

Spiteful Bidding in Sealed-Bid Auctions

Felix Brandt

Computer Science Department
University of Munich
80538 Munich, Germany
brandtf@tcs.ifi.lmu.de

Tuomas Sandholm

Computer Science Department
Carnegie Mellon University
Pittsburgh PA 15213, USA
sandholm@cs.cmu.edu

Yoav Shoham

Computer Science Department
Stanford University
Stanford CA 94305, USA
shoham@cs.stanford.edu

Abstract

We study the bidding behavior of spiteful agents who, contrary to the common assumption of self-interest, maximize a convex combination of their own profit and their competitors' losses. The motivation for this assumption stems from inherent spitefulness or, for example, from competitive scenarios such as in closed markets where the loss of a competitor will likely result in future gains for oneself. We derive symmetric Bayes Nash equilibria for spiteful agents in 1st-price and 2nd-price sealed-bid auctions. In 1st-price auctions, bidders become "more truthful" the more spiteful they are. Surprisingly, the equilibrium strategy in 2nd-price auctions does not depend on the number of bidders. Based on these equilibria, we compare the revenue in both auction types. It turns out that expected revenue in 2nd-price auctions is higher than expected revenue in 1st-price auctions in the case of even the most modestly spiteful agents, provided they still care at least a little for their own profit. In other words, revenue equivalence only holds for auctions in which all agents are either self-interested or completely malicious. We furthermore investigate the impact of common knowledge on spiteful bidding. Divulging the bidders' valuations reduces revenue in 2nd-price auctions, whereas it has the opposite effect in 1st-price auctions.

1 Introduction

Over the last few years, game theory has widely been adopted as a tool to formally model and analyze interactions between rational agents in the field of AI. One of the fundamental assumptions in game theory is that agents are self-interested, *i.e.*, they maximize their own utility without considering the utility of other agents. However, there is some evidence that certain types of behavior in the real world as well as in artificial societies can be better explained by models in which agents have *other-regarding preferences*. While there are settings where agents exhibit altruism, there are also others where agents intend to degrade competitors in order to improve their own standing. This is typically the case in competitive situations such as in closed markets where the loss of

a competitor will likely result in future gains for oneself (*e.g.*, when a competitor is driven out of business) or, more generally, when agents intend to maximize their relative rather than their absolute utility. To give an example, consider the popular Trading Agent Competition (TAC) where an agent's goal (in order to win the competition) should be to accumulate more revenue than his competitors instead of maximizing his own revenue. Other examples include the German 3G mobile phone spectrum license auction¹ in 2000, where one of the network providers (German Telekom) kept raising the price in an (unsuccessful) attempt to crowd out one of the weaker competitors (Grimm *et al.*, 2002), and sponsored search auctions where similar behavior has been reported (Zhou and Lukose, 2006). Clearly, a reduced number of competitors is advantageous for the remaining companies because it increases their market share.

Sealed-bid auctions are an example of well-understood competitive economic processes where questioning the assumption of self-interest is particularly pertinent. For instance, in a 2nd-price auction, it is a dominant strategy for a self-interested agent to truthfully submit his valuation, even if he is informed about the other bids. However, a competitive agent who *knows* he cannot win might feel tempted to place his bid right below the winning bid in order to minimize the winner's profit. In order to account for the behavior of such spiteful agents, which are interested in minimizing the returns to their competitors as well as in maximizing their own profits, we incorporate other-regarding preferences in the utility function. To this end, we define a new utility measure for agents with negative externalities (Section 3). The trade-off between both goals is controlled by a parameter α called *spite coefficient*. Setting α to zero yields self-interested agents whereas a spite coefficient of one defines completely malicious agents, whose only goal is to reduce others' profit. We find that the well-known equilibria for 1st- and 2nd-price auctions no longer apply if $\alpha > 0$. In Sections 4 and 5, respectively, we derive symmetric Bayes Nash equilibria of both auction types in the case of spiteful agents. With respect to these equilibria, we obtain further results on auction revenue and the impact of common knowledge in Section 6. The paper concludes with Section 7.

¹With respect to the revenue generated (50.8 billion Euro), this auction is one of the most successful auctions to date.

2 Related Work

Numerous authors in experimental economics (Fehr and Schmidt, 2006; Saijo and Nakamura, 1995; Levine, 1998), game theory (Sobel, 2005), social psychology (Messick and Sentis, 1985; Loewenstein *et al.*, 1989), and multiagent systems (Brainov, 1999) have observed and explored other-regarding preferences, usually with an emphasis on altruism. Levine (1998) introduced a model in which utility is defined as a linear function of both the agent’s monetary payoff and his opponents’ payoff, controlled by a parameter called “altruism coefficient”. This model was used to explain data obtained in ultimatum bargaining and centipede experiments. One surprising outcome of that study was that an overwhelming majority of individuals possess a *negative* altruism coefficient, corresponding to spiteful behavior. He concludes that “one explanation of spite is that it is really ‘competitiveness’, that is, the desire to outdo opponents” (Levine, 1998). Most papers, including Levine’s, also consider elements of fairness in the sense that agents are willing to be more altruistic/spiteful to an opponent who is more altruistic/spiteful towards them. Brainov (1999) defines a generic type of “antisocial” agent by letting $\frac{\partial U_i}{\partial u_j} < 0$ for any $j \neq i$ (using the notation defined in Section 3.1). A game-theoretic model in which buyers have negative identity-dependent externalities which “can stand for expected profits in future interaction” has been studied by Jehiel *et al.* (1996).

We extend our previous work on spitefulness in auctions (Brandt and Weiß, 2001), where we have already given an equilibrium strategy for spiteful agents in 2nd-price auctions with *complete information*. Recently, other authors have studied the effects of negative externalities in auctions (Morgan *et al.*, 2003; Maasland and Onderstal, 2003). Perhaps closest to our work is the work by Morgan *et al.* (2003). Although we derived our results independently, there are some similarities as well as differences, and so we should make both clear. In contrast to Morgan *et al.*, we model spite as a convex combination of utilities which allows us to capture malicious agents who possess no self-interest at all (permitting results like Corollary 1). Morgan *et al.*’s definition approaches this case in the limit, but such limits are not considered in their paper. The main benefit of our definition of spite is the re-translation of bidding equilibria from bulky integrals to more intuitive conditional expectations which in turn greatly facilitates the proof of the main result (Theorem 3). Furthermore, we quantify the difference of revenue in both auctions types (Theorem 4), analyze the impact of common knowledge by computing equilibria for the complete information setting (Section 6.2), and find that the difference in revenue stems from the uncertainty about others’ valuations. Morgan *et al.*, on the other hand, also provide an equilibrium strategy for English auctions and discuss risk aversion, interpersonal comparisons, and “the love of winning” as alternative explanations for overbidding in auctions.

Interestingly, a special case of our results—the Bayes Nash equilibrium for two spiteful bidders in Vickrey auctions with a uniform prior (see Corollary 3)—was found independently by an algorithmic best-response solver (Reeves and Wellman, 2004).

3 Preliminaries

In this section, we define the utility function of rational spiteful agents and the framework of our auction setting.

3.1 Spiteful Agents

A spiteful agent i maximizes the weighted difference of his own utility u_i and his competitors’ utilities u_j for all $j \neq i$. In general, it would be reasonable to take the average or maximum of the competitors’ utilities. However, since we only consider single-item auctions where all utilities except the winner’s are zero, we can simply employ the sum of all remaining agents’ utilities.

Definition 1 *The utility of a spiteful agent is given by*

$$U_i = (1 - \alpha_i) \cdot u_i - \alpha_i \cdot \sum_{j \neq i} u_j$$

where $\alpha_i \in [0, 1]$ is a parameter called spite coefficient.

In the following, we speak of “utility” when referring to spiteful utility U_i and use the term “profit” to denote conventional utility u_i . Obviously, setting α_i to zero yields a self-interested agent (whose utility equals his profit) whereas $\alpha_i = 1$ defines a completely *malicious* agent whose only goal is to minimize the profit of other agents. When $\alpha_i = \frac{1}{2}$, we say an agent is *balanced spiteful*.²

As mentioned in Section 2, other authors have suggested utility functions with a linear trade-off between self-interest and others’ well-being. In contrast to these proposals, our definition differs in that the weight of one’s own utility is not normalized to 1, allowing us to capture malicious agents who have no self-interest at all. This opens interesting avenues for future research like the possibility to analyze the robustness of mechanisms in the presence of worst-case adversaries.

3.2 Auction Setting

Except for a preliminary result in Section 6.3, we assume that bidders are symmetric, in particular they all have the *same spite coefficient* α . Before each auction, private values v_i are drawn independently from a commonly known probability distribution over the interval $[0, 1]$ defined by the *cumulative distribution function (cdf)* $F(v)$. The *cdf* is defined as the probability that a random sample V drawn from the distribution does not exceed v : $F(v) = Pr(V \leq v)$. Its derivative, the *probability density function (pdf)*, is denoted by $f(v)$.

Once the auction starts, each bidder submits a bid based on his private value. The bidder who submitted the highest bid wins the auction. In a 1st-price auction, he pays the amount he bid whereas in a 2nd-price (or Vickrey) auction he pays the amount of the second highest bid. Extending the notation of Krishna (2002), we will denote equilibrium strategies of 1st- and 2nd-price auctions by $b_\alpha^1(v)$ and $b_\alpha^2(v)$, respectively. When bidders are self-interested ($\alpha = 0$), there are well-known equilibria for both auction types. The unique Bayes Nash equilibrium strategy for 1st-price auctions is to bid at the expectation of the second highest private value, conditional

²In the case of only two balanced spiteful agents, the game at hand becomes a zero-sum game.

on one's own value being the highest, $b_0^1(v) = E[X | X < v]$ where X is distributed according to $G(x) = F(x)^{n-1}$ (Vickrey, 1961; Riley and Samuelson, 1981). 2nd-price auctions are strategy-proof, *i.e.*, $b_0^1(v) = v$ for any distribution of values (Vickrey, 1961). Vickrey also first made the observation that expected revenue in both auction types is identical which was later generalized to a whole class of auctions in the revenue equivalence theorem (Myerson, 1981; Riley and Samuelson, 1981).

4 First-Price Auctions

As is common in auction theory, we study symmetric equilibria, that is, equilibria in which all bidders use the same bidding function (mapping from valuations to bids). Symmetric equilibria are considered the most reasonable equilibria, but in principle need not be the only ones (we will later provide an asymmetric equilibrium in auctions with malicious bidders). Furthermore, we guess that the bidding function is strictly increasing and differentiable over $[0, 1]$. These assumptions impose no restriction on the general setting. They are only made to reduce the search space.

Theorem 1 *A Bayes Nash equilibrium for spiteful bidders in 1st-price auctions is given by the bidding strategy*

$$b_\alpha^1(v) = E[X | X < v]$$

where X is drawn from $G_\alpha^1(x) = F(x)^{\frac{n-1}{1-\alpha}}$.

Proof: We start by introducing some notation. Let $W_i = (b_i(v_i) > b_{(1)}(v_{-i}))$ be the event that bidder i wins the auction, $b_{(1)}(v_{-i})$ be the highest of all bids except i 's, $v_{(1)}$ be the highest private value, and $\bar{v}_i(b)$ denote the inverse function of $b_i(v)$. We will use the short notation \bar{v} for $\bar{v}_{(1)}(b_i(v_i))$ to improve readability. It is important to keep in mind that \bar{v} is a function of $b_i(v_i)$, *e.g.*, when taking the derivative of the expected utility.

Recall that agent i knows his own private value v_i , but only has probabilistic beliefs about the remaining $n-1$ private values (and bids). Thus, the expected utility of a spiteful agent in a 1st-price auction is given by

$$E[U_i(b_i(v_i))] = (1-\alpha) \cdot Pr(W_i) \cdot (v_i - b_i(v_i)) - \alpha \cdot (1 - Pr(W_i)) \cdot (E[v_{(1)} | \neg W_i] - E[b_{(1)}(v_{-i}) | \neg W_i]). \quad (1)$$

We can ignore ties in this formulation because they are zero probability events in the continuous setting we consider. By definition, the probability that any private value is lower than i 's value is given by $F(v_i)$. Since all values are independently distributed, the probability that bidder i has the highest private value is $F(v_i)^{n-1}$. Thus, the probability that i submits the highest *bid* can be expressed by using the inverse bid function

$$Pr(W_i) = F(\bar{v})^{n-1}. \quad (2)$$

The *cdf* of the highest of $n-1$ private values is $F_{(1)}(v) = F(v)^{n-1}$. The associated *pdf* is $f_{(1)}(v) = (n-1)F(v)^{n-2} \cdot f(v)$. Using standard formulas for the conditional expectation (see

Appendix A), this allows us to compute both expectation values on the right-hand side of Equation 1. The expectation of the highest private value is

$$E[v_{(1)} | \neg W_i] = \frac{1}{1 - F(\bar{v})^{n-1}} \int_{\bar{v}}^1 t \cdot (n-1)F(t)^{n-2} \cdot f(t) dt \quad (3)$$

whereas the expectation of the highest bid is

$$E[b_{(1)}(v_{-i}) | \neg W_i] = \frac{1}{1 - F(\bar{v})^{n-1}} \int_{b_i(v_i)}^{b_i(1)} t \cdot (n-1) \cdot F(\bar{v}(t))^{n-2} \cdot f(\bar{v}(t)) \cdot \bar{v}'(t) dt.$$

Inserting these expectations in Equation 1 and simplifying the result yields

$$E[U_i(b_i(v_i))] = (1-\alpha)(F(\bar{v})^{n-1}v_i - F(\bar{v})^{n-1}b_i(v_i)) - \alpha(n-1) \left(\int_{\bar{v}}^1 t \cdot F(t)^{n-2} \cdot f(t) dt - \int_{b_i(v_i)}^{b_i(1)} t \cdot F(\bar{v}(t))^{n-2} \cdot f(\bar{v}(t)) \cdot \bar{v}'(t) dt \right).$$

When taking the derivative with respect to $b_i(v_i)$, both integrals vanish due to the fundamental theorem of calculus and the observation that

$$\frac{\partial \int_{\bar{v}(b)}^1 g(t) dt}{\partial b} = \frac{\partial (G(1) - G(\bar{v}(b)))}{\partial b} = 0 - g(\bar{v}(b)) \cdot \bar{v}'(b). \quad (4)$$

In order to obtain the strategy that generates maximum utility we take the derivative and set it to zero. Thus,

$$0 = (1-\alpha) \left((n-1)F(\bar{v})^{n-2} \cdot f(\bar{v}) \cdot \bar{v}' \cdot v_i - (n-1)F(\bar{v})^{n-2} \cdot f(\bar{v}) \cdot \bar{v}' \cdot b_i(v_i) - F(\bar{v})^{(n-1)} \right) - \alpha(n-1) \left((0 - \bar{v} \cdot F(\bar{v})^{n-2} \cdot f(\bar{v}) \cdot \bar{v}') - (0 - b_i(v_i) \cdot F(\bar{v})^{n-2} \cdot f(\bar{v}) \cdot \bar{v}') \right).$$

From this point on, we treat v_i as a variable (instead of $b_i(v_i)$) and assume that all bidding strategies are identical, *i.e.*, $\bar{v} = \bar{v}_{(1)}(b_i(v_i)) = v_i$. Using the fact that the derivative of the inverse function is the reciprocal of the original function's derivative ($\bar{v}'(b_i(v_i)) = \frac{1}{b_i'(v_i)}$), we can rearrange terms to obtain the differential equation

$$b(v) = v - \frac{(1-\alpha) \cdot F(v) \cdot b'(v)}{(n-1) \cdot f(v)}. \quad (5)$$

It follows that $b(v) \leq v$ because the fraction on the right-hand side is always non-negative (recall that the bidding function is strictly increasing). Since we assume that there are no negative bids, this yields the boundary condition $b(0) = 0$. The solution of Equation 5 with boundary condition $b(0) = 0$ is

$$b(v) = \frac{1}{F(v)^{\frac{n-1}{1-\alpha}}} \int_0^v t \cdot \frac{n-1}{1-\alpha} \cdot F(t)^{\frac{n-1}{1-\alpha}-1} \cdot f(t) dt.$$

Strikingly, the right-hand side of this equation is a conditional expectation (see Appendix A). More precisely, it is the expectation of the highest of $\frac{n-1}{1-\alpha}$ private values below v (ignoring the fact $\frac{n-1}{1-\alpha}$ is not necessarily an integer), *i.e.*,

$$b(v) = E[X | X < v]$$

where the *cdf* of X is given by $G_\alpha^1(x) = F(x)^{\frac{n-1}{1-\alpha}}$. It remains to be shown that the resulting strategy is indeed a mutual best response. We omit this step for reasons of limited space. \square

In 1st-price auctions, bidders face a tradeoff between the probability of winning and the profit conditional on winning. An intuition behind the equilibrium for spiteful agents is that the more spiteful a bidder is, the less emphasis he puts on his expected profit. Whereas a self-interested bidder bids at the expectation of the highest of $n - 1$ private values below his own value, a balanced spiteful agent bids at the expectation of the highest of $2(n - 1)$ private values below his value. Interestingly, agents are “least truthful” when they are self-interested. Any level of spite makes them more truthful. Furthermore, parameter α defines a continuum of Nash equilibria between the well-known standard equilibria of 1st-price and 2nd-price auctions. Even though $G_\alpha^I(x)$ is not defined for $\alpha = 1$, it can easily be seen from Equation 5, that $b_1^I(v) = v$.

Corollary 1 *The 1st-price auction is (Bayes Nash) incentive-compatible for malicious bidders ($\alpha = 1$).*

This result is perhaps surprising because one might expect that always bidding 1 is an optimal strategy for malicious bidders. The following consideration shows why this is not the case. Assume that all agents are bidding 1. Agent i 's expected utility depends on the tie resolution policy. If another bidder is chosen as the winner, i 's expected utility is positive. If he wins the auction, his utility is zero. By bidding less than 1, he can ensure that his expected utility is always positive.

Curiously, there are other, *asymmetric*, equilibria for malicious bidders, e.g., a “threat” equilibrium where one designated bidder always bids 1 and everybody else bids some value below his private value. It is well-known that asymmetric equilibria like this exist in 2nd-price auctions (see Blume and Heidhues, 2004, for a complete characterization). However, asymmetric equilibria in 2nd-price auctions are (weakly) *dominated* whereas the one given above is not, making it more reasonable.

One way to gain more insight in the equilibrium strategy is to instantiate $F(v)$ with the uniform distribution.

Corollary 2 *A Bayes Nash equilibrium for spiteful bidders in 1st-price auctions with uniformly distributed private values is given by the bidding strategy*

$$b_\alpha^I(v) = \frac{n-1}{n-\alpha} \cdot v.$$

Whereas one can get full intuition in the extreme points of the strategy ($\alpha \in \{0, 1\}$), the fact that the scaling between both endpoints of the equilibrium spectrum is not linear in α , even for a uniform prior, is somewhat surprising.

5 Second-Price Auctions

In this section, we derive an equilibrium strategy for spiteful agents in 2nd-price auctions using the same set of assumptions made in Section 4.

Theorem 2 *A Bayes Nash equilibrium for spiteful bidders in 2nd-price auctions is given by the bidding strategy*

$$b_\alpha^{II}(v) = E[X | X > v]$$

where X is drawn from $G_\alpha^{II}(x) = 1 - (1 - F(x))^\frac{1}{\alpha}$.

Proof: We use the same notation introduced in the proof of Theorem 1. The expected utility of spiteful agent i in 2nd-price auctions can be described as follows. There are two general cases depending on whether bidder i wins or loses. In the former case, the utility is simply v_i minus the expected highest bid (except i 's). In the latter case, we have to compute the expectations of the winner's private value and the selling price. In order to specify the selling price, we need to distinguish between two subcases: If bidder i submitted the second highest bid, the selling price is his bid b_i . Otherwise, i.e., if the second highest of all remaining bids is greater than b_i , we can again give a conditional expectation. Thus, the overall expected utility of agent i is

$$\begin{aligned} E[U_i(b_i(v_i))] &= (1 - \alpha) \cdot Pr(W_i) \cdot (v_i - E[b_{(1)}(v_{-i}) | W_i]) - \\ &\alpha \cdot ((1 - Pr(W_i)) \cdot E[v_{(1)} | \neg W_i] - \\ &Pr((b_i(v_i) < b_{(1)}(v_{-i})) \wedge (b_i(v_i) > b_{(2)}(v_{-i}))) \cdot b_i(v_i) - \\ &Pr(b_i(v_i) < b_{(2)}(v_{-i})) \cdot E[b_{(2)}(v_{-i}) | b_i < b_{(2)}(v_{-i})]). \end{aligned} \quad (6)$$

According to the formula given in Appendix A, the conditional expectation of the remaining highest bid, in case bidder i wins, is

$$\begin{aligned} E[b_{(1)}(v_{-i}) | W_i] &= \\ \frac{1}{F(\bar{v})^{n-1}} \int_{b_i(0)}^{b_i(v_i)} t \cdot (n-1)F(\bar{v}(t))^{n-2} \cdot f(\bar{v}(t)) \cdot \bar{v}'(t) dt. \end{aligned} \quad (7)$$

We have already given a formula for $E[v_{(1)} | \neg W_i]$ in Equation 3. The probability that b_i is the second highest bid equals the probability that exactly one bid is greater than b_i and $n - 2$ bids are less than b_i . Depending on who submitted the highest bid, there are $n - 1$ different ways in which this can occur, yielding

$$\begin{aligned} Pr((b_i(v_i) < b_{(1)}(v_{-i})) \wedge (b_i(v_i) > b_{(2)}(v_{-i}))) &= \\ (n-1)F(\bar{v})^{n-2} \cdot (1 - F(\bar{v})). \end{aligned}$$

The *cdf* of the second highest private value (of $n - 1$ values) can be derived by computing the probability that the second highest value is less than or equal to a given v . Either all $n - 1$ values are lower than v , or $n - 2$ values are lower and one is greater than v . As above, there are $n - 1$ different possibilities in the latter case. Thus,

$$\begin{aligned} F_{(2)}(v) &= F(v)^{n-1} + (n-1)F(v)^{n-2}(1 - F(v)) = \\ &(n-1)F(v)^{n-2} - (n-2)F(v)^{n-1}. \end{aligned}$$

It follows that the *pdf* is $f_{(2)}(v) = (n-1) \cdot (n-2) \cdot (1 - F(v)) \cdot F(v)^{n-3} \cdot f(v)$. Finally, the conditional expectation of the second highest bid times the probability of this bid being higher than b_i is

$$\begin{aligned} Pr(b_i(v_i) < b_{(2)}(v_{-i})) \cdot E[b_{(2)}(v_{-i}) | b_i < b_{(2)}(v_{-i})] &= \\ (n-1) \cdot (n-2) \cdot \int_{b_i(v_i)}^{b_i(1)} t \cdot (1 - F(\bar{v}(t))) \cdot F(\bar{v}(t))^{n-3} \cdot f(\bar{v}(t)) \cdot \bar{v}'(t) dt. \end{aligned}$$

Inserting both expectations and the probability of winning (see Equation 2) into Equation 6 yields

$$\begin{aligned}
E[U_i(b_i(v_i))] &= (1 - \alpha) \cdot \left(F(\bar{v})^{n-1} v_i - \right. \\
&(n-1) \cdot \int_{b_i(0)}^{b_i(v_i)} t \cdot F(\bar{v}(t))^{n-2} \cdot f(\bar{v}(t)) \cdot \bar{v}'(t) dt - \\
&\alpha \cdot (n-1) \cdot \left(\int_{\bar{v}}^1 t \cdot F(t)^{n-2} \cdot f(t) dt - \right. \\
&F(\bar{v})^{n-2} \cdot (1 - F(\bar{v})) \cdot b_i(v_i) - \\
&(n-2) \cdot \int_{b_i(v_i)}^{b_i(1)} t \cdot (1 - F(\bar{v}(t))) \cdot F(\bar{v}(t))^{n-3} \cdot f(\bar{v}(t)) \cdot \bar{v}'(t) dt \left. \right).
\end{aligned}$$

As in the previous section, we now take the derivative with respect to $b_i(v_i)$ and set it to zero. All integrals vanish due to the Fundamental Theorem of Calculus and the formula given in Equation 4. We get

$$\begin{aligned}
0 &= (1 - \alpha) \cdot \left((n-1) F(\bar{v})^{n-2} \cdot f(\bar{v}) \cdot \bar{v}' \cdot v_i - \right. \\
&(n-1) (b_i(v_i) \cdot F(\bar{v})^{n-2} \cdot f(\bar{v}) \cdot \bar{v}') - \\
&\alpha \cdot (n-1) \cdot \left((0 - \bar{v} \cdot F(\bar{v})^{n-2} \cdot f(\bar{v}) \cdot \bar{v}') - \right. \\
&((n-2) \cdot F(\bar{v})^{n-3} \cdot f(\bar{v}) \cdot \bar{v}' \cdot b_i(v_i) + F(\bar{v})^{n-2} - \\
&(n-1) \cdot F(\bar{v})^{n-2} \cdot f(\bar{v}) \cdot \bar{v}' \cdot b_i(v_i) - F(\bar{v})^{n-1}) - \\
&(n-2) \cdot (0 - b_i(v_i) \cdot (1 - F(\bar{v})) \cdot F(\bar{v})^{n-3} \cdot f(\bar{v}) \cdot \bar{v}') \left. \right).
\end{aligned}$$

Using the fact that the derivative of the inverse function is the reciprocal of the original function's derivative ($\bar{v}'(b_i(v_i)) = \frac{1}{b_i'(v_i)}$) and $\bar{v} = v_i$, we can simplify and rearrange terms to obtain the differential equation

$$b(v) = v + \frac{\alpha \cdot (1 - F(v)) \cdot b'(v)}{f(v)}. \quad (8)$$

It turns out that $b(0) = 0$ does not hold for 2nd-price auctions. However, a boundary condition can easily be obtained by letting $v = 1$. By definition, $F(1) = 1$ which yields $b(1) = 1$. Given this boundary condition, the solution of Equation 8 is

$$b(v) = \frac{1}{(1 - F(v))^{\frac{1}{\alpha}}} \int_v^1 \frac{t \cdot (1 - F(t))^{\frac{1}{\alpha} - 1} \cdot f(t)}{\alpha} dt. \quad (9)$$

Like in proof of Theorem 1, the right-hand side of Equation 9 resembles a conditional expectation. In fact, the bidding strategy can be reformulated as the expectation of some random variable X , given that $X > v$,

$$b(v) = E[X \mid X > v]$$

where the *cdf* of X is given by $G_\alpha^{\text{II}}(x) = 1 - (1 - F(x))^{\frac{1}{\alpha}}$.

It can easily be checked that $G_\alpha^{\text{II}}(x)$ is indeed a valid *cdf* ($G_\alpha^{\text{II}}(0) = 0$, $G_\alpha^{\text{II}}(1) = 1$, and $G_\alpha^{\text{II}}(x)$ is non-decreasing and differentiable). By inserting this *cdf* in Equation 12, we obtain the equilibrium bidding strategy. The resulting expectation is the expected value of the lowest of $\frac{1}{\alpha}$ values above v . $G_\alpha^{\text{II}}(x)$ is not defined for $\alpha = 0$, but the correct, well-known, equilibrium can quickly be read from Equation 8. It remains to be shown that the resulting strategy is indeed a mutual best response. We omit this step for reasons of limited space. \square

Remarkably, the resulting equilibrium strategy is independent of the number of bidders n (though it *does* depend on the prior distribution of private values). For example, a balanced spiteful bidder bids at the expectation of the lowest of two private values above his own value. As in the previous section, we try to get more insight in the equilibrium by instantiating the uniform distribution.

Corollary 3 *A Bayes Nash equilibrium for spiteful bidders in 2nd-price auctions with uniformly distributed private values is given by the bidding strategy*

$$b_\alpha^{\text{II}}(v) = \frac{v + \alpha}{1 + \alpha}$$

For example, given a uniform prior, the optimal strategy for balanced spiteful agents is $b(v) = \frac{2}{3} \cdot v + \frac{1}{3}$, regardless of the number of bidders. As in the 1st-price auction setting, the surprising equilibrium strategies are those for $0 < \alpha < 1$. There is no linear scaling between both extreme points of the equilibrium spectrum. As we will see in the following section, this leads to important consequences on auction revenue.

6 Consequences

In order to obtain instructive results from these equilibria, we compare a key measure in auction theory—the seller's revenue—and investigate the impact of common knowledge on bidding and revenue.

6.1 Revenue Comparison

The well-known revenue equivalence theorem, which states that members of a large class of auctions all yield the same revenue under certain conditions, does not hold when agents are spiteful. Figure 1 shows the expected revenue in both auction types when agents are balanced spiteful and private values are uniformly distributed.

It can be shown that the revenue gap visible in the figure exists for any prior and spite coefficient as long as agents are neither self-interested nor malicious.

Theorem 3 *For the same spite coefficient $0 < \alpha < 1$, the 2nd-price auction yields more expected revenue than the 1st-price auction. When $\alpha \in \{0, 1\}$, expected revenue in both auction types is equal in the symmetric equilibrium.*

Proof: The statement can be deduced from the following three observations:

- *When agents are malicious, expected revenue in both auction types is identical.*

In the 1st-price auction, truthful bidding is in equilibrium. In the 2nd-price auction, the second highest bidder bids at the expectation of the highest private value. In both cases, revenue equals the expectation of the highest value.

- *$b_\alpha^{\text{I}}(v)$ and $b_\alpha^{\text{II}}(v)$ are strictly increasing in α .*

In the 1st-price auction, bidders bid at the (conditional) expectation of the highest value of a number of private values that *increases* as α grows. In the 2nd-price auction, bidders bid at the expectation of the lowest




Figure 1: Expected revenue (interpolated for non-integer n)

value of a number of private values that *decreases* as α grows. Obviously, both expectations are increasing in α . More formally, G_α^I stochastically dominates G_β^I and G_α^{II} stochastically dominates G_β^{II} for any $\alpha > \beta$.

- $b_\alpha^I(v)$ is convex in α . $b_\alpha^{II}(v)$ is concave in α .
Since both equilibria are symmetric, we just need to consider the curvature of expectations distributed according to $G_\alpha^I(X)$ and $G_\alpha^{II}(X)$ for variable α . Bids in 1st-price auctions are the (conditional) expectation of the highest of $\frac{1}{1-\alpha}$ values. The slope of this expectation increases as α rises. In 2nd-price auctions, bids are the expectation of the *lowest* of $\frac{1}{\alpha}$ values. If it were the highest value, the slope would be increasing too. However, since it is the expectation of the lowest value, the slope is strictly decreasing in α .

Let $E[R_\alpha^I]$ and $E[R_\alpha^{II}]$ be the expected revenue in 1st- and 2nd-price auctions, respectively, and consider these as functions of α . So far, we know that both functions are equal for $\alpha \in \{0, 1\}$ and strictly increasing. Furthermore, $E[R_\alpha^I]$ is convex and $E[R_\alpha^{II}]$ is concave. These facts imply that $E[R_\alpha^{II}] > E[R_\alpha^I]$ for any $0 < \alpha < 1$ (see Figure 2). \square

Naturally, more revenue for the seller results in less profit for the bidders. However, if you look at (spiteful) *utility*, the utility of winning bidders in 2nd-price auctions is lower than in 1st-price auctions, whereas the utility of losing bidders is higher in 2nd-price auctions. As a consequence, social welfare (if one is willing to consider such a notion in a setting of spitefulness) is higher in 2nd-price auctions than in 1st-price auctions if the number of bidders is sufficiently large.

Revenue inequalities for other special conditions such as when bidders or the seller are risk-averse have been used to argue in favor of one auction form over another. Hence, Theorem 3 can be interpreted as an advantage of the 2nd-price auction (from the perspective of a seller) because it yields more revenue than the 1st-price auction whenever bidders exhibit the slightest interest in reducing their competitors' profit

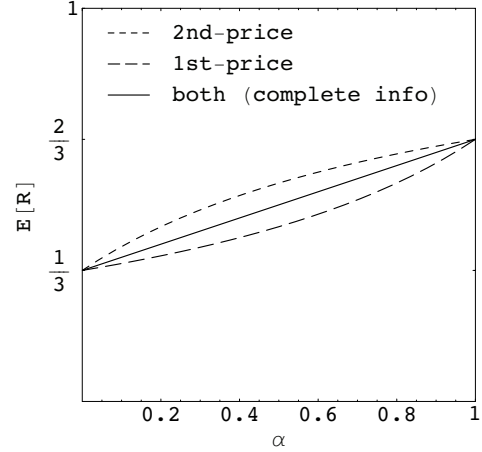


Figure 2: Expected revenue for $n = 2$ and varying α

(and they still care about their own profit). On the other hand, the difference in expected revenue is relatively small, even for just few bidders. For example, the difference in expected revenue for ten bidders with uniformly distributed private values is less than 2% for any α . Interestingly, the revenue difference is maximal for some α slightly below 0.5.

Theorem 4 *The difference in expected revenue between 2nd-price and 1st-price auctions is maximal for some $\alpha \leq 0.5$ that approaches $\frac{1}{1+\sqrt{2}} \approx 0.4142$ in the limit as n rises, when private values are uniformly distributed.*

Proof: By definition, the revenue difference is the difference of the expectation of the second highest bid in 2nd-price auctions minus the highest expected bid in 1st-price auctions. Instantiating with the uniform distribution, we get

$$\begin{aligned} E[R_\alpha^{II}] - E[R_\alpha^I] &= b_\alpha^{II}(E[v_{(2)}]) - b_\alpha^I(E[v_{(1)}]) \\ &= \frac{\frac{n-1}{n+1} + \alpha}{(1-\alpha) \cdot \alpha} - \frac{n-1}{n-\alpha} \cdot \frac{n}{n+1} \quad (10) \\ &= \frac{1}{(\alpha+1) \cdot (n-\alpha)} \geq 0. \end{aligned}$$

In order to obtain the maximal revenue difference, we take the derivative of the expression given in Equation 10 with respect to α and set it to zero:

$$\begin{aligned} 0 &= \frac{\partial}{\partial \alpha} \cdot \frac{(1-\alpha) \cdot \alpha}{(\alpha+1) \cdot (n-\alpha)} = \frac{\partial}{\partial \alpha} \cdot \frac{\alpha^2 - \alpha}{\alpha - n + \alpha^2 - \alpha \cdot n} = \\ &= \frac{(\alpha^2 - \alpha)(2 \cdot \alpha - n + 1)}{(\alpha^2 + \alpha - n - \alpha \cdot n)^2} + \frac{2 \cdot \alpha - 1}{\alpha^2 + \alpha - n - \alpha \cdot n} \\ \Rightarrow \alpha_{max} &= \frac{n}{n + \sqrt{2} \cdot \sqrt{n \cdot (n-1)}}. \end{aligned}$$

When there are only two bidders, $\alpha_{max} = 0.5$. α_{max} is strictly decreasing as n increases and

$$\lim_{n \rightarrow \infty} \alpha_{max} = \frac{1}{1 + \sqrt{2}} \approx 0.4142.$$

\square

6.2 Complete Information

In this section, we give bidding equilibria for spiteful agents in a model with complete information, *i.e.*, all private values are common knowledge. This allows us to examine the influence of uncertainty on spiteful bidding.

Theorem 5 *A Nash equilibrium for both 1st- and 2nd-price auctions in a model of complete information is given by the bidding strategy profile*

$$\begin{aligned} b_{i,\alpha}^I = b_{i,\alpha}^{II} &= v_2 + \alpha(v_1 - v_2) \text{ for } i \in \{1, 2\} \text{ and} \\ b_{i,\alpha}^I = b_{i,\alpha}^{II} &< v_2 + \alpha(v_1 - v_2) \text{ for } i \in \{3, 4, \dots, n\} \end{aligned}$$

where 1 and 2 are the indices of the bidder with the highest and second highest valuation, respectively.³

Proof: Let us first consider why bidder 1 and 2 would not deviate from the given strategy profile in a 1st-price auction. Bidder 1's utility for bidding b_1 , given that bidder 2 bids according to the equilibrium strategy is

$$\begin{aligned} U_\alpha^I(b_1) &= \begin{cases} -\alpha(v_2 - b_{2,\alpha}^I) & \text{if } b_1 \leq b_{1,\alpha}^I \\ (1 - \alpha)(v_1 - b_1) & \text{if } b_1 \geq b_{1,\alpha}^I \end{cases} \\ &= \begin{cases} \alpha^2(v_1 - v_2) & \text{if } b_1 \leq b_{1,\alpha}^I \\ (1 - \alpha)(v_1 - b_1) & \text{if } b_1 \geq b_{1,\alpha}^I \end{cases}. \end{aligned}$$

Bidder 1 cannot increase his utility by deviating from the equilibrium strategy. If he bids less, his utility stays the same. If he bids more, his utility is diminishing (it is less than $(1 - \alpha)^2(v_1 - v_2)$). The same holds for bidder 2 whose utility is

$$\begin{aligned} U_\alpha^I(b_2) &= \begin{cases} -\alpha(v_1 - b_{1,\alpha}^I) & \text{if } b_2 \leq b_{2,\alpha}^I \\ (1 - \alpha)(v_2 - b_2) & \text{if } b_2 \geq b_{2,\alpha}^I \end{cases} \\ &= \begin{cases} -(1 - \alpha)\alpha(v_1 - v_2) & \text{if } b_2 \leq b_{2,\alpha}^I \\ (1 - \alpha)(v_2 - b_2) & \text{if } b_2 \geq b_{2,\alpha}^I \end{cases}. \end{aligned}$$

The equilibrium point is exactly the strategy for which bidder 2 is indifferent between winning and losing since both payoffs are equal. It also coincides with his maximin strategy, *i.e.*, the strategy that guarantees the highest payoff regardless of other players' rationality (see also Brandt and Weiß, 2001). It follows that the remaining bidders have no incentive to interfere (by bidding at least as much as bidder 1 and 2) because their utility would only decrease.

In 2nd-price auctions, the argumentation is analogous. When the other bidders employ the equilibrium strategy, bidder 1's utility is

$$\begin{aligned} U_\alpha^{II}(b_1) &= \begin{cases} -\alpha(v_2 - b_1) & \text{if } b_1 \leq b_{1,\alpha}^{II} \\ (1 - \alpha)(v_1 - b_{2,\alpha}^{II}) & \text{if } b_1 \geq b_{1,\alpha}^{II} \end{cases} \\ &= \begin{cases} -\alpha(v_2 - b_1) & \text{if } b_1 \leq b_{1,\alpha}^{II} \\ (1 - \alpha)^2(v_1 - v_2) & \text{if } b_1 \geq b_{1,\alpha}^{II} \end{cases}. \end{aligned}$$

³Handling ties introduces some unnecessary complications to the equilibrium strategy (involving the minimum bid increment ϵ). We brush aside these complications by assuming that whenever $\alpha < 0.5$, bidder 1 wins and whenever $\alpha > 0.5$, bidder 2 wins in the case of a tie. If $\alpha = 0.5$, ties can be resolved either way.

Bidding more will not change anything and bidding less results in less utility ($U_\alpha^{II}(b_1) < \alpha^2(v_1 - v_2)$). Like above, bidder 2, whose utility is

$$\begin{aligned} U_\alpha^{II}(b_2) &= \begin{cases} -\alpha(v_1 - b_2) & \text{if } b_2 \leq b_{2,\alpha}^{II} \\ (1 - \alpha)(v_2 - b_{1,\alpha}^{II}) & \text{if } b_2 \geq b_{2,\alpha}^{II} \end{cases} \\ &= \begin{cases} -\alpha(v_1 - b_2) & \text{if } b_2 \leq b_{2,\alpha}^{II} \\ (1 - \alpha)\alpha(v_1 - v_2) & \text{if } b_2 \geq b_{2,\alpha}^{II} \end{cases}, \end{aligned}$$

is indifferent between winning and losing in equilibrium. \square

Since the main purpose of considering the complete information model is a comparison with the equilibria given in Sections 4 and 5, we just provided an equilibrium for symmetric spite. Computing equilibria for any given profile of spite coefficients is straightforward in the complete information model.

Apparently, equilibria for 1st- and 2nd-price auctions are identical and scale linearly between the second highest and highest valuation (see Figure 2). This has interesting consequences on the availability of information and expected revenue: Whereas revealing private values increases expected revenue in 1st-price auctions, it decreases revenue in 2nd-price auctions, whenever $0 < \alpha < 1$. This effect is quite surprising because expected revenue in a setting with self-interested bidders is identical in the incomplete and complete information model (see *e.g.*, Osborne, 2004).

6.3 Asymmetries

An important extension of our setting is one that deals with asymmetries in spitefulness. For example, it would be very desirable to extend the revenue inequality (Theorem 3) to arbitrary profiles of spite coefficients $(\alpha_1, \alpha_2, \dots, \alpha_n)$ or a general prior from which each α_i is drawn. A first step towards this direction can be made by observing that the equilibrium strategies of self-interested bidders are in a sense "robust" against spiteful bidding.

Proposition 1 *Rational self-interested bidders will stick with their bidding strategy when other agents bid according to the strategies given in Theorems 1 and 2, respectively, and private values are uniformly distributed.*

Proof: The statement for 1st-price auctions follows from a result by Porter and Shoham (2003) who proved that bidders in 1st-price auctions will stick with their equilibrium strategy even when other bidders bid constant fractions of their private value larger than $\frac{n-1}{n} \cdot v$ (in the case of a uniform prior). This holds for a certain class of probability distributions, including the uniform distribution.

The statement for 2nd-price auctions trivially follows from the fact that bidding truthfully is a dominant strategy for self-interested agents and therefore holds for any given prior. \square

The previous proposition can be interpreted as a setting in which there are self-interested and spiteful agents participating in the same auction. Self-interested agents are aware of this asymmetry whereas spiteful agents believe that everybody is spiteful.

7 Conclusion

We studied the bidding behavior of spiteful agents who, contrary to the common assumption of self-interest, maximize a convex combination of their own profit and their competitors' losses. We derived symmetric Bayes Nash equilibria for spiteful agents in 1st-price and 2nd-price sealed-bid auctions. The main results are as follows. In 1st-price auctions, bidders become "more truthful" the more spiteful they are. When bidders are completely malicious, truth-telling is in Nash equilibrium. Surprisingly, the equilibrium strategy in 2nd-price auctions does not depend on the number of bidders. Based on these equilibria, we compared the revenue in both auction types. It turned out that revenue equivalence breaks down for this setting. Expected revenue in 2nd-price auctions is higher than revenue in 1st-price auctions whenever the spite coefficient α satisfies $0 < \alpha < 1$. However, revenue equivalence holds at each extreme: auctions where all agents are self-interested ($\alpha = 0$) and auction where all agents are malicious ($\alpha = 1$). We showed that the difference in revenue stems from the uncertainty about others' valuations. Whereas revealing private values increases expected revenue in 1st-price auctions, it decreases revenue in 2nd-price auctions if $0 < \alpha < 1$.

There are several open problems left for future work. Most importantly, we intend to extend the revenue inequality (Theorem 3) to settings with *asymmetric spite* and investigate the *mechanism design* problem for spiteful agents.

Acknowledgements

We thank Christina Fong, Kate Larson, Chris Luhrs, Ryan Porter, Michael Rothkopf, Moshe Tennenholtz, Gerhard Weiß, and Michael Wellman for valuable comments.

This material is based upon work supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/1-1 and by the National Science Foundation under ITR grants IIS-0121678, IIS-0427858, and IIS-0205633, and a Sloan Fellowship.

A Conditional Expectations

Let X be a random variable drawn from the interval $[0, 1]$ according to the cumulative distribution function $F(x)$. The expectation of X is $E[X] = \int_0^1 t \cdot f(t) dt$. The conditional expectation that X is smaller or greater than some constant x , respectively, is given by

$$E[X | X < x] = \frac{1}{F(x)} \int_0^x t \cdot f(t) dt, \quad \text{and} \quad (11)$$

$$E[X | X > x] = \frac{1}{1 - F(x)} \int_x^1 t \cdot f(t) dt. \quad (12)$$

References

A. Blume and P. Heidhues. All equilibria of the Vickrey auction. *Journal of Economic Theory*, 114(1):170–177, 2004.

S. Brainov. The role and the impact of preferences on multi-agent interaction. In *Intelligent Agents VI*, volume 1757 of *Lecture Notes in Artificial Intelligence (LNAI)*. Springer-Verlag, 1999.

F. Brandt and G. Weiß. Antisocial agents and Vickrey auctions. In J.-J. Ch. Meyer and M. Tambe, editors, *Intelligent Agents VIII*, volume 2333 of *Lecture Notes in Artificial Intelligence (LNAI)*, pages 335–347. Springer-Verlag, 2001.

E. Fehr and K. Schmidt. The economics of fairness, reciprocity and altruism - Experimental evidence and new theories. In *Handbook of the Economics of Giving, Altruism and Reciprocity*, pages 615–692. North-Holland, 2006.

V. Grimm, F. Riedel, and E. Wolfstetter. The third generation (UMTS) spectrum auction in Germany. *ifo Studien*, 48(1), 2002.

P. Jehiel, B. Moldovanu, and E. Stacchetti. How (not) to sell nuclear weapons. *American Economic Review*, 86:814–829, 1996.

V. Krishna. *Auction Theory*. Academic Press, 2002.

D. K. Levine. Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1:593–622, 1998.

G. Loewenstein, L. Thompson, and M. Bazerman. Social utility and decision making in interpersonal contexts. *Journal of Personality and Social Psychology*, 57(3):426–441, 1989.

E. Maasland and S. Onderstal. Auctions with financial externalities. Working Paper, 2003.

D. M. Messick and K. P. Sents. Estimating social and nonsocial utility functions from ordinal data. *European Journal of Social Psychology*, 15:389–399, 1985.

J. Morgan, K. Steiglitz, and G. Reis. The spite motive and equilibrium behavior in auctions. *Contributions to Economic Analysis & Policy*, 2(1):1102–1127, 2003.

R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

M. Osborne. *An Introduction to Game Theory*. Oxford University Press, 2004.

R. Porter and Y. Shoham. On cheating in sealed-bid auctions. In *Proceedings of the 4th ACM Conference on Electronic Commerce (ACM-EC)*, pages 76–84. ACM Press, 2003.

D. Reeves and M. Wellman. Computing best-response strategies in infinite games of incomplete information. In *Proceedings of the 20th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 470–478. AUAI Press, 2004.

J. Riley and W. Samuelson. Optimal auctions. *American Economic Review*, 71:381–392, 1981.

T. Saijo and H. Nakamura. The "spite" dilemma in voluntary contribution mechanism experiments. *Journal of Conflict Resolution*, 39(3):535–560, 1995.

J. Sobel. Interdependent preferences and reciprocity. *Journal of Economic Literature*, 43(2):392–436, 2005.

W. Vickrey. Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

Y. Zhou and R. Lukose. Vindictive bidding in keyword auctions. In *Proceedings of the 2nd Workshop on Sponsored Search Auctions (in conjunction with ACM-EC)*, 2006.