

Spiteful Bidding in the Dollar Auction

Marcin Waniek

University of Warsaw
vua@mimuw.edu.pl

Tomasz Michalak

University of Oxford
and University of Warsaw
tomasz.michalak@cs.ox.ac.uk

Agata Nieścieruk

Polish-Japanese Academy of
Information Technology
agata.niescieruk@pjwstk.edu.pl

Talal Rahwan

Masdar Institute of Science and Technology
trahwan@gmail.com

Abstract

Shubik’s (all-pay) dollar auction is a simple yet powerful auction model that aims to shed light on the motives and dynamics of conflict escalation. Common intuition and experimental results suggest that the dollar auction is a trap, inducing conflict by its very design. However, O’Neill [1986] proved the surprising fact that, contrary to the experimental results and the intuition, the dollar auction has an immediate solution in pure strategies, *i.e.*, theoretically it should not lead to conflict escalation. In this paper, inspired by the recent literature on spiteful bidders, we ask whether the escalation in the dollar auction can be induced by meanness. Our results confirm this conjecture in various scenarios.

1 Introduction

On the surface, many social situations appear to be a trap, where it seems a bad idea to invest further resources, but also seems bad to retract and lose already-invested resources. Such dilemmas are often faced by lobbyists who battle with each other in a costly and seemingly endless process of acquiring a public contract [Fang, 2002], by oligopolistic companies pressured to invest in R&D only because the competitors have done so [Dasgupta, 1986], or by many ready-to-marry people who feel trapped in their relationships [Rhoades *et al.*, 2010].

Mimicking human societies, conflicts also abound in AI and in Multi-Agent Systems (MAS) [Müller and Dieng, 2000; Vasconcelos *et al.*, 2009; Castelfranchi, 2000]. While such conflicts often used to be considered in the literature as failures or synchronisation problems [Weyns and Holvoet, 2003; Tessier *et al.*, 2002], it has been recently argued that there is a need for more advanced studies, including analyses of conflict generation, escalation, or detection in MAS [Campos *et al.*, 2013].

In this paper, we study Shubik’s [1971] *dollar auction*—a simple yet powerful all-pay auction model that aims to shed light on the motives and dynamics of conflict escalation. In this auction, two bidders i and j compete for a dollar; the highest bidder wins the prize, but both the winner and the

loser have to pay their bids to the auctioneer. One might argue that it is best not to participate in this auction. However, this is not always possible. Furthermore, the possibility that a player may choose not to bid creates a clear incentive for the other player to bid and get the prize. Matter-of-factly, this reasoning is the centrepiece of the entire dollar auction mechanism that ultimately pushes players towards conflict escalation.

The above dollar auction game has become an influential abstraction of conflict escalation processes. It shows that conflicts may reach irrational levels despite the fact that, locally, every single participant makes a rational decision. Similar patterns of behaviour are observed in “clinical” experiments with the dollar auction—more often than not, a dollar bill is sold for considerably more than a dollar [Shubik, 1971; Kagel and Levin, 2008]. One of the key reasons behind this “paradox of escalation” is that a rational strategy to play this game is far from obvious; it is difficult to make an optimal choice between when “to quit” and when “to bid” (which is also true in many real-life situations).

In his beautiful paper, O’Neill [1986] offered a surprising solution to the dollar auction—he proved that, assuming finite budgets of players, in all equilibria in pure strategies, only one player bids and wins the prize. The exact amount of such a “golden” bid is a non-trivial function of the stake, the budgets, and the minimum allowable increment.

Does O’Neill’s result mean that the conflict in the dollar auction does not escalate after all? The issue was revisited by Leininger [1989], who showed that the escalation can be justified in this game because there exist equilibria with escalation in mixed strategies. Later on, Demange [1992] proved that, if there is some uncertainty about the strength of the players, then the only stable equilibrium may entail escalation.

In this paper, we reconsider O’Neill’s results in pure equilibria from a different perspective. Following recent literature on spiteful bidders (*e.g.* [Brandt *et al.*, 2005]), we ask whether the escalation in the dollar auction may actually be caused by the meanness of some participants. Do some of us put others in an inauspicious position simply because of spite, rather than greed? Do we allow ourselves to be dragged along simply because we do not expect a spiteful opponent? Our results confirm this conjecture in various scenarios.

2 Preliminaries

In this section, we formally introduce the notation and rules of the dollar auction and the concept of spitefulness.

The Dollar Auction: The auction setting proposed by Shubik [1971] consists of two players, $N = \{1, 2\}$. We will often refer to one of them as i and the other as j . The players can declare bids x_i and x_j , respectively, that are multiples of a unit $c \in \mathbb{N}$. The winner of the auction receives the stake $s \in \mathbb{N}$. Without loss of generality, unless stated otherwise, we assume that $c = 1$. Some arbitrary mechanism chooses the player who places the first bid (from this moment on we call him player 1). Player 1 can either place a bid or pass the turn to player 2. If player 2 does not want to bid either, the auction ends. Otherwise, the turn moves back to player 1 and so on and so forth. At any moment, if a player quits, he receives nothing, while the other receives the stake s . After the auction, both players must pay their bids to the auctioneer, regardless of the identity of the winner. Players cannot make deals, form a coalition or make threats.

O'Neill [1986] extended the above model by explicitly assuming that the players have finite budgets $b_1, b_2 \in \mathbb{N}$, where $b_1 > s > c$ and $b_2 > s > c$. These budgets naturally constrain the players, e.g., if $x_i \geq b_j$ then player j cannot outbid player i , in which case i wins. Figure 2 presents an example of the dollar auction where $b_1 = b_2 = 7$. Each node represents a pair of bids (x_1, x_2) . In the black nodes, player 1 chooses his bid and each solid arrow corresponds to one of his possible decisions. Analogically, the white nodes are the decision nodes of player 2 and the dotted arrows represent his possible bids. The nodes without an outgoing arrow represent situations in which the corresponding player has spent his entire budget. The auction ends either when one of the players chooses not to make his move, or when a node is reached that has no outgoing arrows.

O'Neill [1986] showed that in an auction with limited budgets, the starting player has a winning strategy; he should make a bid of $(b-1) \bmod (s-1) + 1$ and his opponent should pass, as he will never win the stake. We sketch a proof of this result in the proof of Theorem 1.

Spitefulness: Following, e.g. Brandt *et al.* [2005], we assume that a spiteful agent is interested in increasing his own profit while at the same time decreasing the profit of his opponent. More formally, the utility of player i is:

$$u_i = (1 - \alpha_i)p_i - \alpha_i p_j,$$

where p_i and p_j are the profits of the respective players, and $\alpha_i \in [0, 1]$ is the *spite coefficient* of player i . This coefficient indicates how important to player i is the loss of the opponent. When $\alpha_i = 0$, player i simply maximizes his own profit, in which case i is said to be a *non-spiteful* player. When $0 < \alpha_i \leq 1$, player i is said to be a *spiteful* player. Finally, when $\alpha_i = 1$, player i is said to be a *malicious* player, since he is solely interested in maximizing the loss of the opponent.

The concept of spitefulness is somewhat related to inequity aversion [Fehr and Schmidt, 1999], since a spiteful player suffers from any advantage of his opponent.

Let us define the profit of player i in the dollar auction as the difference between his initial budget b_i and his balance after

the auction. More formally, let us denote by (x_1, x_2) the final bids of the players, where $x_1 + x_2 > 0$.¹ Then, assuming that $x_i > x_j$, the profit of the winner of the auction is $p_i = s - x_i$, while the “profit” of the loser is $p_j = -x_j$. Formally, the utility of player i in the dollar auction is:

$$u_i = \begin{cases} \alpha_i x_j + (1 - \alpha_i)(s - x_i) & \text{if } x_i > x_j, \\ \alpha_i(x_j - s) - (1 - \alpha_i)x_i & \text{if } x_i < x_j. \end{cases}$$

3 Auction Settings

We studied various ways in which spitefulness can be introduced to the dollar auction. In particular, we considered:

- auctions in which one player is non-spiteful ($\alpha = 0$) and the other is spiteful/malicious ($\alpha \in (0, 1]$).
- auctions in which both players are spiteful/malicious;

Moreover, we considered two alternative assumptions about knowledge/rationality of the players. First, for setting (a) we assumed that the non-spiteful player does not suspect the spitefulness of his opponent, and so uses a strategy that is optimal against a non-spiteful opponent. After that, for both settings (a) and (b) we considered the case in which both players are aware of each others' spite coefficients. Finally, we considered all cases assuming both equal and unequal budgets. For all of the aforementioned settings, we consider subgame perfect equilibria.

Due to space constraints we focus below on setting (a) assuming limited knowledge of the non-spiteful player, results of which we found most interesting. However, at the end we briefly summarise the results for alternative settings.

4 An Auction with Equal Budgets

In this section, we analyze the dollar auction between players with equal budgets, i.e., $b_i = b_j = b$ where $b > s$. A spiteful player has a spite coefficient ranging from 1, if he is malicious, to (almost) 0. On the other hand, a non-spiteful player has a spite coefficient of 0.

4.1 A Malicious Player ($\alpha_j = 1$)

If $\alpha_i = 0$ and $\alpha_j = 1$, then player i (whom we assume to be non-spiteful and does not suspect the spitefulness of his opponent), is challenged by player j who is actually malicious, meaning that the goal of j is to maximize the loss of i , no matter what the cost. In the theorem below, we show that the optimal strategy of the malicious player is to lure the non-spiteful player to continue bidding as long as possible. Still, it is the malicious bidder who gets the stake at the end.

Theorem 1. *Let i be a non-spiteful player ($\alpha_i = 0$) who follows the strategy by O'Neill [1986], and let player j be malicious ($\alpha_j = 1$). The optimal strategy of j is to bid:*

$$x_j = \begin{cases} x_i + 1 & \text{if } x_i < b - (s - 1), \\ b & \text{otherwise.} \end{cases}$$

¹The assumption that $x_1 + x_2 > 0$ guarantees that the bidding has actually started, i.e., at least one bid has been made.

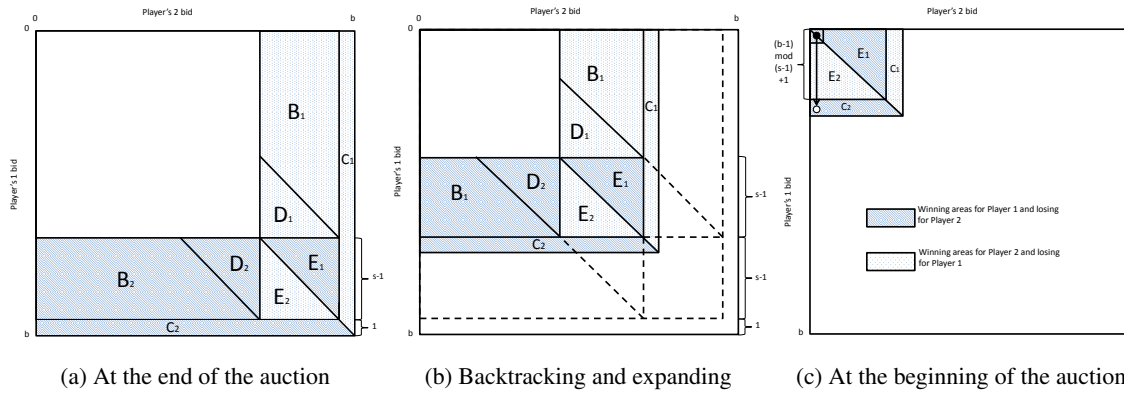


Figure 1: Winning and losing areas during different moments of the auction.

Proof. We begin the proof by presenting the optimal strategy of a non-spiteful player who does not suspect spitefulness from the other player. This strategy was derived by O'Neill [1986]. To this end, consider the graph-based representation of the dollar auction, an example of which was discussed earlier in Section 2 and illustrated in Figure 2. Recall that every node in this graph represents a pair of bids. We call node (x_1, x_2) a *winning node* for player i if, by starting the game from this node, player i is guaranteed to eventually win the stake s . Importantly, while the bid of player i may exceed s , we note that it would be irrational to do so (the only way to increase a non-spiteful player's utility is to win the auction and get the stake s ; increasing the bid by more than s outweighs any possible profit). Finally, recall that the game ends if one player exceeds his budget.

We will now analyze each node of the graph to determine whether it is a winning node for player i . Figure 1a highlights the parts of the graph that consist of winning nodes. Specifically, the width of area C_i is one node, while the widths of areas B_i , D_i and E_i are $s - 1$ nodes each. Any node in C_j is a winning one for player i , since player j cannot bid more than b and his only choice is to pass. Moreover, area E_i consists of winning nodes for i , since by raising his bid by less than s , he can move to area C_j , win the stake s and finish the auction. Nodes in area D_i are losing nodes for i , since by raising his bid by less than s he can only move to area E_j , which is a winning area for j . Thus, i 's optimal move in area D_i is to pass. Finally, in area B_i , the only valid bid increments are those greater than, or equal to s , so it is also optimal for i to pass. The analysis for areas C_j , B_j , D_j , E_j is identical.

We can now move on to the analysis of the areas of the graph that are gradually closer to $(0, 0)$, as illustrated in Figure 1b. The areas that we have just analyzed now play the role of areas C_1 and C_2 . One could, of course, bid higher, but moving into C_i always optimizes cost. All arguments previously stated for areas B_i , D_i , E_i still stand.

We can repeat this process until we reach an area whose dimensions are less than $s \times s$, as illustrated in Figure 1c. In this case, areas C_1 and C_2 are winning areas for players 2 and 1, respectively. Consequently, area E_i is winning for player i , where his optimal choice is to make the bid $(b - 1) \bmod (s - 1) + 1$ and move to area C_j . As a result, the non-spiteful player

who makes the first move is always able to ensure his own victory in the auction against the other non-spiteful player.

Figure 3 depicts the winning (shaded) areas for player 1 and the winning moves for each such area, as well as the losing (white) areas for player 1, where the optimal choice for him is to pass.

Let us now assume that player i is non-spiteful and player j is malicious. Recall that the utility of the malicious player is: $u_j^{mal} = -p_i$. Naturally, the malicious player cannot make the first move and arrive at an area that is outside of the winning areas of the non-spiteful player. If he did so, then the non-spiteful player would simply pass, leaving both players with zero utility.

Next, we will show that the malicious player maximizes his utility if he plays as the second player. Hence, given the chance to move first, he will always pass the move to the opponent. Thus, throughout the remainder of this proof, player 2 will be the malicious one (*i.e.*, $i = 1$ and $j = 2$).

Since the utility of the malicious player j is: $u_j^{mal} = -p_i$, it is maximized when the game ends in area E_1 , or in the adjacent nodes of area C_1 in Figure 1a. However, player i will not stay in area E_1 , but will rather make a move towards C_2 . On the other hand, nodes from C_1 can only be reached via a move from B_1 or D_1 . As visible in Figure 3, $(b - (s - 1), b)$ is the only node that player j is able to reach. His utility at this node is $b - s + 1$.

The next best group of nodes in terms of utility u_j^{mal} are those in C_2 from Figure 1a; player i wins the stake, but is forced to use all of his budget. The utility of player j is therefore $b - s$, which is lower than in the node $(b - (s - 1), b)$.

A bid by a malicious player j that is always one unit higher than the bid of a non-spiteful player i keeps the game in the nodes that player i considers to be "winning" nodes. When the bid of player i reaches $b - (s - 1)$, player j should bid b , thereby reaching the node in which he achieves his highest possible utility. \square

4.2 A Spiteful Player ($0 < \alpha_j < 1$)

This subsection describes strategies for player j who is spiteful but not malicious (*i.e.*, $0 < \alpha_j < 1$). His opponent is player i , who is non-spiteful (*i.e.*, $\alpha_i = 0$) and does not suspect the spitefulness of his opponent, meaning that he is following the strategy proposed by O'Neill [1986]. We divide our analysis

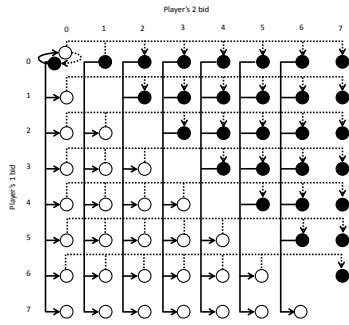


Figure 2: The graph of the dollar auction for $b = 7$.




Figure 3: Winning areas (blue) & winning moves of a non-spiteful player 1.




Figure 4: Non-spiteful player 1, spiteful player 2. End-nodes marked in red.

into three parts: *weakly* spiteful player with $\alpha_j \in (0, \frac{1}{2})$, *moderately* spiteful player with $\alpha_j = \frac{1}{2}$ and *strongly* spiteful player with $\alpha_j \in (\frac{1}{2}, 1)$. First, let us start with the following two lemmas, which hold for all of the aforementioned three cases.

Lemma 1. *Let $i = 1$ be a non-spiteful player ($\alpha_i = 0$) who follows the strategy by O'Neill [1986], and let $j = 2$ be a spiteful player ($0 < \alpha_j < 1$). Furthermore, let x be the optimal initial bid of player i , and:*

$$X_0 = (0, x), X_k = (x + (k-1)(s-1), x + k(s-1)), \\ Y_0 = (x, 0), Y_k = (x + k(s-1), x + (k-1)(s-1) + 1),$$

where $k \in \{1, \dots, n\}$ and $n = \lfloor \frac{b}{s-1} \rfloor$.

The optimal end-node for player j (that maximizes his utility) is among nodes $X_0, \dots, X_n, Y_0, \dots, Y_n$.

Proof. First, let us define areas marked in Figure 4 as:

$$A_k = \{(x_{k,1}, d) : d \in \{x_{k,2}, \dots, b\}\}, \\ B_k = \{(y_{k,1}, d) : d \in \{y_{k,2}, \dots, y_{k,1} - 1\}\},$$

where $(x_{k,1}, x_{k,2}) = X_k$ and $(y_{k,1}, y_{k,2}) = Y_k$. These are the only possible end nodes of the auction. In particular, the auction cannot end in any of the unmarked nodes under the diagonal, because they can only be reached by a move by player i , but he will never make such a move (see the proof of Theorem 1). The auction can end in the marked areas under the diagonal, since player j can choose to pass at these nodes.

The auction cannot end in any of the areas marked E_1 . This is because player i can make moves that he considers winning from these nodes. Conversely, player i passes in all other areas above the diagonal. Now, we observe that the only difference between the nodes in marked areas is the value of the bid made by player j . For any $d \in \mathbb{N}$ we have $u_j((y, x+d)) \leq u_j((y, x))$. \square

Lemma 2. *Let $i = 1$ be a non-spiteful player following the strategy by O'Neill [1986], and let $j = 2$ be a spiteful player. The end-nodes X_1, \dots, X_n and Y_0, \dots, Y_n (defined in Lemma 1) can be reached by player j regardless of who starts the game, while X_0 can be reached by player j if he starts the game.*

Proof. Player j can reach any node Y_0, \dots, Y_n by bidding always one unit more than the non-spiteful player i . This way, in every move player i will make his optimal bid to enter Y_k . Now, player j can pass in node Y_k , or bid $x + s - 1$ in node $Y_{k-1} = (x, y)$ to reach X_k (for $k > 0$). Moreover, if player j starts, he can reach X_0 by making the bid $(b-1) \bmod (s-1) + 1$. As shown in Theorem 1, in every node X_k it is optimal for a non-spiteful player to pass. \square

Weakly-spiteful player with $\alpha_j \in (0, \frac{1}{2})$: As it turns out, a player with a small (*i.e.*, below $\frac{1}{2}$) spite coefficient behaves like a non-spiteful player if he starts the bidding. Otherwise, after his opponent's move, he passes or forces his opponent to pass. Figure 5a presents an example of a utility map for a weakly-spiteful player.

Theorem 2. *Let i be a non-spiteful player following the strategy by O'Neill [1986], and let j be a weakly-spiteful player (*i.e.*, $\alpha_j \in (0, \frac{1}{2})$). If j starts the bidding, it is optimal for him to bid like a non-spiteful player. If i moves first and makes bid x , then the optimal strategy for player j is to make the bid $x + s - 1$ when $x < \frac{\alpha_j s}{1 - \alpha_j} + 1$, and pass otherwise.*

Proof. Let x be the optimal initial move of a non-spiteful player, *i.e.*, $x = (b-1) \bmod (s-1) + 1$. Note that $x \in [1, s-1]$. Based on Lemma 1, the optimal end of an auction for player j is one of the nodes from the set $\{X_0, \dots, X_n, Y_0, \dots, Y_n\}$. Moreover, Lemma 2 states that all of these nodes can be reached by player j . Consider the utility function of player j in nodes X_k and X_{k+1} for $k \geq 1$. If $X_k = (a, b)$, then $X_{k+1} = (a + (s-1), b + (s-1))$. Hence,


$$u_j(X_{k+1}) - u_j(X_k) = (s-1)(2\alpha_j - 1) < 0.$$

Analogously, the utility function of player j is lower in Y_{k+1} than in Y_k . Thus, the optimal solution that can be reached by player j is in one of the nodes: $X_0 = (0, x)$, $X_1 = (x, x + s - 1)$, $Y_0 = (x, 0)$, $Y_1 = (x + s - 1, x + 1)$. Now, we have:

$$u_j(X_0) - u_j(X_1) = (s-1)(1 - \alpha_j) - \alpha_j x > 0, \\ u_j(X_0) - u_j(Y_0) = s - x > 0, \\ u_j(X_1) - u_j(Y_1) = 2 - \alpha_j > 0.$$

Thus, if player j starts the bidding, he prefers to bid x and end the game in X_0 . By comparing X_1 and Y_0 we have:

$$u_j(X_1) - u_j(Y_0) = x(\alpha_j - 1) + \alpha_j s + 1 - \alpha_j.$$



Note: Blue lines represent the same utility, red lines represents zero utility. The darker the color, the higher the utility.

Figure 5: Utility maps for a spiteful player $j = 2$ with different values of spite coefficient.

Thus, if $x < \frac{\alpha_j s}{1-\alpha_j} + 1$, then it is better for player j to bid and end the game in the state X_1 ; otherwise, player j passes. \square

Moderately-spiteful player with $\alpha_j = \frac{1}{2}$: The following theorem implies that a moderately-spiteful player, whose spite coefficient is $\frac{1}{2}$, acts like a non-spiteful player if he starts the auction, but can act like a malicious player otherwise. A sample utility map for such a player is depicted in Figure 5b.

Theorem 3. *Let i be a non-spiteful player (i.e., $\alpha_i = 0$) who is following the strategy by O'Neill [1986], and let j be a moderately-spiteful player, i.e., $\alpha_j = \frac{1}{2}$. If j starts the bidding, it is optimal for him to bid like a non-spiteful player. On the other hand, if i starts the bidding, the optimal strategy for j is to outbid i by 1 for any number of turns and finally raise his bid by $s - 1$.*

Proof. Let us denote by x the optimal initial move of a non-spiteful player, i.e., $x = (b-1) \bmod (s-1) + 1$. Note, that $x \in [1, s-1]$. Now, based on Lemmas 1 and 2, we can consider only nodes $X_1, \dots, X_n, Y_0, \dots, Y_n$ as the optimal end points of the auction. Since $u_j(X_{k+1}) = u_j(X_k)$ and $u_j(Y_{k+1}) = u_j(Y_k)$ for $k \geq 1$, the optimal reachable solution can be found by considering the nodes $X_0 = (0, x), X_1 = (x, x + s - 1), Y_0 = (x, 0), Y_1 = (x + s - 1, x + 1)$.

Now, we have: $u_j(X_0) = \frac{s-x}{2}$, $u_j(X_1) = \frac{1}{2}$, $u_j(Y_0) = \frac{x-s}{2}$, $u_j(Y_1) = -1$. Therefore, the following three inequalities hold: $u_j(X_0) \geq u_j(X_1)$, $u_j(X_1) > u_j(Y_1)$, $u_j(X_1) > u_j(Y_0)$. Thus, if player j starts, it is optimal for him to end the auction in the state X_0 ; otherwise, he should end in any of the nodes X_k . \square

Strongly-spiteful player with $\alpha_j \in (\frac{1}{2}, 1)$: A player with a high (i.e., $> \frac{1}{2}$) spite coefficient can act like a non-spiteful player if he starts the auction for very specific values of b and s . However, in most cases he acts like a malicious player; he forces a non-spiteful player to raise his bids and then wins the stake in the end. Figure 5c presents an example of a utility map for such a strongly-spiteful player.

Theorem 4. *Let i be a non-spiteful player (i.e., $\alpha_i = 0$) who is following the strategy by O'Neill [1986], and let j be a strongly-spiteful player (i.e., $\alpha_j \in (\frac{1}{2}, 1)$). Furthermore, let x be the optimal initial bid of a non-spiteful player. If j starts*

the bidding and $x < \frac{\alpha_j}{1-\alpha_j}(s-1) - \frac{2\alpha_j-1}{1-\alpha_j}b$, then the optimal strategy for him is to make the bid x . Otherwise, if i starts the bidding, then it is optimal for j to follow the optimal strategy of a malicious player, described in Theorem 1.

Proof. Let us denote by x the optimal initial move of a non-spiteful player, i.e., $x = (b-1) \bmod (s-1) + 1$. Note, that $x \in [1, s-1]$. Again, we will limit our analysis to nodes from the set $\{X_1, \dots, X_n, Y_0, \dots, Y_n\}$ based on Lemmas 1 and 2.

Consider the utility function of player j in nodes X_k and X_{k+1} for $k \geq 1$. If $X_k = (a, b)$, then $X_{k+1} = (a + (s-1), b + (s-1))$. Here, unlike the previous cases, we have:

$$u_j(X_{k+1}) - u_j(X_k) = (s-1)(2\alpha_j - 1) > 0.$$

An analogous analysis can be obtained for nodes Y_k and Y_{k+1} . Thus, the optimal reachable solution is one of the following four nodes: $X_0 = (0, x), X_n = (b - (s-1), b), Y_0 = (x, 0), Y_n = (b, b - (s-2))$. Now, we have:

$$u_j(X_0) - u_j(Y_0) = s - x > 0,$$

$$u_j(X_n) - u_j(Y_n) = 2 - \alpha_j > 0,$$

$$u_j(X_n) - u_j(Y_0) = \alpha_j(2b - s - x + 1) + s - b > 0.$$

Thus, it is always better for player j to end the game in node X_0 or in node X_n , rather than in node Y_0 or in node Y_n . Moreover,

$$u_j(X_n) - u_j(X_0) = (2\alpha_j - 1)b + (1 - \alpha_j)x + \alpha_j(1 - s).$$

Therefore, if $x < \frac{\alpha_j}{1-\alpha_j}(s-1) - \frac{2\alpha_j-1}{1-\alpha_j}b$ holds, then it is better for player 1 to end the game with the first bid (if player j starts). Otherwise, X_n is the optimal solution. \square

4.3 Alternative Settings

Here, we report results for the alternative auction settings in which both players know each others' spite coefficients, α_i and α_j . We omit proofs due to space constraints.

Theorem 5. *Let i be a non-spiteful, weakly-spiteful or moderately-spiteful player (i.e., $\alpha_i \leq \frac{1}{2}$) and let j be a spiteful player with $\alpha_j \in (0, 1]$. The optimal strategy for player i is to either follow the strategy proposed by O'Neill (if $\alpha_j \leq \frac{1}{2}$) or pass (if $\alpha_j > \frac{1}{2}$). The optimal strategy for player j when $\alpha_j \leq \frac{1}{2}$ is the same as for player i , and when $\alpha_j > \frac{1}{2}$ is to bid 1 and continue with overbidding if player i bids.*

Theorem 6. *Let i be a strongly-spiteful or malicious player ($\alpha_i > \frac{1}{2}$) and let j be a spiteful player with $\alpha_j \in (0, 1]$. The optimal strategy for player i is to follow the strategy of a malicious player, described in Theorem 1. The optimal strategy for player j when $\alpha_j > \frac{1}{2}$ is the same as for player i , and when $\alpha_j \leq \frac{1}{2}$ is to pass.*

5 An Auction with Unequal Budgets

We now consider auctions where player budgets are unequal.

5.1 Strategies of a non-spiteful player

As shown by O'Neill [1986], a non-spiteful player has a certain strategy for an auction with unequal budgets. If he starts the game with a higher budget, he makes a bid of just one unit and expects a non-spiteful player to pass. If he starts the game with a lower budget, he makes a bid of $s - 1$ and expects a non-spiteful player to pass.

5.2 Strategies of a malicious player

Now let us consider an auction with a malicious player j and a non-spiteful player i (who does not suspect the spitefulness of his opponent). As it turns out, a malicious player with a higher budget can force a non-spiteful opponent to pay $s - 1$ and then force him to lose the stake; in this case the utility of j is $u_j^{mal} = s - 1$. On the other hand, if a malicious player starts an auction with a lower budget b_j , he can force a non-spiteful opponent to pay $b_j + 1$ in order to win the stake; in this case the utility of j is $u_j^{mal} = b_j + 1 - s$. Counter-intuitively, this means that a malicious player with lower budget can actually be more powerful than with higher budget; this is the case when: $b_j > 2s - 2$.

Surprisingly, if a malicious player starts an auction with a lower budget b_j , he can force a non-spiteful opponent to pay $b_j + 1$ and then force him to lose. In other words, counter-intuitively, a malicious player with lower budget can actually be more powerful than a malicious one with higher budget.

Theorem 7. *Let i be non-spiteful, and j be malicious. If $b_j > b_i$, and j moves second, his optimal strategy is to bid s as an answer to i 's bid of $s - 1$. Player i then passes. If j starts the auction, it is better for him to let i move first.*

Proof. (Sketch) Player i makes the bid $s - 1$ because it is an optimal bid against a non-spiteful player. He can gain the utility of 1, when player j passes. If player j bids s or higher, then player i passes, as he cannot win due to having a smaller budget. This yields his minimal utility of $-(s - 1)$, and the maximal utility of a malicious player j .

Suppose j starts the auction and bids at least 1. Since i has a smaller budget, he knows he cannot win. To cut his losses, he passes at the very beginning, giving malicious player j zero utility. On the other hand, if player j lets his opponent move first, he gets positive utility, as shown above. \square

Theorem 8. *Consider the dollar auction with a malicious player j , a non-spiteful player i and budgets $b_j < b_i$. Player j can force player i to pay $b_j + 1$ in order to acquire the stake.*

Proof. (Sketch) Player j prefers to continue the auction, as his utility rises as the auction progresses. The optimal strategy for player j is to always outbid his opponent by 1. Since player i can always win, he can make a minimum raise of 1 to minimize his losses. When the bid of player i exceeds $b_j - (s - 2)$, player j should raise his bid to b_j . Player i can then get the stake s by raising his bid by $s - 1$ to $b_j + 1$. Since it is never rational for player i to bid more than $b_j + 1$, this strategy of player j grants him maximal utility. \square

5.3 Alternative Settings

Here, we report again results for the alternative auction settings in which both players know each others' spiteful coefficients, α_i and α_j . We omit proofs due to space constraints.

Theorem 9. *Let i be a non-spiteful, weakly-spiteful, or moderately-spiteful player (i.e., $\alpha_i \leq \frac{1}{2}$) and let j be a spiteful player with $\alpha_j \in (0, 1]$. The optimal strategy for player i is to either follow the strategy of O'Neill (if $\alpha_j \leq \frac{1}{2}$) or pass (if $\alpha_j > \frac{1}{2}$). The optimal strategy for player j when $\alpha_j \leq \frac{1}{2}$ is the same as for player i , and when $\alpha_j > \frac{1}{2}$ is to bid 1 and then to continue overbidding if player i bids.*

Theorem 10. *Let i be a strongly-spiteful or malicious player ($\alpha_i > \frac{1}{2}$) and let j be a spiteful player with $\alpha_j \in (0, 1]$. When $\alpha_j \geq \frac{1}{2}$ and $b_i < b_j$ the optimal strategy for i is to bid b_i (regardless of whether he starts). When $\alpha_j \geq \frac{1}{2}$ and $b_i > b_j$ the optimal strategy for i is to start bidding with $b_j - 1$ and answer with $b_j + 1$ to the opponent's bid of b_j . When $\alpha_j < \frac{1}{2}$ player j will not enter the auction, so player i should simply bid 1. The optimal strategy for player j when $\alpha_j > \frac{1}{2}$ is the same as for player i , and when $\alpha_j \leq \frac{1}{2}$ is to pass.*

6 Related work

In this section we comment on the bodies of literature related to the key characteristics of our auction setting: spitefulness, all-pay format, and the assumption of the finite budget.

Spitefulness in auctions: On top of some work in game theory [Baye *et al.*, 1996], and experimental game theory [Bolle *et al.*, 2013] in particular, the analysis of spitefulness in simple types of auctions can be found, among others, in the works by Morgan *et al.* [2003], Brandt and Weiß [2002], Brandt *et al.* [2005], Babaioff *et al.* [2007], and Sharma and Sandholm [2010]. An interesting study of vindictive behaviours in auction-like settings (where rivals engage in aggressive retaliatory behaviors) can be found in [Bolle *et al.*, 2013]. Similarly, vindictive bidding in keyword auctions was studied by Zhou and Lukose [2007].

All pay auctions: While all-pay auctions are a relatively rare auction format, they have been extensively studied [DiPalantino and Vojnovic, 2009; Lewenberg *et al.*, 2013] as they model various realistic settings in which the prize is awarded, often implicitly, on the basis of contestants' efforts. These include lobbying, job-promotion competitions, political campaigns, and R&D competitions, to name a few [Lev *et al.*, 2013]. For an overview of experimental research on all-pay auctions see [Dechenaux *et al.*, 2012].

The finite budget: While in many studies on auctions the existence of the budget can be neglected, this is not necessarily so in all-pay auctions. The budget constrains in all-pay auctions were studied by Che and Gale [1996].

7 Conclusions

Our results suggest that the escalation in real-life experiments with the dollar auction could be related not only to the desire to win but also (at least to some extent) to human meanness. In various scenarios in which a non-spiteful bidder unwittingly bids against a spiteful one, the conflict escalates. Not only can the spiteful bidder force the non-spiteful opponent to spend most of the budget but he also often wins the prize. Surprisingly, a malicious player with a smaller budget is likely to plunge the opponent more than a malicious player with a bigger budget. Thus, a malicious player should not only hide his real preferences but also the real amount of his budget. Intuitively, a weak, easy-to-overcome bait may seem more attractive than a stronger one.

For future work, it would be interesting to study how the results of the dollar auction change if bidders are not spiteful but rather altruistic [Chen *et al.*, 2011], to model auctions with communication between players, or to model players with bounded rationality.

Acknowledgements

This work was supported by the Polish National Science Centre grant 2012/05/B/ST6/03364. Tomasz Michalak was supported by the European Research Council under Advanced Grant 291528 (“RACE”).

References

- [Babaioff *et al.*, 2007] M. Babaioff, R. Kleinberg, and C. H. Papadimitriou. Congestion games with malicious players. In *EC'07*, pages 103–112. ACM, 2007.
- [Baye *et al.*, 1996] M. R. Baye, D. Kovenock, and C. G. De Vries. The all-pay auction with complete information. *Economic Theory*, 8(2):291–305, 1996.
- [Bolle *et al.*, 2013] F. Bolle, J. Tan, and D. Zizzo. Vendettas. *American Economic Journal: Microeconomics*, 2013.
- [Brandt and Weiß, 2002] F. Brandt and G. Weiß. Antisocial agents and vickrey auctions. In *Intelligent Agents VIII*, volume 2333 of *Lecture Notes in Computer Science*, pages 335–347. Springer Berlin Heidelberg, 2002.
- [Brandt *et al.*, 2005] F. Brandt, T. Sandholm, and Y. Shoham. Spiteful bidding in sealed-bid auctions. In *Computing and Markets*, 2005.
- [Campos *et al.*, 2013] J. Campos, C. Martinho, and A. Paiva. Conflict inside out: A theoretical approach to conflict from an agent point of view. AAMAS '13, pages 761–768. International Foundation for Autonomous Agents and Multiagent Systems, 2013.
- [Castelfranchi, 2000] C. Castelfranchi. Conflict ontology. In *Computational Conflicts*, pages 21–40. Springer, 2000.
- [Che and Gale, 1996] Y. Che and I. Gale. Expected revenue of all-pay auctions and first-price sealed-bid auctions with budget constraints. *Economics Letters*, 50(3):373–379, 1996.
- [Chen *et al.*, 2011] P. A. Chen, B. de Keijzer, D. Kempe, and G. Schäfer. Robust price of anarchy for atomic games with altruistic players. Technical report, CWI Amsterdam, 2011.
- [Dasgupta, 1986] P. Dasgupta. The theory of technological competition. In *New developments in the analysis of market structure*. Cambridge: MIT Press, Cambridge, 1986.
- [Dechenaux *et al.*, 2012] E. Dechenaux, D. Kovenock, and R. M. Sheremeta. A survey of experimental research on contests, all-pay auctions and tournaments. WZB Discussion Paper SP II 2012-109, Berlin, 2012.
- [Demange, 1992] G. Demange. Rational escalation. *Annales d'économie et de Statistique*, (25/26):pp. 227–249, 1992.
- [DiPalantino and Vojnovic, 2009] D. DiPalantino and M. Vojnovic. Crowdsourcing and all-pay auctions. In *EC'09*, pages 119–128. ACM, 2009.
- [Fang, 2002] H. Fang. Lottery versus all-pay auction models of lobbying. *Public Choice*, 112(3-4):351–371, 2002.
- [Fehr and Schmidt, 1999] Ernst Fehr and Klaus M Schmidt. A theory of fairness, competition, and cooperation. *Quarterly journal of Economics*, pages 817–868, 1999.
- [Kagel and Levin, 2008] J. H. Kagel and D. Levin. Auctions: a survey of experimental research, 1995–2008. 2008.
- [Leininger, 1989] W. Leininger. Escalation and cooperation in conflict situations the dollar auction revisited. *Journal of Conflict Resolution*, 33(2):231–254, 1989.
- [Lev *et al.*, 2013] O. Lev, M. Polukarov, Y. Bachrach, and J. S. Rosenschein. Mergers and collusion in all-pay auctions and crowdsourcing contests. AAMAS '13, pages 675–682. International Foundation for Autonomous Agents and Multiagent Systems, 2013.
- [Lewenberg *et al.*, 2013] Y. Lewenberg, O. Lev, Y. Bachrach, and J. S. Rosenschein. Agent failures in all-pay auctions. IJCAI/AAAI, 2013.
- [Morgan *et al.*, 2003] J. Morgan, K. Steiglitz, and G. Reis. The spite motive and equilibrium behavior in auctions. *Contributions in Economic Analysis & Policy*, 2(1), 2003.
- [Müller and Dieng, 2000] H. J. Müller and R. Dieng. On conflicts in general and their use in ai in particular. In *Computational conflicts*, pages 1–20. Springer, 2000.
- [O'Neill, 1986] B. O'Neill. International escalation and the dollar auction. *Journal of Conflict Resolution*, 30(1):33–50, 1986.
- [Rhoades *et al.*, 2010] G. K. Rhoades, S. M. Stanley, and H. J. Markman. Should i stay or should i go? predicting dating relationship stability from four aspects of commitment. *Journal of family psychology*, 24(5):543, 2010.
- [Sharma and Sandholm, 2010] A. Sharma and T. Sandholm. Asymmetric spite in auctions. In *AAAI*, 2010.
- [Shubik, 1971] M. Shubik. The dollar auction game: A paradox in noncooperative behavior and escalation. *Journal of Conflict Resolution*, pages 109–111, 1971.
- [Tessier *et al.*, 2002] C. Tessier, L. Chaudron, and H. J. Miller. Agents' conflicts: New issues. In *Conflicting Agents, volume 1 of Multiagent Systems, Artificial Societies, And Simulated Organizations*, pages 1–30. Springer US, 2002.
- [Vasconcelos *et al.*, 2009] W. W. Vasconcelos, M. J. Kollingbaum, and T. J. Norman. Normative conflict resolution in multi-agent systems. *AAMAS'09*, 19(2):124–152, 2009.
- [Weyns and Holvoet, 2003] D. Weyns and T. Holvoet. Regional synchronization for simultaneous actions in situated multi-agent systems. pages 497–510, 2003.
- [Zhou and Lukose, 2007] Y. Zhou and R. Lukose. Vindictive bidding in keyword auctions. In *ICEC*, pages 141–146. ACM, 2007.