microeconomics courses includes $u(x) = \sqrt{x}$ since $\sqrt{x} = x^{1/2}$. As a practice, note that the Arrow-Pratt coefficient of absolute risk aversion $r_A(x) = -\frac{u''(x)}{u'(x)}$ for this utility function yields $\frac{1-\alpha}{x}$, confirming that, when $\alpha = 1$, the coefficient of risk aversion becomes zero, but when $0 < \alpha < 1$, the coefficient is positive.

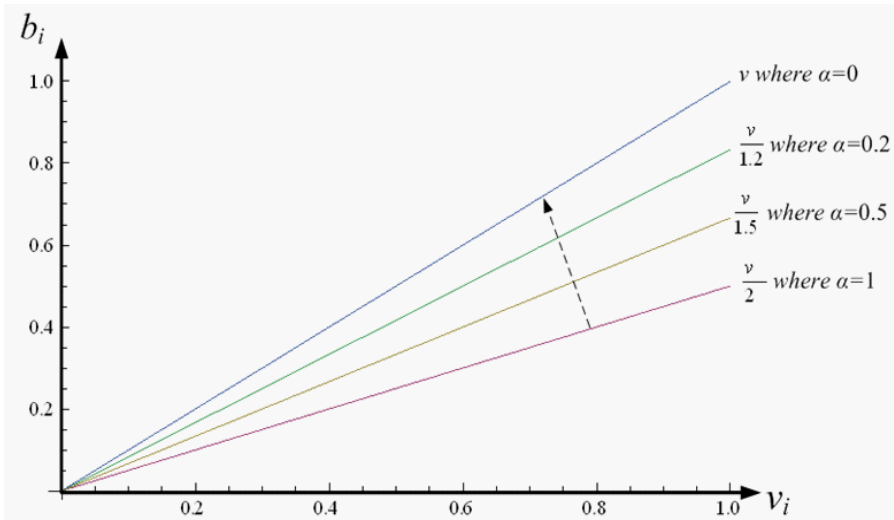

Fig. 8. Optimal bidding function with risk-averse bidders.

Intuitively, a risk-averse bidder submits more aggressive bids than a risk-neutral bidder in order to minimize the probability of losing the auction. In particular, consider that bidder i reduces his bid from b_i to $b_i - \varepsilon$. In this context, if he wins the auction, he obtains an additional profit of ε, since he has to pay a lower price for the object he acquires. However, by lowering his bid, he increases the probability of losing the auction. Importantly, for a risk-averse bidder, the positive effect of slightly lowering his bid, arising from getting the object at a cheaper price, is offset by the negative effect of increasing the probability that he loses the auction. In other words, since the possible loss from losing the auction dominates the benefit from acquiring the object at a cheaper price, the risk-averse bidder does not have incentives to reduce his bid, but rather to increase it, relative to the risk-neutral bidders.

4 Second-price auction

In this class of auctions, bidding your own valuation, i.e., $b_i(v_i) = v_i$, is a weakly dominant strategy for all players. That is, regardless of the valuation you assign to the object, and independently on your opponents' valuations, submitting a bid $b_i(v_i) = v_i$ yields expected profit equal or above that from submitting any other bid, $b_i(v_i) \neq v_i$. In order to show this bidding strategy is an equilibrium outcome of the SPA, let's first examne bidder i's expected payoff from submitting a bid that coincides with his own valuation v_i (which we refer to as the First case below), and then compare it with what he would obtain from deviating to bids below his valuation for the object, $b_i(v_i) < v_i$ (denoted as Second case), or above his valuation, $b_i(v_i) > v_i$ (Third case). Let us next separately analyze the payoffs resulting from each bidding strategy.

First case: If the bidder submits his own valuation, $b_i(v_i) = v_i$, then either of the following

situations can arise (for presentation purposes, the figure below depicts each of the three cases separately):

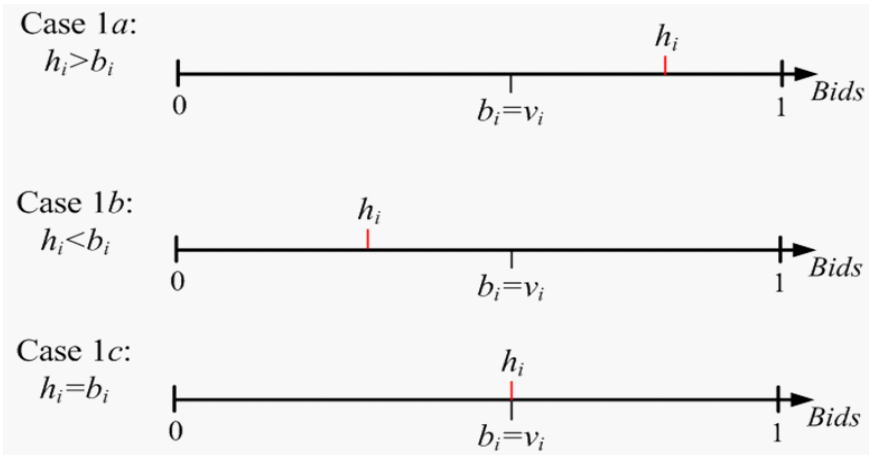

Fig 9. Cases arising when $b_i(v_i) = v_i$.

1a) If his bid lies below the highest competing bid, i.e., $b_i < h_i$ where $h_i = \max_{j \neq i}\{b_j\}$,[9] then bidder i loses the auction, obtaining a zero payoff.

1b) If his bid lies above the highest competing bid, i.e., $b_i > h_i$, then bidder i wins the auction. In this case, he obtains a net payoff of $v_i - h_i$, since in a SPA the winning bidder does not have to pay the bid he submitted, but the second-highest bid, which is h_i in this case since $b_i > h_i$.

1c) If, instead, his bid coincides with the highest competing bid, i.e., $b_i = h_i$, then a tie occurs. For simplicity, ties are normally solved in auctions by randomly assigning the object to the bidders who submitted the highest bids. As a consequence, bidder i's payoff becomes $v_i - h_i$, but with only $\frac{1}{2}$ probability, i.e., his expected payoff $\frac{1}{2}(v_i - h_i)$.[10]

Second case: Let us now compare the above equilibrium payoffs with those bidder i could obtain by deviating towards a bid that shades his valuation, i.e., $b_i < v_i$. In this case, we can also identify three cases emerging (see figure 10), depending on the ranking between bidder i's bid, b_i, and the highest competing bid, h_i.

[9] Intuitively, expression $h_i = \max_{j \neq i}\{b_j\}$ just finds the highest bid among all bidders different from bidder i, $j \neq i$.

[10] Note that, more generally, if $K \geq 2$ bidders coincide in submitting the highest bid, their expected payoff becomes $\frac{1}{K}(v_i - h_i)$.

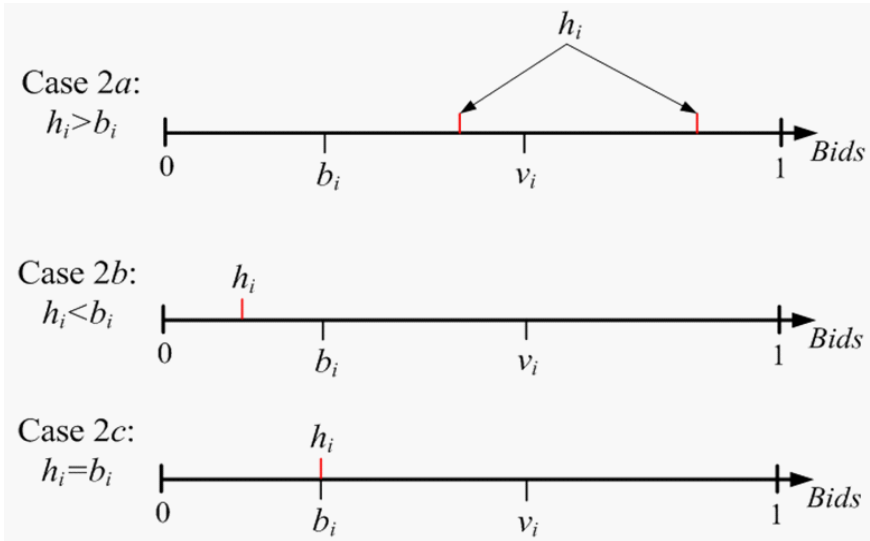

Fig 10. Cases arising when $b_i(v_i) < v_i$.

2a) If his bid lies below the highest competing bid, i.e., $b_i < h_i$, then bidder i loses the auction, obtaining a zero payoff.

2b) If his bid lies above the highest competing bid, i.e., $b_i > h_i$, then bidder i wins the auction, obtaining a net payoff of $v_i - h_i$.

2c) If, instead, his bid coincides with the highest competing bid, i.e., $b_i = h_i$, then a tie occurs, and the object is randomly assigned, yielding an expected payoff of $\frac{1}{2}(v_i - h_i)$.

Hence, we just showed that bidder i obtains the same payoff submitting a bid that coincides with his privately observed valuation for the object ($b_i = v_i$, as in the First case) and shading his bid ($b_i < v_i$, as described in teh Second case). Therefore, he does not have incentives to conceal his bid, since his payoff would not improve from doing so.

Third case: Let us finally examine bidder i's equilibrium payoff from submitting a bid above his valuation, i.e., $b_i(v_i) > v_i$. In this case, three cases also arise (see figure 11).

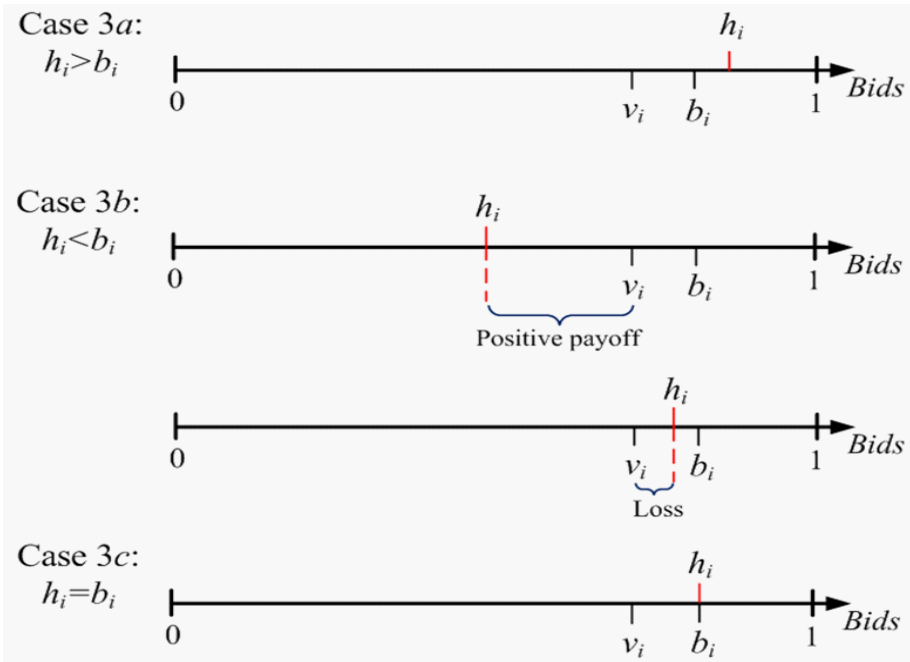

Fig 11. Cases arising when $b_i(v_i) > v_i$.

3a) If his bid lies below the highest competing bid, i.e., $b_i < h_i$, then bidder i loses the auction, obtaining a zero payoff.

3b) If his bid lies above the highest competing bid, i.e., $b_i > h_i$, then bidder i wins the auction. In this scenario, his payoff becomes $v_i - h_i$, which is positive if $v_i > h_i$, or negative otherwise. (These two situations are depicted in case 3b of figure 11.) The latter case, in which bidder i wins the auction but at a loss (negative expected payoff), did not exist in our above analysis of $b_i(v_i) = v_i$ and $b_i(v_i) < v_i$, since players did not submit bids above their own valuation. Intuitively, the possibility of a negative payoff arises because bidder i's valuation can lie below the second-highest bid, as figure 11 illustrates, where $v_i < h_i < b_i$.

3c) If, instead, his bid coincides with the highest competing bid, i.e., $b_i = h_i$, then a tie occurs, and the object is randomly assigned, yielding an expected payoff of $\frac{1}{2}(v_i - h_i)$. Similarly as our above discussion, this expected payoff is positive if $v_i > h_i$, but negative otherwise.

Hence, bidder i's payoff from submitting a bid above his valuation either coincides with his payoff from submitting his own value for the object, or becomes strictly lower, thus nullifying his incentives to deviate from his equilibrium bid of $b_i(v_i) = v_i$. In other words, there is no bidding strategy that provides a strictly higher payoff than $b_i(v_i) = v_i$ in the SPA, and all players bid their own valuation, without shading their bids; a result that differs from the optimal bidding function in FPA, where players shade their bids unless $N \to \infty$.

Remark. The above equilibrium bidding strategy in the SPA is, importantly, unaffected by the number of bidders who participate in the auction, N, or their risk-aversion preferences. In particular, our above discussion considered the presence of N bidders, and an increase in their number does not emphasize or ameliorate the incentives that every bidder has to submit a bid that coincides with his own valuation for the object, $b_i(v_i) = v_i$. Furthermore, the above results remain when bidders evaluate their net payoff, e.g., $v_i - h_i$, according to a concave utility function, such as $u(x) = x^\alpha$, exhibiting risk aversion. Specifically, for a given value of the highest competing bid, h_i, bidder i's expected payoff from submitting a bid $b_i(v_i) = v_i$ would still be weakly larger than from deviating to a bidding strategy above, $b_i(v_i) > v_i$, or below, $b_i(v_i) < v_i$, his true valuation for the object.

4.1 Efficiency in auctions

Auctions, and generally allocation mechanism, are characterized as efficient if the bidder (or agent) with the highest valuation for the object is indeed the person receiving the object. Intuitively, if this property does not hold, the outcome of the auction (i.e., the allocation of the object) would open the door to negotiations and arbitrage among the winning bidder —who, despite obtaining the object, is not the player who assigns the highest value to it— and other bidder/s with higher valuations who would like to buy the object from him. In other words, the auction's outcome would still allow for negotiations that are beneficial for all parties involved, i.e., Pareto improving negotiations, thus suggesting that the initial allocation was not Pareto efficient.

According to this criterion, both the FPA and the SPA are efficient, since the bidder with the highest valuation submits the highest bid, and the object is ultimately assigned to the player who submits the highest bid. Other auction formats, such as the Chinese (or lottery) auction described in the Introduction, are not necessarily efficient, since they may assign the object to an individual who did not submit the highest valuation for the object. In particular, recall that the probability of winning the object in this auction is a ratio of the bid you submit relative to the sum of all players' bids. Hence, a bidder with a low valuation for the object, and who submits the lowest bid, e.g., \$1, can still win the auction. Alternatively, the person that assigns the highest value to the object, despite submitting the highest bid, might not end up receiving the object for sale. Therefore, for an auction to satisfy efficiency, bids must be increasing in a player's valuation, and the probability of winning the auction must be one (100%) if a bidder submits the highest bid.

5 Common-value auctions

The auction formats considered above assume that each bidders privately observes his own valuation for the object, and this valuation is distributed according to a distribution function $F(v)$, e.g., a uniform distribution, implying that two bidders are unlikely to assign the same value to the object for sale. However, in some auctions, such as the government sale of oil leases, bidders (oil companies) might assign the same monetary value to the object (common value), i.e., the profits they would

obtain from exploiting the oil reservoir. Bidders are, nonetheless, unable to precisely observe the value of this oil reservoir but, instead, gather estimates of its value. In the oil lease example, firms cannot accurately observe the exact volume of oil in the reservoir, or how difficult it will be to extract, but can accumulate different estimates from their own engineers, or from other consulting companies, that inform the firm about the potential profits to be made from the oil lease. Such estimates are, nonetheless, imprecise, and only allow the firm to assign a value to the object (profits from the oil lease) within a relatively narrow range, e.g., $v \in [10, 11, \ldots, 20]$ in millions of dollars. Consider that oil company A hires a consultant, and gets a signal (a report), s, as follows

$$s = \begin{cases} v + 2 \text{ with prob. } \frac{1}{2}, \text{ and} \\ v - 2 \text{ with prob. } \frac{1}{2} \end{cases}$$

and, hence, the signal is above the true value to the oil lease with 50% probability, or below its value otherwise. We can alternatively represent this information by examining the conditional probability that the true value of the oil lease is v, given that the firm receives a signal s, is

$$prob(v|s) = \begin{cases} \frac{1}{2} \text{ if } v = s - 2 \text{ (overestimate), and} \\ \frac{1}{2} \text{ if } v = s + 2 \text{ (underestimate)} \end{cases}$$

since the true value of the lease is overestimated when $v = s - 2$, i.e., $s = v + 2$ and the signal is above v; and underestimated when $v = s + 2$, i.e., $s = v - 2$ and the signal lies below v. Notice that, if company A was not participating in the auction, then the expected value of the oil lease would be

$$\underbrace{\frac{1}{2}(s-2)}_{\text{if overestimation}} + \underbrace{\frac{1}{2}(s+2)}_{\text{if underestimation}} = \frac{(s-2)+(s+2)}{2} = s$$

implying that the firm would pay for the oil lease a price $p < s$, making a positive expected profit. But, what if the oil company participates in a FPA for the oil lease against another company B? In this context, every firm uses a different consultant, i.e., can receive different signals, but does not know whether their consultant systematically over- or under-estimates the true value of the oil lease. In particular, consider that every firm receives a signal s from their consultant. Observing its own signal, but not observing the signal received by the other firm, every firm $i = \{A, B\}$ submits a bid from the set $\{1, 2, \ldots, 20\}$, where the upper bound of this interval represents the maximum value of the oil lease according to all estimates.

We will next show that slightly shading your bid, e.g., submitting $b = s - 1$, cannot be optimal for any firm. At first glance, however, such a bidding strategy seems sensitive: the firm bid is increasing in the signal it receives and, in addition, its bid is below the signal, $b < s$, entailing that, if the true value of the oil lease was s, the firm would obtain a positive expected profit from winning. In order to show that bid $b = s - 1$ cannot be optimal, consider that firm A receives a signal $s = 10$, and thus submits a bid $b = s - 1 = 10 - 1 = \$9$. Given such a signal, the true value

of the oil lease is

$$v = \begin{cases} s + 2 = 12 & \text{with prob. } \frac{1}{2}, \text{ and} \\ s - 2 = 8 & \text{with prob. } \frac{1}{2} \end{cases}$$

Specifically, when the true value of the oil lease is $v = 12$, firm A receives a signal of $s_A = 10$ (an underestimation of the true valuation, 12), while firm B receives a signal of $s_B = 14$ (an overestimation). In this setting, firms bid $b_A = 10 - 1 = \$9$ and $b_B = 14 - 1 = \$13$ and, thus, firm A loses the auction. If, in contrast, the true value of the lease is $v = 8$, firm A receives a signal of $s_A = 10$ (an overestimation of the true valuation, 8), while firm B receives a signal $s_B = 6$ (an underestimation). In this context, firms bid $b_A = 10 - 1 = \$9$, and $b_B = 6 - 1 = \$5$, and firm A wins the auction. In particular, firm A's expected profit from participating in this auction is

$$\frac{1}{2}(8 - 9) + \frac{1}{2}0 = -\frac{1}{2}$$

which is negative! This is the so-called "winner's curse" in common-value auctions. In particular, the fact that a bidder wins the auction just means that he probably received an overestimated signal of the true value of the object for sale, as firm A receiving signal $s_A = 10$ in the above example. Therefore, in order to avoid the winner's curse, participants in common-value auctions must significantly shade their bid, e.g., $b = s - 2$ or less, in order to consider the possibility that the signals they receive are overestimating the true value of the object.[11]

The winner's curse in practice. Despite the straightforward intuition behind this result, the winner's curse has been empirically observed in several controlled experiments. A common example is that of subjects in an experimental lab, where they are asked to submit bids in a common-value auction where a jar of nickels is being sold. Consider that your instructor shows up in class with a bid jar plenty of nickels. The monetary value you assign to the jar coincides with that of your classmates, but none of you can accurately estimate the number of nickels in the jar, since you can only gather some imprecise information about its true value by looking at it for a few seconds. In these experiments, it is usual to find that the winner ends up submitting a bid a monetary amount beyond the jar's true value, i.e., the winner's curse emerges.[12]

More surprisingly, the winner's curse has also been shown to arise among oil company executives. Hendricks et al. (2003) analyze the bidding strategies of companies, such as Texaco, Exxon, an British Petroleum, when competing for the mineral rights to properties 3-200 miles off-shore and initially owned by the U.S. government. Generally, executives did not systematically fall prey of the winner's curse, since their bids were about 1/3 of the true value of the oil lease. As a consequence, if their bids resulted in their company winning the auction, their expected profits would become positive. Texaco executives, however, not only fell prey of the winner's curse, but submitted bids above the estimated value of the oil lease. Such a high bid, if winning, would have resulted in

[11] It can be formally shown that, in the case of $N = 2$ bidders, the optimal bidding function is $b_i(v_i) = \frac{1}{2}s_i$, where s_i denotes the signal that bidder i receives. More generally, for N bidders, bidder i's optimal bid becomes $b_i(v_i) = \frac{(N+2)(N-1)}{2N^2}s_i$. For more details, see Harrington (2009), pp. 321-23.

[12] For some experimental evidence on the winner's curse see, for instance, Thaler (1988).

negative expected profits. One cannot help but wonder if Texaco executives were enrolled in a remedial course on auction theory.

6 Suggested exercises

1. Consider an auction with five participants, each of them with the following (privately observed) valuation of the object for sale: Person A ($10), Person B ($6), Person C ($45), Person D ($81), and Person E ($62).

   (a) If the seller organizes a second-price auction, who will be the winner? What will be his winning bid? What price he will pay for the object?

   (b) Suppose now that bidders can observe each other's valuations, but the seller cannot. The seller, however, only knows that bidders' valuations are in the range $\{0, 1, ..., \$90\}$. If players participate in a first-price auction, how will be the winner? What is his winning bid?

2. [**All-pay auction**] Consider the following all-pay auction with two bidders privately observing their valuation for the object. Valuations are uniformly distributed $v_i \sim U[0, 1]$. The player submitting the highest bid wins the object, but all players must pay the bid they submitted. Find the optimal bidding strategy, taking into account that it is of the form $b_i(v_i) = m \cdot v_i^2$, where m denotes a positive constant.

3. [**Third-price auction**] Consider a third-price auction, where the winner is the bidder who submits the highest bid, but he/she only pays the third highest bid. Assume that you compete against two other bidders, whose valuations you are unable to observe, and that your valuation for the object is $10. Show that bidding *above* your valuation (with a bid of, for instance, $15) can be a best response to the other bidders' bid, while submitting a bid that coincides with your valuation ($10) might not be a best response to your opponents' bids.

References

[1] Besanko, D. and R. Braeutigam (2011) *Microeconomics*, Wiley Publishers.

[2] Harrington, J. (2009) *Games, Strategies, and Decision Making*, Worth Publishers.

[3] Hendricks, K., J. Pinske, and R.H. Porter (2003) "Empirical implications of equilibrium bidding in first-price, symmetric, common-value auctions," Review of Economic Studies, 70, pp. 115-45.

[4] Klemperer, P. (2004) *Auctions: Theory and Practice* (The Toulouse Lectures in Economics). Princeton University Press.

[5] Krishna, V. (2002), *Auction Theory.* Academic Press.

[6] McKinsey (2002), "Comparative Assessment of the Licensing Regimes for 3G Mobile Communications in the European Union and their impact on the Mobile Communications Sector," available at the European Commission's website: http://ec.europa.eu/information_society/topics/telecoms/radiospec.

[7] Menezes, F.M. and P.K. Monteiro (2004), *An Introduction to Auction Theory,* Oxford University Press.

[8] Milgrom, P. (2004), *Putting Auction Theory to Work*, Cambridge University Press.

[9] Pindyck, R. and D. Rubinfeld (2012) *Microeconomics*, Pearson Publishers.

[10] Perloff, J.M. (2011) *Microeconomics, Theory and Applications with Calculus*, Addison Wesley.

[11] Thaler, R. (1988) "Anomalies: The Winner's Curse," The Journal of Economic Perspectives, 2(1), pp. 191-202.

[12] Varian, H. (2010) *Intermediate Microeconomics*, Norton Publishing.