Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)

Asymmetric Spite in Auctions*

Ankit Sharma and **Tuomas Sandholm**
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
Email: {ankits, sandholm}@cs.cmu.edu

Abstract

In many auctions, agents bid more aggressively than self-interest would prescribe. This can be explained by spite, where the agent's utility not only increases in the agent's surplus but also decreases as the other bidders' surpluses increase. Spite can stem from long-term benefits from making competitors worse off and from inherent psychological effects. There have been important recent game-theoretic analyses of spiteful bidding assuming all agents are equally spiteful. We present, to our knowledge, the first auction analysis in the more realistic setting where bidders may be spiteful to different extents. We show that the equilibrium bidding function can still be written in the same form—except that the spite factor is replaced by an 'expressed' spite factor. This leads to bidders expressing spites that are higher or lower than their true spite depending on others' spite. Perhaps surprisingly, in the two-bidder case, the mapping from true spite to expressed spite is the same across all common auction mechanisms. Furthermore, even with two bidders, important properties of symmetric-spite settings cease to hold: the allocation can be inefficient and the revenue ranking may reverse between first- and second-price auctions. We also show that in sealed-bid auctions under asymmetric valuation distributions, there can be a "bargaining problem" in selecting bids. Finally, we study the generalization where agents can have different extents of spite toward different other bidders.

Introduction

Auctions have emerged as effective ways of allocating resources and tasks among human and software agents. Most of the auction literature assumes that each agent only cares about her own surplus: what goods she gets and how much she has to pay. However, in reality agents often have *other-regarding* preferences where they care about others' surpluses too. This can take the form of altruism, or more commonly in auctions and similar settings, spite. The spite motive, which is the preference to make others worse off, stems from mainly two reasons. The first reason is strategic. The agent might benefit in the long run by weakening her competitors, for example, driving competitors' market share down or causing them to have to pay more for a given allocation in the auction (as has been observed in spectrum auctions (Grimm, Riedel, and Wolfstetter 2001) and sponsored search auctions (Zhou and Lukose 2006)). Furthermore, in certain competitions such as the Trading Agents Competition, agents might give more weight to relative rankings rather than absolute performance. The second reason is psychological. There is ample evidence from experimental economics and psychology that people behave against their self-interest in strategic settings, and that this can be explained as rational behavior among agents that inherently have other-regarding preferences (Saijo and Nakamura 1995), (Levine 1998), (Loewenstein, Thompson, and Bazerman 1989).

Game-theoretic analysis of spiteful bidding in auctions was initiated relatively recently (Brandt and Weiß 2001). Spite can explain why people bid more aggressively in auctions than theory would predict among self-interested agents (Morgan, Steiglitz, and Reis 2003). (Brandt, Sandholm, and Shoham 2007) and (Morgan, Steiglitz, and Reis 2003) discuss the scenario where each bidder is equally spiteful, and give the equilibrium bidding functions for the first- and second-price one-item auctions. (Vetsikas and Jennings 2007) extend the analysis of the symmetric-spite setting to multi-unit auctions.

This prior literature has assumed that all bidders are equally spiteful. A priori, however, there is no reason to believe that each bidder would be equally spiteful (Brainov 2000). Different bidders can care to a different extent about the surplus of other bidders. Moreover, a bidder might care more for the surplus of a particular bidder than for the surplus of some other bidder. In this paper we present, to our knowledge, the first auction analysis of the broader setting where bidders can be asymmetrically spiteful.

Model

In this paper we study 1-item auctions. Making the standard assumptions of quasilinear utility functions and that losers in the auction pay nothing, we have that in the absence of spite, the utility function of bidder X is

$$u_X = \begin{cases} v_X - p_X & \text{if } X \text{ wins the auction} \\ 0 & \text{if } X \text{ loses the auction} \end{cases}$$

where v_X is the bidder's valuation of the item and p_X is the amount the bidder has to pay.

*This work was supported by NSF grant IIS-0905390.

In presence of spite, the utility function is

$$u_X = \begin{cases} v_X - p_X & \text{if } X \text{ wins the auction} \\ -\alpha_X^t \cdot (v_Y - p_Y) & \text{if } Y \neq X \text{ wins the auction} \end{cases}$$

$(0 \leq \alpha_X^t \leq 1)$

where α_X^t is a measure of the spite of agent X. The subscript X in α_X^t is to emphasize that the spite factor depends on the bidder X. The superscript t, where t stands for 'true', is there to distinguish α_X^t from symbol α_X^e which we will introduce later in the paper. Higher α_X^t means greater spite. In the symmetric model of spite, which has been considered in the prior work, all agents are equally spiteful ($\forall X, \alpha_X^t = \alpha^t$).

The utility can also be expressed in terms of surplus:

$$u_X = \begin{cases} surplus(X) & \text{if } X \text{ wins the auction} \\ -\alpha_X^t \cdot (surplus(Y)) & \text{if } Y \neq X \text{ wins the auction} \end{cases}$$

So, conditional on losing the auction, the bidder would like to minimize the surplus of the winning bidder. We assume $0 \leq \alpha_X^t \leq 1$, that is, bidders care at least as much for their own surplus as the negation of anyone else's surplus.

Among self-interested agents, it is weakly dominant for each bidder to bid her true valuation in a second-price sealed-bid auction (Vickrey 1961). The following example shows that this ceases to be the case among (even symmetrically) spiteful bidders. Let there be two bidders, A and B. Let $v_A = 5$ and $v_B = 10$, and $\alpha_A^t = \alpha_B^t = 0.1$. If both bid truthfully, B wins and pays 5 (second highest bid). B's surplus is $10 - 5 = 5$, so $u_A = -0.1 \cdot 5 = -0.5$. But, for example, A can get higher utility $u_A = -0.1 \cdot 2 = -0.2$ by bidding 8, thus causing B to pay 8.

We consider the four common auction mechanisms: first-price sealed-bid, second-price sealed-bid, English, and Dutch. In the first-price (second-price) sealed-bid auction, all bidders submit their bid in a sealed envelope and the bidder with the highest bid wins the auction and pays her bid price (second-highest bid price). The English auction is modeled with a clock displaying the current bid price. The clock price increases continuously and each bidder has a button which she releases when she wants to drop out. When the second-to-last bidder releases her button, the auction ends and the remaining bidder wins at the current price. The Dutch auction is modeled with a clock where the price decreases continuously. The first bidder to release her button wins and pays the price on the clock at that point.

Spite in the discrete valuations setting

Before discussing spiteful bidding in the case where the bidders draw their valuation from a continuous distribution, we first discuss the discrete case which has some characteristics worth noting and which gives insight into the issues that arise in spiteful bidding.

Complete information setting

Consider the example above, but now in a first-price sealed-bid auction. If A were to bid too low, say 3, then B could bid 4 and win the item with surplus $10 - 4 = 6$. A's utility would be $-0.1 \cdot 6 = -0.6$. So A must bid higher in order to force B to bid higher, thereby reducing B's surplus. How high can A go? Once A's bid overshoots his true valuation of 5, there is the risk that B does not bid higher which means A wins the item at a price greater than his own valuation thereby ending up with a negative utility. So, let us try to calculate the bid price at which even if A wins the item, he would get the same utility as in the case he loses.

Let the winning price be ρ. If A wins, his utility would be $(5 - \rho)$ and if he loses, it would be $-0.1 \cdot (10 - \rho)$. Equating these two yields $\rho \approx 5.5$. Hence, A is indifferent between winning and losing at that bid price. At any bid price above that, A would prefer to lose. At any bid price below that, A would prefer to win. We call 5.5 the *crossover point* for A.

Similarly, we can calculate the crossover point for B, which is approximately 9.5. Above that price B would prefer to lose, and below that price, B would prefer to win.

There is a range of bid prices from 5.5 to 9.5 wherein A prefers to lose and B prefers to win. Each bid price between these two values constitutes an equilibrium. If A is adamant to bid at least 8, it is in B's best interest to bid at least $8 + \epsilon$ and win. Similarly, if B is adamant not to bid above 6, it is in A's best interest to lose by bidding $6 - \epsilon$. Conditioned on B winning, the closer the winning bid is to 5.5, the higher is the B's utility and the lower is A's utility. Similarly, the closer the winning bid is to 9.5, the higher is A's utility and the lower is B's utility. Hence, there is a 'bargaining problem' in equilibrium selection here in the case of asymmetric valuations where the two agents bargain for the equilibrium bidding price.

In the English and Dutch auction there is no such bargaining problem. In the English auction the equilibrium is at the higher crossover point (9.5), because the lower-valuation bidder can safely bid up to that point because the higher-valuation bidder would not want to drop out before then. Analogously, in the Dutch auction the equilibrium is at the lower crossover point (5.5).

Incomplete information setting

We now discuss the more realistic setting where bidders have incomplete information about each others' valuations. Let bidders A and B have the joint distribution of their valuations given in Table 1. If A has valuation M, we say that A is of type M. In this example, we intentionally set the probability that both bidders have same type to zero because tied bids make the analysis more involved. Let the bidders have spite $\alpha_A^t = \alpha_B^t = 0.1$.

Table 1: *Joint probability distribution over valuations.*

A's type	B's type	Probability	A's type	B's type	Probability
50	100	1/6	100	50	1/6
50	200	1/6	100	200	1/6
200	50	1/6	200	100	1/6

Clearly A and B are symmetric so we can look for symmetric bidding strategies that will constitute an equilibrium.

Furthermore, we make the natural assumption that the equilibrium bids at type 50, 100, and 200 are in increasing order so a bidder of type 50 (if one exists) always loses and a bidder of type 200 (if one exists) always wins.

Let Q be the bid made by bidder of type 100 in equilibrium. Now a type-50 bidder can bid anywhere up to Q since she knows that the other bidder (regardless of whether he is of type 100 or 200) will bid at least Q. Similarly, a type-200 bidder would like to bid as close to (but higher than) Q as possible since that would maximize her utility. So, in equilibrium all types bid basically the same amount (but they prefer ties to be broken in favor of higher types).

Q clearly has to be such that a bidder of type 50 prefers to lose at that bid and a bidder of type 200 prefers to win at that bid. Let T be the crossover point of the type-50 bidder. Then, $50 - T = -0.1 \cdot (100 - T + 200 - T)/2$, where the left hand side is her utility if she wins at T and the right hand side is her expected utility if she loses. Solving this yields T=59.1. So, Q must be at least 59.1. Similarly, denoting by W the crossover point of the type-200 bidder, we must have $200 - W = -0.1 \cdot (50 - T + 100 - T)/2$. This yields $W = 188.6$, so Q must be at most 188.6. If Q is the bidding price, then the expected utility of a type-100 bidder is $((100 - Q) - 0.1 \cdot (200 - Q))/2$. (Here we have used the fact that a type-50 other bidder always loses and a type-200 other bidder always wins.) If Q is above 88.9, then the expected utility is negative. Hence, Q must be below 88.9. Q can therefore lie between 59.1 and 88.9. So again, like in the complete information setting, there is a range of bid values Q that constitute an equilibrium. Thus there is a bargaining problem in this setting as well.

Prior results

Prior work has provided equilibrium analysis for settings where bidders draw their valuations from the same distribution and have *equal* spite values α^t $(\forall X, \alpha^t_X = \alpha^t)$ (Morgan, Steiglitz, and Reis 2003; Brandt, Sandholm, and Shoham 2007). We now summarize some of those prior results in order to provide a comparison point to the results we will derive. Table 2 summarizes the symmetric equilibrium bidding strategies for the settings where the bidder's valuation are drawn uniformly from $[0, 1]$.

Table 2: *Bidding functions under symmetric spite.*

Auction	2-bidders	n-bidders
First-price sealed-bid and Dutch	$\left(\frac{1+\alpha^t}{2+\alpha^t}\right) v$	$\left(\frac{n-1}{n-\frac{\alpha^t}{1+\alpha^t}}\right) v$
Second-price sealed-bid	$\left(\frac{1+\alpha^t}{1+2\alpha^t}\right) v + \frac{\alpha^t}{1+2\alpha^t}$	$\left(\frac{1+\alpha^t}{1+2\alpha^t}\right) v + \frac{\alpha^t}{1+2\alpha^t}$
English	$\left(\frac{1+\alpha^t}{1+2\alpha^t}\right) v + \frac{\alpha^t}{1+2\alpha^t}$	*text explains strategy*

So, in the first-price sealed-bid 2-bidder case , agents bid higher under spite than under self-interest: $\frac{2v}{3}$ when $\alpha^t = 1$ and $\frac{v}{2}$ when $\alpha^t = 0$. This is also the case in the second-price 2-bidder auction: as α^t varies from 0 to 1, the bid of the spiteful bidder varies from v to $\frac{2}{3}v + \frac{1}{3}$. Incidentally, in the first-price auction, the bidding function of the 2-bidder case can be transformed to the n-bidder case by replacing $(1 + \alpha^t)$ by $(n - 1)(1 + \alpha^t)$. Furthermore, for the second-price sealed-bid auction, the bidding function is the same in 2-bidder and n-bidder settings.

In the English auction with n bidders, the bidding strategy differs from the 2-bidder case and is the following (Morgan, Steiglitz, and Reis 2003).

- If three or more bidders are present, each bidder drops out as the price reaches her valuation.
- If only two bidders remain, each bidder drops out when the price reaches $b(v)$, where $b(v)$ is the bid she would have submitted in a 2-bidder second-price sealed-bid auction.

Asymmetric spite results

We now move to the setting where bidders can be spiteful to different extents. Throughout the rest of the paper we will assume that the bidder's valuations are drawn uniformly and independently from $[0, 1]$. Further, as in prior research, we assume that for a given set of spite factors of the bidders, the equilibrium bidding function is strictly increasing in the agent's valuation. As in prior research, we will focus on studying symmetric equilibria, that is, equilibria where the form of the bidding function is the same for every agent. That, of course, does not mean that the agents' bids are the same because they have different valuations and different spite factors.

The 2-bidder case

We first analyze the setting with two bidders, A and B. Denote by $b_B(\cdot)$ the equilibrium bidding function of bidder B, that is, if B has valuation v_B, she bids $b_B(v_B)$ in equilibrium. Similarly, denote by b_A the bid of bidder A when she has valuation v_A.

First-price sealed-bid auction and Dutch auction In the first-price sealed-bid auction (and its strategic equivalent, the Dutch auction), the expected utility of bidder A is

$$\int_0^{b_B^{-1}(b_A)} [v_A - b_A]\mathrm{d}v_B - \alpha^t_A \int_{b_B^{-1}(b_A)}^1 [v_B - b_B(v_B)]\mathrm{d}v_B \quad (1)$$

The first term above is for the case where A wins and the second term is for the case where she loses. To solve this for A's bid b_A, our high-level approach is to differentiate the above equation with respect to b_A and solve for b_A by equating the differential to 0. We will now present the derivation in detail. We guess that the bidding function in symmetric equilibrium is a linear functions of the bidder's valuation (as it was in the symmetric-spite setting). Hence we can write $b_B(v_B) = r_B(\alpha^t_A, \alpha^t_B)\cdot v_B$ and differentiate (1) with respect to b_A. Equating the differential to 0, we get

$$b_A = \left(\frac{1}{2 + \alpha^t_A(1 - 1/r_B)}\right) v_A$$

We observe that this bidding function for A is of the linear form we guessed. This confirms the guess that these bidding functions constitute a symmetric equilibrium.

This can be written as $b_A = r_A \cdot v_A$, where

$$r_A = \left(\frac{1}{2+\alpha_A^t(1-1/r_B)}\right)$$

We have an exactly analogous equation for r_B. Here, r_A and r_B are written as functions of each other, while we would like to express them as functions of the spite coefficients only. Hence we solve the equations for r_A and r_B simultaneously to get

$$r_A = \frac{1-\alpha_A^t\alpha_B^t}{2-\alpha_A^t-\alpha_A^t\alpha_B^t}$$

This can be put in a better-looking form by introducing symbols α_A^e and α_B^e, so b_A becomes

$$b_A(v) = \left(\frac{1+\alpha_A^e}{2+\alpha_A^e}\right)v \tag{2}$$

where

$$\alpha_A^e = \frac{\alpha_A^t}{1+\alpha_B^e-\alpha_A^t} = \alpha_A^t\left(\frac{1-\alpha_B^t}{1-\alpha_A^t}\right) \tag{3}$$

We get analogous equations for α_B^e.

With the above transformation, we can observe that the bidding function in this asymmetric-spite settings looks like the bidding function in the symmetric-spite setting (Table 2)—except that there is now the symbol α_A^e in place of α^t. Because of their close connection, we call α^t *(t)rue* α and α^e *(e)xpressed* α, though no semantics behind these names are intended here.

Note that the first expression for α_A^e in Equation 3 is in terms of α_B^e and α_A^t, while the second expression is in terms of α_B^t and α_A^t. Depending on the situation, either of these forms can be useful.

Second-price sealed-bid auction In the second-price sealed-bid auction, the expected utility of bidder A is

$$\int_0^{b_B^{-1}(b_A)}[v_A - b_B(v_B)]\mathrm{d}v_B - \alpha_A^t\int_{b_B^{-1}(b_A)}^1[v_B-b_A]\mathrm{d}v_B \tag{4}$$

We guess that b_B is of the form $r_B(\alpha_A^t,\alpha_B^t)\cdot v_B + s_B(\alpha_A^t,\alpha_B^t)$ just as it was in the symmetric-spite second-price sealed-bid setting. Note that the guess here also includes an additive term $s_B(\alpha_A^t,\alpha_B^t)$ unlike in the first-price sealed-bid setting. Using this form of b_B, we differentiate (4) with respect to b_A. We then equate the differential to 0 to get

$$b_A = \left(\frac{r_B}{r_B-\alpha_A^t(1-2r_B)}\right)v_A + \frac{\alpha_A^t(r_B^2+(r_B-1)s_B)}{r_B-\alpha_A^t(1-2r_B)}$$

This bidding function is of the form we had guessed with

$$r_A = \frac{r_B}{r_B-\alpha_A^t(1-2r_B)},\quad s_A = \frac{\alpha_A^t(r_B^2+(r_B-1)s_B)}{r_B-\alpha_A^t(1-2r_B)}$$

We have analogous equations for bidder B. This proves that the guess is correct, that is, these functions indeed constitute a symmetric equilibrium.

Solving the above equations for r_A and s_A simultaneously with the analogous ones for r_B and s_B, we get

$$r_A = \frac{1-\alpha_A^t\alpha_B^t}{1+\alpha_A^t-2\alpha_A^t\alpha_B^t},\quad s_A = \frac{\alpha_A^t-\alpha_A^t\alpha_B^t}{1+\alpha_A^t-2\alpha_A^t\alpha_B^t}$$

These equations can be again put in a nice form by introducing symbols α_A^e and α_B^e. We get the equilibrium bidding function for A as

$$b_A(v) = \left(\frac{1+\alpha_A^e}{1+2\alpha_A^e}\right)v + \frac{\alpha_A^e}{1+2\alpha_A^e} \tag{5}$$

where

$$\alpha_A^e = \frac{\alpha_A^t}{1+\alpha_B^e-\alpha_A^t} = \alpha_A^t\left(\frac{1-\alpha_B^t}{1-\alpha_A^t}\right) \tag{6}$$

We get analogous equations for bidder B.

We observe that the bidding function is exactly of the form as in the symmetric-spite case—except that α^e has replaced α^t. Furthermore, the expression of α^e is the same as in the asymmetric-spite *first*-price sealed-bid auction.

English Auction In the English auction, when the clock price is z, the expected utility of bidder A as a function of her bid b_A is

$$\frac{1}{1-z}\Big(\int_{b_B^{-1}(z)}^{b_B^{-1}(b_A)}[v_A-b_B(v_B)]\mathrm{d}v_B-\alpha_A^t\int_{b_B^{-1}(b_A)}^1[v_B-b_A]\mathrm{d}v_B\Big) \tag{7}$$

The methodology to solve for the equilibrium bid function remains the same. We guess that b_B is of the form $r_B(\alpha_A^t,\alpha_B^t)\cdot v_B + s_B(\alpha_A^t,\alpha_B^t)$ just as it was in the symmetric-spite case, and it turns out that we get the same bidding function as in the second-price sealed-bid auction.

We summarize the results in Table 3. We observe that in all cases, the bidding functions in is of the same form as in the symmetric-spite setting (Table 2)—except with α^e occupying the place of α^t.

Table 3: *Equilibrium in the 2-bidder asymmetric-spite setting.*

Auction type	Bidding function for bidder A	Expression for α_A^e
First-price sealed-bid and Dutch	$\left(\frac{1+\alpha_A^e}{2+\alpha_A^e}\right)v$	$\frac{\alpha_A^t}{1+\alpha_B^e-\alpha_A^t}$
Second-price sealed-bid and English	$\left(\frac{1+\alpha_A^e}{1+2\alpha_A^e}\right)v+\frac{\alpha_A^e}{1+2\alpha_A^e}$	

Furthermore, these are the only linear (in case of first-price sealed-bid and Dutch auctions) and affine (in case of second-price sealed-bid and English auctions) equilibrium bidding functions. This is because guessing these forms for one bidder yielded unique bidding functions of the same form for the other bidder.

Comparison of α^t and α^e In this section we discuss α^e as compared to α^t. Although α^t always lies between 0 and 1 (by assumption), α^e is bounded below by 0 and is unbounded from above. Table 4 lists values of α^e_A for some combinations of α^t_A and α^t_B.

Table 4: *The values in the table are bidder A's α^e. The rows correspond to various values of bidder A's α^t and the columns correspond to values of bidder B's α^t.*

	0	**0.1**	**0.3**	**0.5**	**0.7**	**0.9**
0	0	0	0	0	0	0
0.1	0.11	0.1	0.08	0.06	0.03	0.01
0.3	0.43	0.39	0.3	0.21	0.13	0.04
0.5	1	0.9	0.7	0.5	0.3	0.1
0.7	2.33	2.1	1.63	1.17	0.7	0.23
0.9	9	8.1	6.3	4.5	2.7	0.9

From Equation 3 and Table 4 we see the following.

1. In the symmetric case, $\alpha^e = \alpha^t$ as we should expect from comparing the equations of the symmetric and asymmetric case.
2. In the asymmetric case,
   - For a given α^t_A, the value of α^e_A decreases linearly with increasing α^t_B.
   - For a given α^t_B, the value of α^e_A increases with increasing α^t_A.

   This implies that for a fixed α^t_A, bidder A might bid higher or lower in the asymmetric-spite case than in the symmetric-spite case depending on α^t_B. In fact, the mapping between true and expressed spite factors is such that if $\alpha^t_A \neq \alpha^t_B$, then the difference between α^e_A and α^e_B is greater than the difference between α^t_A and α^t_B. So, in equilibrium, the more spiteful bidder expresses an over-exaggerated spite and the less spiteful bidder expresses an under-exaggerated spite.

Revenue (Brandt, Sandholm, and Shoham 2007) show that in symmetric-spite settings, second-price auctions yield higher expected revenue than first-price auctions. In contrast, we show that in asymmetric-spite settings there is no revenue-dominant auction. For a given set of α^e's, we give in Table 5 the expected revenue for a first-price sealed-bid/Dutch auction and a second-price sealed-bid/English auction.

Table 5: *Expected revenue for given α^e's in 2-bidder setting.*

Auction type	Expected revenue	Notation
First-price sealed-bid and Dutch	$\frac{1}{3}(\frac{p_A^2}{p_B} + \frac{p_B^2}{p_A})$	$p_X = \frac{1+\alpha^e_X}{2+\alpha^e_X}$
Second-price sealed-bid and English	$\frac{1}{2}(\frac{q_B}{q_A} + \frac{q_A}{q_B}) - \frac{1}{3}(\frac{q_B^2}{q_A} + \frac{q_A^2}{q_B})$	$q_X = \frac{1+\alpha^e_X}{1+2\alpha^e_X}$

For $\alpha^e_A = 8.1$ and $\alpha^e_B = 0.01$, corresponding to $\alpha^t_A = 0.9$ and $\alpha^t_B = 0.1$, the first-price auction yields expected revenue 0.63 while the second-price auction yields expected revenue 0.49. For $\alpha^e_A = 0.39$ and $\alpha^e_B = 0.08$, corresponding to $\alpha^t_A = 0.3$ and $\alpha^t_B = 0.1$, the first-price auction yields expected revenue 0.37 while the second-price auction yields expected revenue 0.42. So here, neither auction mechanism beats the other in expected revenue in general. By substituting various values of spite into the formulas of Table 5, we found that when bidders have comparable true spite, the second-price auction yields higher expected revenue while if the true spite values differ largely, the first-price auction yields higher revenue, as can be seen in Figure 1.

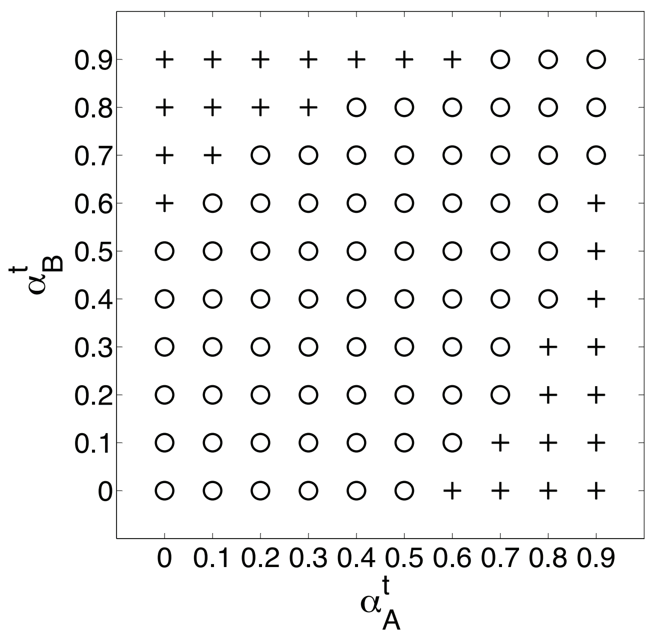

Figure 1: *The circles denote the points where the second-price auction yields higher expected revenue than the first-price auction. The pluses denote the points where the reverse occurs.*

Allocation In symmetric-spite settings, the bidder with the highest valuation wins the item. In asymmetric-spite settings, this is no longer true in general. For example, in a first-price sealed-bid auction, consider the case where bidder A has valuation 0.9 and $\alpha^t_A = 0.1$ while bidder B has valuation 0.7 and $\alpha^t_B = 0.7$. For this pair of α^t's, we get $\alpha^e_A = 0.03$ and $\alpha^e_B = 2.1$, so A bids 0.46 and B bids 0.53. Hence, the lower-valuation bidder (B) wins.

Figures 2 and 3 plot a bidder's bid against her α^e for various valuations she may have in first-price and second-price auctions. For low valuations, $v = 0.1$ or 0.2, the curve is nearly flat in the first-price auction. This means that as α^e increases, her bid does not increase much. However, in the second-price auction, for the same set of low valuations, the bid increases steeply as α^e increases from 0 to 2, and keeps increasing after that. In contrast, consider high valuations. For $v = 0.9$ or 0.8, in the first-price auction, the bid increases rapidly as α^e increases from 0 to 2, and keeps increasing steadily after that. In the second-price auction, however, for valuations 0.9 and 0.8, there is not much change in the bid as α^e increases from 0 to 10.

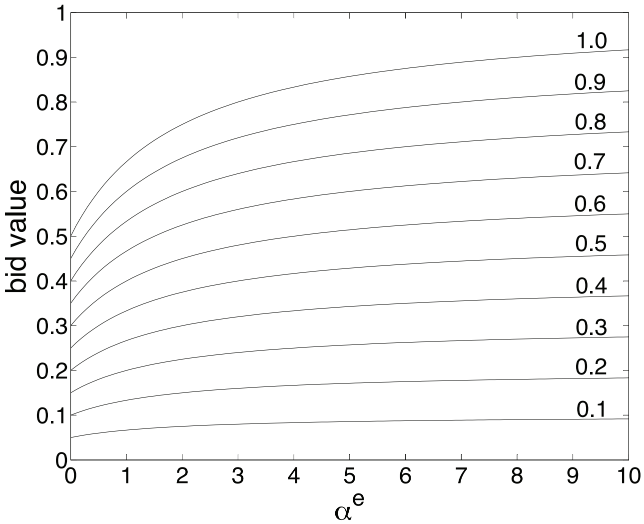

Figure 2: *The bid in a first-price auction varies with the expressed spite α^e. The curves are for different valuations v, with the lowest curve corresponding to $v = 0.1$ and the highest to $v = 1$.*

These observation imply that if A has a high valuation in the second-price auction, it is unlikely that B can win if B has a slightly lower valuation, say 0.8, no matter how high α^e_B is. In the first-price auction however, if A has a high valuation, say 0.9, but low α^e_A (< 1), then bidder B can win even with valuation 0.7 but with a high α^e_B, say 5. Similar statements can be made for the low valuation case, but with the first-price and second-price auction switching roles.

The n-bidder setting with directed spite

We now extend our analysis of asymmetric-spite auctions to n bidders. With more than two bidders, there is the possibility that some bidder(s) have different extents of spite toward different other bidders. We call this *directed* spite. For example, bidder A can have spite factor α^t_{AB} toward bidder B and spite factor α^t_{AC} toward bidder C, so her utility would be

$$u_A = \begin{cases} v_A - p_A & \text{if } A \text{ wins} \\ -\alpha^t_{AB} \cdot (v_B - p_B) & \text{if } B \text{ wins} \\ -\alpha^t_{AC} \cdot (v_C - p_C) & \text{if } C \text{ wins} \end{cases}$$

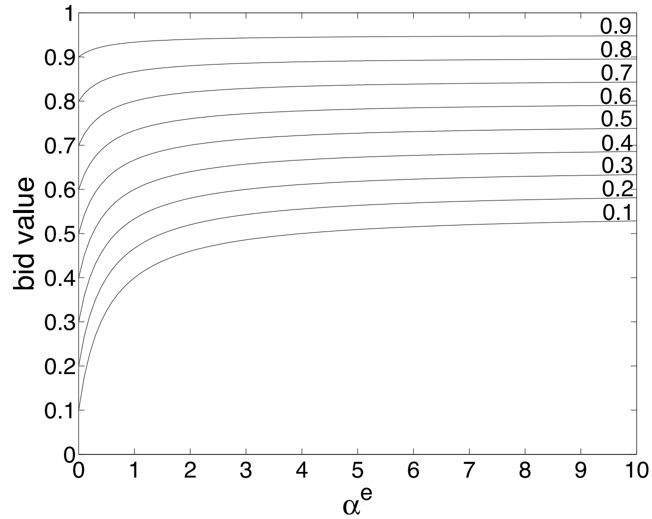

Figure 3: *The bid in a second-price auction varies with the expressed spite α^e. The curves are for different valuations v, with the lowest curve corresponding to $v = 0.1$ and the highest to $v = 0.9$. For $v = 1$, the bidder always bids 1.*

First-price sealed-bid and Dutch auction The expected utility of A is

$$\begin{aligned} &\int_0^{b_C^{-1}(b_A)} \int_0^{b_B^{-1}(b_A)} [v_A - b_A] \mathrm{d}v_B \mathrm{d}v_C \\ &- \alpha^t_{AB} \int_{b_B^{-1}(b_A)}^{1} \int_0^{b_C^{-1}(b_B(v_B))} [v_B - b_B(v_B)] \mathrm{d}v_C \mathrm{d}v_B \\ &- \alpha^t_{AC} \int_{b_C^{-1}(b_A)}^{1} \int_0^{b_B^{-1}(b_C(v_C))} [v_C - b_C(v_C)] \mathrm{d}v_B \mathrm{d}v_C \end{aligned} \tag{8}$$

where b_B is the bidding function of B and b_C is the bidding function of C. The same formula extends to any number of agents in the obvious way. We again guess that b_A, b_B and b_C are linear functions of the valuation, and we indeed find an equilibrium, verifying the guess. The results for the n-agent case are stated in the last row of Table 6. In the row before that, we also show how these results specialize to the case of undirected spite, that is, where each agent X has a spite factor α_X toward each other bidder.

English auction The analysis of the English auction is a bit more intricate. We analyze directed spite; the undirected-spite setting is a special case.

Proposition 1 *In the English auction, an equilibrium strategy for any bidder A is to stay in if the clock price is lower than $\max\{v_A, \max_{X \in S} B(A, X)\}$ and to drop out otherwise. Here S is the set of other bidders who are still in, and $B(A, X)$ is the bidding function for A if X were the only other bidder (and we know this form from our analysis of the two-bidder case, Table 3).*

Proof: Any n-bidder auction would in the end reduce to a 2-bidder setting. Until what clock price should A stay in?

Table 6: *Bidding function for n-bidder asymmetric-spite first-price sealed-bid and Dutch auctions.*

Case	Bidding function for bidder A	Expression for $\frac{\alpha_A^e}{1+\alpha_A^e}$
Undirected spite	$\left(\frac{n-1}{n-\frac{\alpha_A^e}{1+\alpha_A^e}}\right)v$	$\frac{\alpha_A^t}{n-1}\left(\sum_{X:X\neq A}\frac{1}{1+\alpha_X^e}\right)$
Directed spite	$\left(\frac{n-1}{n-\frac{\alpha_A^e}{1+\alpha_A^e}}\right)v$	$\frac{1}{n-1}\left(\sum_{X:X\neq A}\frac{\alpha_{AX}^t}{1+\alpha_X^e}\right)$

As long as there is at least one bidder, Z, still in to whom A would bid higher in a 2-bidder setting than the current clock price, A should stay in. This is because if A leaves before the clock reaches that price and all other bidders except Z also exit before the clock reaches that price, then Z will win at a price lower than if A had stayed in. So, A would end up with a lower utility due to leaving early.

What if there is no such person Z still in? If the clock price has exceeded A's valuation v_A, then A no longer wants to win, so it is best for her to exit. If, on the other hand, the clock price has not yet reached v_A, then A should stay in until she wins or the price exceeds v_A. □

Conclusions and future research

We game-theoretically analyzed the four common auction mechanisms when bidders have asymmetric spite. A noteworthy feature is that the symmetric equilibrium bidding function continues to be the same as with symmetric spite—except that the true spite is replaced by 'expressed' spite. Unlike in the symmetric-spite setting, bidders express spites that are higher or lower than their true spite depending on others' spites. Moreover, the equation for expressed spite does not depend on the auction mechanism in the 2-bidder case. Furthermore, we found that the allocation can be inefficient and that the revenue ranking may reverse between first- and second-price auctions. We also studied the generalization in the n-bidders setting where agents can have different extents of spite toward different other bidders. We also showed that in sealed-bid auctions under asymmetric valuation distributions, there can be a "bargaining problem" in selecting bids.

Future work includes solving for the equilibrium of the second-price auction in the n-bidder case. We also plan to study valuation priors that are not uniform.

We assumed the bidders know each others' true spite. In settings where they do not know, we have to understand whether the equilibrium will be reached, and how. We conducted experiments that indicate that in a repeated-game setting the equilibrium can be learned as long as the bidders are able to infer each others' *expressed* spites. In our simulation, each bidder adjusts her own expressed spite given the others' expressed spites (using the expression for α^e in Table 3 or 8). The bidders rapidly converged to the equilibrium values of α^e's regardless of the initial values of the α^e's. Future work includes proving bounds on this convergence. Another interesting direction would be to solve for the equilibrium in settings where bidders do not know each others' true spite coefficients but have a joint prior over them.

References

Brainov, S. 2000. The role and the impact of preferences on multiagent interaction. In Jennings, N., and Lespérance, Y., eds., *Intelligent Agents VI, LNAI 1757*. Springer-Verlag. 349–363.

Brandt, F., and Weiß, G. 2001. Antisocial agents and Vickrey auctions. In Meyer, J.-J. C., and Tambe, M., eds., *Intelligent Agents VIII, LNAI 2333*.

Brandt, F.; Sandholm, T.; and Shoham, Y. 2007. Spiteful bidding in sealed-bid auctions. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*. Early version in GTDT-05.

Grimm, V.; Riedel, F.; and Wolfstetter, E. 2001. The third generation (UMTS) spectrum auction in Germany. CESifo Working Paper Series CESifo Working Paper No. 584, CESifo Group Munich.

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

Loewenstein, George F.; hompson, Leigh T.; and Bazerman, Max H. 1989. Social utility and decision making in interpersonal contexts. *Journal of Personality and Social Psychology*, 57: 426–441.

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

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

Vetsikas, I., and Jennings, N. 2007. Outperforming the competition in multi-unit sealed bid auctions. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8–37.

Zhou, Y., and Lukose, R. 2006. Vindictive bidding in keyword auctions. In *Workshop on Sponsored Search Auctions*.