

Antisocial Agents and Vickrey Auctions

Felix Brandt and Gerhard Weiß

Institut für Informatik, Technische Universität München
80290 München, Germany
Tel. +49-89-28928416 / 28922606
{brandtf, weissg}@in.tum.de

Abstract. In recent years auctions have become more and more important in the field of multiagent systems as useful mechanisms for resource allocation and task assignment. In many cases the Vickrey (second-price sealed-bid) auction is used as a protocol that prescribes how the individual agents have to interact in order to come to an agreement. We show that the Vickrey auction, despite its theoretical benefits, is inappropriate if “antisocial” agents participate in the auction process. More specifically, an antisocial attitude for economic agents that makes *reducing the profit of competitors* their main goal besides maximizing their own profit is introduced. Under this novel condition, agents need to deviate from the dominant truth-telling strategy. This paper presents a strategy for bidders in repeated Vickrey auctions who are intending to inflict losses to fellow agents in order to be more successful, not in absolute measures, but relatively to the group of bidders. The strategy is evaluated in a simple task allocation scenario.

1 Introduction

The area of multiagent systems [20], which is concerned with systems composed of technical entities called agents that interact and in some sense can be said to be intelligent and autonomous, has achieved steadily growing interest in the past decade. Two key problems to be addressed in this area are automated resource allocation and task assignment among the individual agents. As a solution to these problems it has become common practice to apply well known results and insights from auction theory (e.g., [9, 10]) and well understood auction protocols like the English auction, the Dutch auction, and the Vickrey auction. Among the different protocols, the Vickrey auction [18] (also known as second-price sealed-bid auction) has received particular attention within the multiagent community and has been applied in a variety of contexts like e-commerce, operating systems, and computer networks (e.g., [17, 6, 5, 19, 8, 4]). The Vickrey auction is favored because of three main characteristics:

- it requires low bandwidth and time consumption
- it possesses a dominant strategy, namely, to bid one’s true valuation (see Section 2)
- it is a sealed-bid auction (bids (expressing private values) remain secret)

These characteristics make the Vickrey auction protocol particularly appealing from the point of view of automation. The Vickrey auction, in its original formulation and as it is used for *selling goods or resource allocation*, works as follows: each bidder makes a

sealed bid expressing the amount he is willing to pay, and the bidder who submits the highest bid wins the auction; the price to be paid by the winner is equal to the second highest bid. In *task assignment scenarios* the Vickrey auction works exactly the other way round (and for that reason is often referred to as reverse Vickrey auction): each bidder willing to execute a task makes a bid expressing the amount he wants to be paid for task execution, and the bidder submitting the lowest bid wins the auction; the winner receives an amount equaling the second lowest bid (and his payoff thus is the second lowest bid minus his prime costs for execution). If there is more than one winning bid, the winner is picked randomly from the highest bidders. This paper concentrates on the use of the Vickrey auction for task assignment scenarios; it should be noted, however, that all presented considerations and results do also hold for Vickrey auctioning in its original formulation. It is an important fact, that bids are sealed. Bidders that did not win an auction do not necessarily get to know who placed the lowest bid and particularly how low this bid was. The only information that is revealed to all bidders is the selling price (see [2] for a more detailed discussion of privacy and security issues in Vickrey auctions).

A general assumption often made in applications of multiagent systems is that an individual agent intends to maximize his profit without caring for the profits of others. This, however, is by no means the case in real-world settings. What can be often observed here is that a company, besides maximizing its own profit, deliberately inflicts losses to rivaling companies by minimizing their profits. In fact, it is real-world practice that a company accepts a lower profit or is even willing to sell goods at a loss if this financially damages a competing company or at least helps to bind available and gain new costumers. Such a company obviously exhibits an *antisocial* behavior and attitude in the sense that it considers (at least for some period of time or in some cases) its absolute profit not as important as its profit relative to other companies' profits. (It is worth to point out that an "antisocial company" behaves rationally given its objective, even if it could be said to act irrationally under the condition that its objective were to maximize its absolute profit.) This paper investigates the performance of Vickrey auctioning in multiagent systems in which the presence of antisocial agents can not be excluded. A related problem that has been discussed in the economic literature are "externalities" between bidders [7]. Externalities describe preferences of a bidder specifying who he wants to win an auction (besides himself), e.g., when selling patents or nuclear weapons. However, our approach is solely profit-oriented. Antisocial agents do not care who buys a good or who is awarded a task contract. They are just interested in decreasing their fellows' profits to strengthen their own market position. This is a major difference.

The paper is structured as follows. Section 2 explains why agents aiming at a maximization of their absolute profits should bid their true valuations in Vickrey auctions. Section 3 formally captures the antisocial attitude sketched above and Section 4 shows its impact on the bidding strategy. Section 5 introduces and analyzes an antisocial strategy for repeated Vickrey auctions. Section 6 presents experimental results that illustrate the implications of this strategy. Finally, Section 7 concludes the paper with a brief overview of advantages and disadvantages of the agents' "bad attitude".

2 The dominant Strategy

The Vickrey auction has a dominant strategy, which means that if an agent applies this strategy he receives the highest possible payoff, no matter which strategies are used by the other bidders. The dominant strategy is to bid one's true valuation of the task. Even if an agent knows all the other bids in advance, he still does best by bidding his private valuation. The conditions for the existence of the dominant strategy equilibrium are that the bidders are symmetric (i.e., they can not be distinguished) and have independent valuations of the tasks¹. This implies that tasks cannot be recontracted.

Given two agents A and B , their corresponding private values v_a and v_b , and their submitted bids b_a and b_b , the profit for agent A is defined by the following equation.

$$profit_a(b_a, b_b) = \begin{cases} b_b - v_a & \text{if } b_a \leq b_b \\ 0 & \text{if } b_a \geq b_b \end{cases} \quad (1)$$

It can easily be seen why bidding the task value is the optimal strategy. We consider agent A and investigate the possible profits he would make by not bidding his private value. It suffices to model only one opposing agent B representing the entire competition because A only cares whether he wins or loses. He does not draw distinctions between his fellow bidders.

If A bids more than his prime costs ($b_a > v_a$) there are three subcases conditional on agent B 's bid b_b :

- i) $b_b < v_a < b_a$: B wins the auction and receives more money than if A had bid v_a .
- ii) $v_a < b_b < b_a$: A loses the auction, instead of winning it by bidding v_a .
- iii) $v_a < b_a < b_b$: A still wins the auction, but does not gain anything, because the task price remains b_b and his payoff is still $b_b - v_a$.

If he bids $b_a < v_a$ the following cases describe the resulting situations:

- iv) $b_a < v_a < b_b$: A wins, but is paid the same amount of money (b_b) as if he bid v_a .
- v) $b_a < b_b < v_a$: A wins, but gets less money than his own prime costs, i.e. he is losing $v_a - b_b$.
- vi) $b_b < b_a < v_a$: A still loses and reduces B 's payoff by $v_a - b_a$.

Concluding, bidding anything else than v_a cannot yield more profit than bidding the true valuation v_a . Obviously, this extremely simplifies the bid preparation, due to the absence of wasteful counter-speculation, that is needed e.g., in first-price or Dutch auctions.

3 The Antisocial Attitude

In most multiagent applications it is assumed that the objective of an agent (or of a team of agents) is to maximize his absolute profit without caring for the profits made by the

¹ If the bidders don't have to estimate their private values, the dominant equilibrium exists independently of risk neutrality [11].

other agents. However, in many real-world applications it is more realistic to assume that agents may be present that try to gain as much money as possible *relative* to others (their competitors). In other words, in many scenarios it is wise to take into consideration the availability of “antisocial agents,” that is, agents who accept small losses if they can inflict great losses to other agents. To make this more precise, we develop a formal description of this antisocial attitude. As a starting point for this formalization, it appears to be reasonable to assume that an antisocial agent wants to maximize the difference between his profit and the gain of his competitors; this means that the own profit on the one hand and the other agents’ losses on the other hand are considered to be of equal importance from the point of view of this antisocial agent. In a two-player scenario, this view captures the antisocial agent’s intention to be better than his rival. To achieve a higher degree of flexibility in describing and analyzing antisocial agents, it is useful to think of different degrees of anti-sociality like “aggressive anti-sociality” (where it is an agent’s objective to harm competitors at any cost) and “moderate anti-sociality” (where an agent puts somewhat more emphasis on his own profit rather than the loss of other agents). These considerations lead to our formal specification of an antisocial agent (or an agent’s antisocial attitude) as an agent who tries to maximize the weighted difference of his own profit and the profit of his competitors. Precisely, an antisocial agent i intends to maximize his *payoff*² that is given by

$$payoff_i = (1 - d_i)profit_i - d_i \sum_{j \neq i} profit_j \quad , \quad (2)$$

where $d_i \in [0, 1]$ is a parameter called *derogation rate*. The derogation rate is crucial because it formally captures, and allows to modify, an agent’s degree of antisocial behavior. It is obvious that this formula covers “regular” agents by setting $d = 0$. If d is higher than 0.5, hurting others has greater priority than helping yourself. A purely destructive agent is defined by $d = 1$. We say an agent is *balanced antisocial* if $d = 0.5$, e.g., his own profit and the profit of his competitors are of equal importance.³ Please note that the described antisocial attitude, which is feasible in all kinds of competitive markets, is rational as the agents intend to benefit from their rivals’ losses at a later date (e.g., by shortening a competitor’s budget for future negotiations or by putting a rival completely out of business).

4 Antisociality and Vickrey Auctions

The implications of this formalized notion of an antisocial attitude on the Vickrey auction is enormous. In the remaining of this section, these implications are theoretically investigated. By combining equations 1 and 2 we get the (antisocial) payoff for agent

² The payoff for a non-antisocial agent is simply his profit.

³ When considering other types of auctions like uniform-price auctions, it might be wise to scale the sum in equation 2. In a Vickrey auction, however, all profits but one are zero which makes the scaling redundant.

A as a function of the bids b_a and b_b .

$$payoff_a(b_a, b_b) = \begin{cases} (1 - d_a)(b_b - v_a) & \text{if } b_a \leq b_b \\ -d_a(b_a - v_b) & \text{if } b_a \geq b_b \end{cases} \quad (3)$$

Consider case vi) of Section 2. Agent A is not able to effectively win the auction, but the price agent B receives completely depends on A 's bid. So, if A carefully adjusts his bid downwards, he is capable of reducing B 's profit. Supposing that A knows B 's private value v_b , his optimal strategy would be to bid $v_b + \epsilon$ (see Figure 1), which reduces B 's profit to the absolute minimum of ϵ .




Fig. 1. A reduces B 's profit to a minimum

However, if B knows A 's prime costs v_a as well and he also prefers inflicting great losses to A over gaining small profits (i.e., his derogation rate is positive) he can bid $v_b + 2\epsilon$, rejecting a possible gain of ϵ and making A lose $v_a - v_b + 2\epsilon$. As a result, if A 's derogation rate d_a is 0.5, A should only bid $v_b + \frac{v_a - v_b}{2} + \epsilon$ to be safe from being overbid by B (see Figure 2). If B still tops A 's bid, he is renouncing more money than A loses.




Fig. 2. Careful, aggressive bidding

If $d_b = 0.5$ as well, B 's best strategy is to bid $v_b + \frac{v_a - v_b}{2}$.

Concluding, the following bidding strategy seems to be "safe" for an antisocial agent i (still under the unrealistic assumption that an agent knows private values of other agents: v_1 is the lowest private value, v_2 the second lowest):

$$b_i = \begin{cases} v_i - d_i(v_i - v_1) + \epsilon & \text{if } v_i > v_1 \\ v_i + d_i(v_2 - v_i) & \text{else} \end{cases} \quad (4)$$

It is possible to omit the margin ϵ . We included it to avoid randomized experimental results that occur when two or more bidders share the same minimum bid. However, we will use the strict ϵ -less version in the theorems stated below.


Theorem: In Vickrey auctions with balanced antisocial bidders ($\forall i : d_i = 0.5$), the strategy defined by equation 4 is in *Nash equilibrium*.

Proof: The assumption states that under the supposition that all agents apply this strategy, there is no reason for a single agent to deviate from it.

We consider the payoff of agent A . If v_a is the lowest private value, let B be the second lowest bidder. Otherwise, B represents the lowest bidder. It suffices to consider B as A cannot differentiate between the individual bidders and the Vickrey auction has a sole victor. There are only two possibilities for an antisocial agent to cause harm. The cheapest provider can hurt the second highest bidder and an inferior agent can reduce the profit of the lowest bidder. Assume B applies the antisocial strategy defined in 4.

$$b_b = \begin{cases} v_b - \frac{1}{2}(v_b - v_a) & \text{if } v_a \leq v_b \\ v_b + \frac{1}{2}(v_a - v_b) & \text{if } v_a \geq v_b \end{cases} = \frac{v_a + v_b}{2}$$

$$\Rightarrow \text{payoff}_a(b_a, b_b) = \text{payoff}_a\left(b_a, \frac{v_a + v_b}{2}\right) = \begin{cases} \frac{v_b - v_a}{4} & \text{if } b_a \leq \frac{v_a + v_b}{2} \\ -\frac{b_a - v_b}{2} & \text{if } b_a \geq \frac{v_a + v_b}{2} \end{cases}$$



$$\max_{b_a} \text{payoff}_a\left(b_a, \frac{v_a + v_b}{2}\right) = \frac{v_b - v_a}{4} \Rightarrow b_a \leq \frac{v_a + v_b}{2}$$

Concluding, if A bids less than $\frac{v_a + v_b}{2}$, he only receives equal payoff; if he bids more, his payoff is diminishing. \square


Theorem: The strategy defined by equation 4 is in *Maximin equilibrium* independently of the derogation rates of the bidders.

Proof: It is claimed that the strategy is an optimal strategy to reduce the possible losses that occur in worst case encounters. “Worst-case” means that the other bidders (represented by a single agent B again) try to reduce agent A 's payoff as much as possible.

$$\begin{aligned} \min_{b_b} \text{payoff}_a(b_a, b_b) &= \min_{b_b} \begin{cases} (1 - d_a)(b_b - v_a) & \text{if } b_b \geq b_a \\ -d_a(b_a - v_b) & \text{if } b_b \leq b_a \end{cases} \\ &= \min \left\{ \underbrace{(1 - d_a)(b_a - v_a)}_{f(b_a)}, \underbrace{-d_a(b_a - v_b)}_{g(b_a)} \right\} \end{aligned}$$

f yields the minimum profit if A wins and g yields the minimum profit if he loses the auction. In the following, we consider the maximum of these minima.

$$\max_{b_a} \min_{b_b} \text{payoff}_a(b_a, b_b) = \max_{b_a} \min \{ f(b_a), g(b_a) \}$$



Due to the fact that f is increasing and g is decreasing, the Maximin equilibrium point can be computed by setting $f(b_a) = g(b_a)$.

$$\begin{aligned} f(b_a) = g(b_a) &\Leftrightarrow (1 - d_a)(b_a - v_a) = -d_a(b_a - v_b) \\ &\Leftrightarrow b_a - v_a - d_a b_a + d_a v_a = -d_a b_a + d_a v_b \\ &\Leftrightarrow b_a = v_a + d_a(v_b - v_a) \quad \square \end{aligned}$$

5 Antisocial Bidding in Repeated Auctions

On the basis of the theoretical foundations of the previous Section, we now develop an antisocial bidding strategy that can actually be used in realistic environments.

In the general case, an agent does not know the private value of other bidders. However, in principle an agent has several possibilities to figure out that value, for instance by means of *espionage*, careful *estimation* (e.g., by assuming a uniform distribution of private values), *colluding* with the auctioneer (e.g., bribing him) or by *learning* from previous auctions. This paper deals with the latter technique.

5.1 Revealing Private Values by Underbidding

We consider the auctioning of a fixed number of tasks, that repeats for several rounds. Now, suppose a balanced antisocial agent loses an auction in the first round. When the same task is auctioned for the second time, he bids zero. As a consequence, he wins the auction⁴, and receives an amount equaling the second lowest bid, which is the private value of the cheapest agent (supposing this agent applied the dominant strategy). Thus, he is able to figure out the needed private value and can place his next bid right in the middle between the two private values. Using this technique, he loses the difference between both values once, but can safely cut off 50% of the competitor's profit for all following auction rounds. If the total number of rounds is high enough, the investment pays.

In a scenario where all other agents definitely follow the dominant bidding strategy and no counter-speculation is needed, an effective, antisocial bidding strategy looks like this:

1. Bid 0 (p =received price)
2. Bid
$$\begin{cases} v_i + d_i(p - v_i) & \text{if } v_i < p \\ v_i - d_i(v_i - p) + \epsilon & \text{else} \end{cases}$$

Bidding zero is elegant but dangerous, especially if more than one agent is applying this strategy. In this case, one of the zero-bidding agents wins the auction, but is paid no money at all (because the second lowest bid is zero as well), thus producing a huge deficit.

5.2 Step by Step Approach

It's safer, but possibly not as efficient, to reduce a bid from round to round by a small margin s until the best agent's bid is reached. Figure 3 displays the modified strategy. If the step size s equals the private value ($s = v$), this algorithm emulates the aggressive zero-bidding strategy. The algorithm works somewhat stable in dynamic environments where agents can vanish and new ones appear from time to time. However, the strategy is not perfect, e.g., if two balanced antisocial agents apply this strategy, the more expensive agent is only able to reduce the winning agent's profit by 25% because he is usually

⁴ unless some other agent bids zero as well.

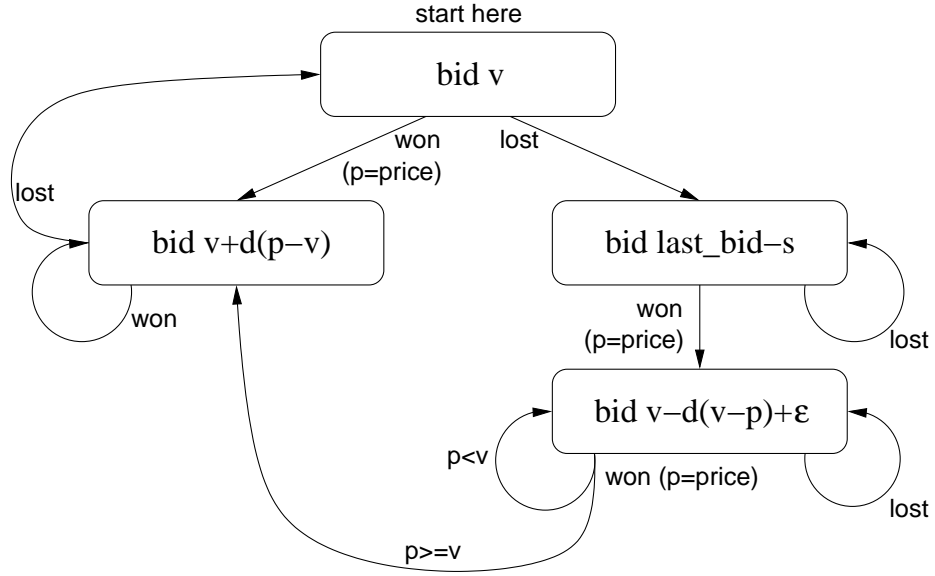


Fig. 3. Antisocial strategy for repeated Vickrey auctions

not able to figure out the real private value of the cheaper agent in time. Nevertheless, this is still tolerable, because in a group of antisocial agents a bidder never takes *more* risk than his derogation rate prescribes.

Generally, a careful agent should use a small step size s in order to be safe that the competitor already suffered huge losses before he makes negative profit himself. A reasonable setting of s depends on the number of rounds, the distribution of private values and his derogation rate (an upper bound for s is specified in [1]).

5.3 Leveled Commitment Contracting

If the task execution contracts are not binding and can be breached by paying a penalty (leveled commitment contracting [15, 16, 3]), the unavoidable loss an agent produces by underbidding the cheapest competitor can be reduced by breaking the negative contract. Due to the fact that the only reason for closing that deal is to figure out the private value of another agent, the agent has no incentive to really accomplish the task. Therefore, a contractee will break the contract if the loss he makes by accepting the contract is greater than the penalty he pays by breaking the deal. Supposing the common definition of a penalty as a fraction of the contract value, agent i is better off breaching the contract if

$$p \leq \frac{v_i}{pr + 1} \quad (5)$$

with p being the actual task price and $pr \in [0; 1]$ the penalty rate. To give an example, under the assumption that $pr = 0.25$, an agent should break a contract if the task price

is less or equal than $\frac{4}{5}$ of his private value. When the distribution of prime costs is uniform, this is true in 80% of all possible cases.

6 Experimental Results

The experimental setting investigated in this paper is similar to the one used in [3]. Due to reasons of limited space, we can only present some of the experiments we conducted. Please see [1] for more detailed results including a randomized cost table.

There is a number of buyers or *contractees* (CE_i) who are willing to execute tasks. Contractees associate prime costs with task execution and are interested in tasks whose prices are higher than their own costs. All prices and bids are integer values ($\epsilon = 1$). Whenever the selling of a task is announced, each interested contractee calculates and submits one sealed bid. The contractee who submitted the lowest bid is declared as the winner of the auction, and the *second lowest* bid is taken as the price of the announced task; the contractee is paid this price and executes the task. If there are two or more equal winning bids, the winner is picked randomly. As a contractee wants to earn money for handling tasks, his private value of a task is his prime costs plus ϵ . It is assumed that each contractee can execute as many tasks as he wants during one round (full commitment contracting).

Table 1 contains the prime costs of three contractees with exactly identical abilities. Each contractee has one task, that he can handle for the cheapest price. If all three truly bid their private values for 100 rounds, each one would gain $((50 - 30) + 1) \times 100 = 2100$. Figure 4 shows the profits accumulated by the contractees in 100 rounds. CE_1

	Task 1	Task 2	Task 3
CE_1	70	50	30
CE_2	50	30	70
CE_3	30	70	50

Table 1. Fair cost table

and CE_2 apply the dominant strategy and bid their prime costs plus one. CE_3 , however, is antisocial and tries to harm his competitors by reducing their profits to a minimum. As CE_3 is the only antisocial agent and because his derogation rate is 1, he could use a very large step size, e.g., $s_3 = v_3$. We chose a careful step size setting ($s_3 = \epsilon$) for two reasons. First of all, CE_3 may not know he is the only antisocial bidder, and secondly, this setting superiorly visualizes how the antisocial strategy works.

Despite the dominant strategy, CE_1 and CE_2 are incapable of gaining their regular profit (2100). CE_3 outperforms his rivals by losing only 60. The summed up profit of the entire group of contractees is reduced by more than 50% by a single agent using an aggressive strategy. This emphasizes the particular vulnerability of Vickrey auctions to “irrational” bidding. There is no counter-strategy available for CE_1 or CE_2 . As they apply a dominant strategy, they already make the highest possible payoff (which is almost nil) according to their definition of payoff. However, if they decided to become antisocial as well, they would be able to increase their payoff. In certain scenarios, this




Fig. 4. Identical contractees (1/3 antisocial)


might lead to an antisocial chain reaction.

It might appear confusing at the first glance that an agent who does not care for his own profit at all ($d_3 = 1$) nevertheless makes the highest profit. This odd effect can be explained by the peaceableness of the fellow bidders and the ϵ that we added in equation 4. CE_3 risks his entire profit in order to hurt CE_1 and CE_2 , but as both are completely harmless, he keeps his gain.

If all three contractees are antisocial, overall performance breaks down as expected (see Figure 5). The agents cut off more than a quarter of profits of their rivals due to a mutual derogation rate of 0.5.

7 Conclusions

The antisocial attitude for agents and its formalization introduced in this paper leads to a significant need for important changes in strategic behavior of agents. As argued above, and as it is obviously implied by many real-world applications, it is necessary to take the existence of antisocial agents into consideration. The paper focused on the Vickrey auction and showed that this auction protocol, in addition to other known deficiencies [14, 13, 12], is vulnerable to antisocial bidders. This is an effect of the second-price policy which enables easy price manipulation and emphasizes the required confidentiality of



All contractees are antisocial ($d_i = 0.5, s_i = \epsilon$).

Fig. 5. Identical contractees (3/3 antisocial)

private values and the need for a secure auction protocol that hides those values from all parties, including the auctioneer [2]. As the common *English* auction for private value bidders is equivalent to the Vickrey auction [9, 11], all strategies in this paper work for English auctions as well. The inability to prevent profit reduction can be regarded as a major disadvantage of those two auction types as *Dutch* and *first-price sealed-bid* auctions do not suffer from antisocial strategies.

One problem that arises using the new strategy in repeated Vickrey auctions is that if there is more than one inferior, antisocial agent, it would be desirable if only the one that intends to cut off the cheapest agent's profit by the highest margin reduces his bid. All other bidders should stay with bidding their private value, since they would lose money once without harming anyone in the following rounds. This would require the antisocial agents to make a special arrangement, which may not be feasible in certain settings. The behavior described in this paper can be seen as an opposite of *bidder collusion* where bidders coordinate their bids in order to help each other (harming the auctioneer). In contrast, antisocial agents bid with the intention to harm fellow bidders.

As a next research step we want to explore "anti-sociality" and strategies for antisocial agents in more detail. To do so, we plan to examine these strategies in more complex environments where the number of contractees varies over time and tasks can be recontracted. Furthermore, we intend to investigate types of collaboration of antisocial agents that allow them to act more efficiently.

References

1. F. Brandt. Antisocial Bidding in Repeated Vickrey Auctions. Technical Report FKI-241-00, Institut für Informatik, Technische Universität München, 2000.
2. F. Brandt. Cryptographic protocols for secure second-price auctions. In M. Klusch and F. Zambonelli, editors, *Cooperative Information Agents V*, volume 2182 of *Lecture Notes in Artificial Intelligence*, pages 154–165, Berlin et al., 2001. Springer.
3. F. Brandt, W. Brauer, and G. Weiß. Task assignment in multiagent systems based on Vickrey-type auctioning and leveled commitment contracting. In M. Klusch and L. Kerschberg, editors, *Cooperative Information Agents IV*, volume 1860 of *Lecture Notes in Artificial Intelligence*, pages 95–106, Berlin et al., 2000. Springer.
4. K. Danielsen and M. Weiss. User control modes and IP allocation. <http://www.press.umich.edu/jep/works/DanieContr.html>, 1995. presented at MIT Workshop on Internet Economics.
5. K.E. Drexler and M.S. Miller. Incentive engineering for computational resource management. In B.A. Huberman, editor, *The Ecology of Computation*. The Netherlands, 1988.
6. B. Huberman and S.H. Clearwater. A multiagent system for controlling building environments. In *Proceedings of the 1st International Conference on Multiagent Systems (ICMAS-95)*, pages 171–176, Menlo Park, CA, 1995. AAAI Press.
7. P. Jehiel, B. Moldovanu, and E. Stacchetti. How (not) to sell nuclear weapons. *American Economic Review*, 86:814–829, 1996.
8. J.K. MacKie-Mason and H.R. Varian. Pricing the internet. In *Proceedings of the Public Access to the Internet Conference*. JFK School of Government, 1993.
9. R.P. McAfee and J. McMillan. Auctions and Bidding. *Journal of Economic Literature*, 25:699–738, 1987.
10. P.R. Milgrom and R.J. Weber. A Theory of Auctions and Competitive Bidding. *Econometrica*, 50:1089–1122, 1982.
11. E. Rasmusen. *Games and Information*. Basil Blackwell, 1995.
12. M.H. Rothkopf and R.M. Harstad. Two models of bid-taker cheating in Vickrey auctions. *Journal of Business*, 68(2):257–267, 1995.
13. M.H. Rothkopf, T.J. Teisberg, and E.P. Kahn. Why are Vickrey auctions rare? *Journal of Political Economy*, 98(1):94–109, 1990.
14. T.W. Sandholm. Limitations of the Vickrey auction in computational multiagent systems. In *Proceedings of the 2nd International Conference on Multiagent Systems (ICMAS-96)*, Menlo Park, CA, 1996. AAAI Press.
15. T.W. Sandholm and V.R. Lesser. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, pages 328–335, 1995.
16. T.W. Sandholm and V.R. Lesser. Advantages of a leveled commitment contracting protocol. In *Proceedings of the 13th National Conference on Artificial Intelligence (AAAI-96)*, pages 126–133, 1996.
17. Sun Microsystems. Webmart: Managing shared resources by market mechanisms. <http://www.sun.com:80/960201/cover/webmart.html>, 1996.
18. W. Vickrey. Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.
19. C.A. Waldspurger, T. Hogg, B. Huberman, J.O. Kephart, and W.S. Stornetta. Spawn: A distributed computational economy. *IEEE Transactions on Software Engineering*, 18(2):103–117, 1992.
20. G. Weiß, editor. *Multiagent Systems. A Modern Approach to Distributed Artificial Intelligence*. The MIT Press, Cambridge, MA, 1999.