

# A Simple Shill Bidding Agent

Jarrod Trevathan

*School of Maths, Physics and IT*

*James Cook University*

*Email: jarrod.trevathan@jcu.edu.au*

Wayne Read

*School of Maths, Physics and IT*

*James Cook University*

*Email: wayne.read@jcu.edu.au*

**Abstract**—Shill bidding is where fake bids are introduced into an auction to drive up the final price for the seller, thereby defrauding legitimate bidders. Although shill bidding is strictly forbidden in online auctions such as eBay, it is still a major problem. This paper presents a software bidding agent that follows a shill bidding strategy. The malicious bidding agent was constructed to aid in developing shill detection techniques. The agent incrementally increases an auction's price, forcing legitimate bidders to submit higher bids in order to win the item. The agent ceases bidding when the desired profit from shilling has been attained, or in the case that it is too risky to continue bidding without winning the auction. The agent's ability to inflate the price has been tested in a simulated marketplace and experimental results are presented. This is the first documented bidding agent that perpetrates auction fraud. We do not condone the use of the agent outside the scope of this research.

## I. INTRODUCTION

Online auction fraud is rampant and every year increasingly costs victims millions in stolen money and sanity. Shill bidding is the act of introducing spurious bids into an auction on the seller's behalf, with the intent to artificially inflate the item's price. Bidders who engage in shilling are referred to as 'shills'. To win the item, a legitimate bidder must outbid a shill's price. If the shill accidentally wins, then the item is re-sold in a subsequent auction. Shill bidding defrauds legitimate bidders as they are forced into paying significantly more for the item.

There are many auction types (e.g., Vickrey, Dutch, etc.). The most popular is the English auction. In an English auction, bidders outbid each other in an attempt to win an item. The winner is the bidder with the highest bid. English auctions are employed in online auctions such as those offered by eBay<sup>1</sup> and Ubid<sup>2</sup>. The English auction is particularly susceptible to shill bidding practices. Shill bidding is strictly forbidden by commercial online auctioneers, and is a prosecutable offence (see [6]). However, the online environment makes shilling easy as bidders are anonymous.

There is often much confusion regarding what constitutes shill behaviour. Bidding behaviour that might seem suspicious, could in fact turn out to be innocent. Furthermore, a shill can engage in what seems to be a limitless number of strategies. This makes it difficult to detect shill bidding. While the online auctioneers monitor their auctions for shilling, there is little academic material available on proven shill detection techniques.

The Research Auction Server (RAS)<sup>3</sup> at James Cook University, is an online server for conducting research into security issues regarding online auctions (see [7]). We are developing methods to detect fraudulent bidding behaviour such as shilling (see [9]). Both real and simulated auctions are conducted to test the effectiveness of the detection methods. To aid in testing, we developed a software bidding agent that bids in a manner consistent with a shill.

A *software bidding agent* is a program that bids on a human bidder's behalf. Agents follow a predetermined strategy, typically with the goal of winning an auction for the minimal amount. In an English auction, a bidding agent is permitted to outbid any bid until the bidding price exceeds a maximum amount specified by the human bidder. In auctions that can last days or weeks, bidding agents remove the need for a bidder to constantly observe an auction. A bidding agent monitors the auction proceedings for any price activity, and responds in accordance with its programmed strategy.

Numerous bidding agents have been proposed in literature (see [1], [3], [4], [5]). The Trading Agent Competition (TAC)<sup>4</sup> (see [10]) pits bidding agents against each other in an elaborate economic game. The TAC server allows bidders to write their own agents using an application programming interface. Agents are assessed on their ability to acquire resources in the most efficient manner. TAC is mainly concerned with furthering the performance aspects of (non-fraudulent) bidding agents. It does not focus exclusively on agent security. All TAC agents are required to behave strictly within the auction's rules, and are disqualified if they act with a malicious intent to influence the auction proceedings.

To our knowledge no literature exists for a type of bidding agent we refer to as a *malicious bidding agent*. Similar to a virus or worm, a malicious bidding agent is a bidding agent that behaves with an intent to do an auction harm in some manner. This might be in the form of inflating the final price by shilling, attacking the cryptographic protocols of a "secure" auction system, or launching a denial of service attack against the Auctioneer.

Automated agents for conducting electronic commerce are becoming more common. The idea has been touted that all human input in auctions will eventually be done using autonomous bidding agents. However, such an environment

<sup>1</sup><http://www.ebay.com>

<sup>2</sup><http://www.ubid.com>

<sup>3</sup><http://auction.math.jcu.edu.au>

<sup>4</sup><http://www.sics.se/tac/>

would definitely spawn undesirable behaviour (e.g., cheating, stealing payments, etc.). (See [8].) Furthermore, undesirable groups such as terrorists can obtain funds through fraudulent activities in auctions. In an extreme scenario, a malicious agent could hinder the world's stock exchanges, in an attempt to undermine the financial system. Therefore, the threat posed by malicious bidding agents to electronic commerce is very serious. As a result, no one is willing to totally trust agent-based negotiation.

Existing agents operate in a controlled and near perfect environment. As it is not permitted to use malicious agents in commercial online auctions, or in TAC, we have created an agent interface for RAS. RAS allows the agents to be tested in a controlled (and legal) manner.

This paper presents a software bidding agent that follows a shill bidding strategy. The agent incrementally increases the auction's price, forcing legitimate bidders to submit higher bids in order to win. The agent ceases bidding when the desired profit from shilling has been attained, or in the case that it is too risky to continue bidding without winning the auction. The agent's performance has been tested against non-fraudulent bidding agents in a simulated market place. Experimental results show that the agent is able to successfully increase the average winning price. This paper *does not condone* the use of these agents in any manner outside the scope of this research. By developing malicious bidding agents, we hope to better understand the characteristics of such agents, and how to protect against the damage they inflict.

This paper is organised as follows: Section II describes general shill behaviour. Section III presents a shill bidding agent that shills in a single auction. Section IV evaluates the agent's ability to shill, contrasting safe and risky approaches. Section V provides some concluding remarks and avenues for future work.

## II. SHILL BEHAVIOUR

This section provides an insight into general shill behaviour. It describes a shill's characteristics and strategies, and contrasts shilling with another type of bidding behaviour referred to as *sniping*.

The main goal for shilling is to artificially inflate the price for the seller beyond what legitimate bidders would otherwise require to win the item. The pay-off for the seller is the difference between the final price and the uninflated price. A shill's goal is to lose each auction. A shill is not constrained by a budget, but rather a profit margin. If the shill wins, the item is resold in a subsequent auction. However, there is a limit on how many times this can be done. For each auction a shill wins, the seller incurs auction listing fees and is required to invest more time. Continual wins erode the profit from shilling on the item.

The shill faces a dilemma for each bid they submit. Increasing a bid could marginally increase the revenue for the seller. However, raising the price might also result in failure if it is not outbid before the auction terminates. The shill must decide whether to 'take the deal', or attempt to increase the pay-off.

On the contrary, a bidder's goal is to win. A bidder has a finite budget and is after the lowest price possible. Increasing a bid for a legitimate bidder decreases the money saved, but increases the likelihood of winning. The following outlines typical shill behaviour and characteristics:

- 1) A shill tends to bid exclusively in auctions only held by one (or a few) particular seller(s).
- 2) A shill generally has a high bid frequency. An aggressive shill will continually outbid legitimate bids to inflate the final price, until the seller's expected pay-off for shilling has been reached, or if the shill risks winning the auction (e.g., near the termination time or during slow bidding).
- 3) A shill has few or no winnings for the auctions participated in.
- 4) It is advantageous for a shill to bid within a small time period after a legitimate bid. Generally a shill wants to give legitimate bidders as much time as possible to submit a new bid before the auction's closing time.
- 5) A shill usually bids the minimum amount required to outbid a legitimate bidder. If the shill bids an amount that is much higher than the current highest bid, it is unlikely that a legitimate bidder will submit any more bids and the shill will win the auction.
- 6) A shill's goal is to try and stimulate bidding. As a result, a shill will tend to bid more near the auction's beginning. This means a shill can influence the entire auction process compared to a subset of it. Furthermore, bidding towards the auction's end is risky as the shill could accidentally win.

*Bid sniping* is a type of bidding behaviour that is typically deemed undesirable. A bidder using a sniping strategy will only bid in the auction's closing seconds. The goal is to prevent other bidders from outbidding the sniper's bid, as they do not have time to respond. Sniping behaviour is shilling's exact opposite. A sniper's goal is to win the auction for the lowest price, whereas a shill's goal is not to win and to create the highest price. Sniping is often used as a preventative measure against shilling. A sniper cannot prevent shilling occurring during an auction. However, the sniper can prevent itself from being shilled. Unlike shilling, sniping is permitted on eBay although its use is discouraged.

## III. A SHILL BIDDING AGENT

This section presents a shill bidding agent that uses bogus bids to inflate an auction's price. We describe the agent's goals and strategic directives. Each directive plugs into the agent interface, which dictates its bidding behaviour.

### A. Components of an English Auction

In order to participate in an auction, a bidder must *register*. They are provided with a unique bidder id,  $b_{id}$ , which they use to submit bids. During the *initialisation* stage, the Auctioneer sets up the auction and advertises it (i.e., item description, starting time, etc). An auction is given a unique number,  $a_{id}$ , for identification purposes. In the *bidding* stage, a bidder computes his/her bid and submits it to the Auctioneer. The

```

A - Directive D1
float D1(current price p) {
    return minimum price
}

B - Directive D3
boolean D3(risk limit θ) {
    calculate d
    calculate tS
    if (tc ≤ tS) then
        return true
    else
        return false
}

```

Fig. 1. Pseudocode for Shill Bidding Agent Directives 1 and 3

agent can place a bid in auction  $a_{id}$ , for price  $p'$ , by invoking the **submit bid**( $a_{id}$ ,  $p'$ ) function.

The Auctioneer must *supply intermediate information* to the agent pertinent to the auction's current state. The agent can request a price quote for a particular auction by invoking the **obtain price quote**( $a_{id}$ ) function. This includes the start, end and current time for the auction, and the starting bid (if one exists). It is assumed that the agent has access to the entire bid history up to the current time in the auction. The history can be considered as an ordered set  $\mathcal{H} = \{h_1, h_2, \dots, h_n\}$ ,  $|\mathcal{H}| = n$ , that contains price quote triples  $h_i = (time, price, b_{id})$ , where  $1 \leq i \leq n$ . The last element is the latest price quote for the auction (i.e.,  $h_n$  is the current highest bid).

Finally, during the *winner determination* stage, the Auctioneer chooses the winner according to the auction rules (e.g., who has the highest bid, whether the reserve has been met, etc.).

### B. Shill Agent Directives

The agent's goal is to maximise the profit from shilling, while avoiding winning the auction. A shill that wins the auction is deemed to have failed. We propose a set of directives that a shill must adhere to. If these directives are satisfied, the agent submits a bid. Each directive is described in turn.

**D<sub>1</sub> - Bid minimum amount required.** As part of the price quote, the Auctioneer also provides the minimum amount required to outbid the highest bid. This is usually calculated as a percentage of the current high bid, or determined according to a scalable amount depending on the current high bid.  $D_1(p)$  is a function that takes the current price,  $p$ , and returns the minimum amount the shill should bid. (See Figure 1A.)

**D<sub>2</sub> - Bid quickly after a rival bid.** The agent must bid immediately in order to influence the other bidders for the maximum time. The agent's location affects its response speed to rival bids. There are two main placement options:

- *On the Auctioneer.* The agent can respond instantly to a rival bid;
- *As a client of the Auctioneer.* The agent must periodically poll the Auctioneer to check if a new rival bid has been submitted. The length of the polling interval limits shilling.

The first approach uses the Auctioneer's computer to run the agent program. For example, eBay's proxy bidding system functions in this manner. However, this places a drain on the Auctioneer's computational and storage resources. As in TAC, the second approach allows bidders to host the agent on their own computer and interact with the Auctioneer via an application programming interface. This distributes the computational burden among the bidders. However, this requires a permanent connection to the Auction server and results in communication overhead. Furthermore, network delays and security threats can influence the agent's operating speed and integrity. RAS facilitates both types of agent placement.

**D<sub>3</sub> - Don't bid too close to the auction's end.** If the shill bids too close to the auction's end, it risks winning. To avoid this, the agent has a risk limit,  $\theta$ ,  $0 \leq \theta \leq 1$ . The agent is prohibited from bidding if the auction is more than  $(\theta \times 100)\%$  complete. We refer to this as *shill time limit*, and denote it as  $t_S$ .  $t_S$  is the absolute bound after which a shill is prevented from bidding. Let  $t_0$  and  $t_e$  be the starting and ending times respectively for an auction. The auction's duration,  $d$ , is calculated as  $d = t_e - t_0$ . The shill time limit,  $t_S$  is calculated as  $t_S = \theta d$ . Let  $t_c$  be the current time in an auction. The agent can only submit bids when  $t_c \leq t_S$ . Larger values of  $\theta$  increase the risk that a shill might win an auction.  $D_3(\theta)$  is a function that takes the risk limit  $\theta$ , and returns true or false regarding whether the agent should continue to bid. (See Figure 1B.)

**D<sub>4</sub> - Bid until the target price has been reached.** The reserve price can influence a bidder's strategy. It has been argued that auctions with lower reserve prices attract more bids, whereas bidders are deterred by high reserve prices. There are three factors that influence an auction's final price: the *reserve price*,  $r$ ; the *seller's true valuation*,  $sv$ ; and a *bidder's true valuation*,  $bv$ .

If the seller lists  $r$  below  $sv$ , s/he risks selling the item at a loss. This is shown in Figure 2A. If the seller lists  $r$  above  $sv$ , this may potentially increase the seller's profit (see Figure 2B). However, this may also deter bidders if it is above  $bv$  (see Figure 2C), in which case the item is not sold, or if there is no reserve, it is sold for a loss. Strategies regarding the choice of reserve price, touches on the complex topic of




Fig. 2. Pricing and Strategies

reserve price shilling, which is beyond the scope of this paper. For simplicity, we assume that the seller's best strategy is to choose  $sv = r$  (see Figure 2D). In this case, it is clear to see that profit occurs where  $bv > r$  (see Figure 2E).

The shill agent's primary goal is to inflate the price to at least  $r$ , (and by assumption  $sv$ ). Its secondary goal is to etch out further profit in the case where  $bv > r$ . In order to know when to stop bidding, the shill is given a *target price*, which is denoted by  $\alpha$ . The higher  $\alpha$  is, the more profit the shill can attain. However, this increases the chances of failure (i.e., winning the auction). Ideally, the choice of  $\alpha$  should be  $r \leq \alpha \leq bv$  (see Figure 2F). When choosing  $\alpha$ , the shill must ensure that it is greater than or equal to  $r$ . Otherwise, the item might be sold at a loss. If  $bv = r$ , then the seller can be sure that shilling will inflate the price to at least to  $r$  (see Figure 2G).  $bv$  must be greater than  $r$ , for the shill agent to be able to etch out profit. Profit in this case is measured as the difference between the inflated price and  $r$  (see Figure 2H).

Dumas *et al* [2] performed an analysis of datasets from eBay and Yahoo<sup>5</sup> and showed that the final prices of a set of auctions for a given item are likely to follow a normal distribution. This is due to a given item having a more or less well-known value, around which most of the auctions should finish. Figure 3 illustrates a typical closing distribution.  $bv$  can typically be deemed to be at the distribution's centre. However, this should only be used as a short to medium term indicator, as permanent price changes over time do occur.

$D_4(p, \alpha)$  is a function that takes the current price  $p$ , the shill's target price  $\alpha$ , and returns true or false regarding

<sup>5</sup>www.yahoo.com/auction




Fig. 3. Final price bid distribution over a series of relatively concurrent auctions

whether the agent should continue to bid. The agent will only bid when the current price  $p$ , is less than or equal to  $\alpha$ . (See Figure 4A.)

**$D_5$  - Only bid when the current bidding volume is high.**

The agent should preferably bid more towards the auction's beginning and slow down towards the shill time limit, unless the bidding activity is high. That is, the bidding volume must increase throughout the auction for the shill to maintain the same bid frequency. The agent uses the bid history  $\mathcal{H}$ , to analyse the current bid volume and decide whether to submit a bid. The agent observes the previous number of bids for a time interval. If the number of bids for the period is below a threshold, then the agent does not submit a bid.

First, we must determine each bid's normalised time in terms of the current time,  $t_c$ . This is represented as  $\delta h_i$ ,  $1 \leq i \leq n$ . Let  $t_i$  be the time for  $h_i$ .  $\delta h_i = \frac{t_i - t_0}{t_c - t_0}$ , where  $0 \leq \delta h_i \leq 1$ . The risk value  $\mu$ ,  $0 \leq \mu \leq 1$ , represents how far back in history from the current time that the agent will observe. For example, if  $\mu = 0.2$ , then the agent will only look for bids that were submitted in the final 20% of the normalised time period of the auction thus far (i.e.,  $\delta h_i \geq (1 - \mu)$ ).  $\kappa$  denotes the number of bids submitted in the last  $(\mu \times 100)\%$  of elapsed time from  $t_c$ . Increasing  $\mu$ , increases the level of risk for the shill (as the agent is influenced by bids further in the past).  $\kappa$  is calculated as follows:

$$\kappa = \sum_{i=1}^{|\mathcal{H}|} j \text{ where } j = \begin{cases} 1, & \text{if } \delta h_i \geq (1 - \mu) \\ 0, & \text{otherwise} \end{cases}$$

where  $0 \leq \kappa \leq |\mathcal{H}|$ .

Next,  $\kappa$  must be weighted depending on the current time in relation to the shill time limit,  $t_S$ . An increase in the trading volume is required towards  $t_S$ , in order for the agent to continue bidding at the same rate it did earlier in the auction. The normalised current time  $\delta t_c$ , in relation to  $t_S$  is calculated as  $\delta t_c = \frac{t_c - t_0}{t_S - t_0}$ , where  $0 \leq \delta t_c \leq 1$ . If  $\delta t_c < 0.5$ , then the agent will only require one bid to be submitted by a rival before it bids. When  $\delta t_c \geq 0.5$ , the agent requires at least two bids before it will submit a bid. This ensures that later in the auction, the agent will only respond to more aggressive rival behaviour.

When no bids have been submitted for an auction, the shill agent will attempt to stimulate bidding by submitting the first

```

B - Directive D5
boolean D5(risk limit  $\mu$ ) {
    if ( $|\mathcal{H}| = 0$ ) then
        return true
    calculate  $\delta h_i$  for all  $i, 1 \leq i \leq n$ 
    calculate  $\kappa$ 
    calculate  $\delta t_c$ 
    if ( $\kappa > 0$ ) then
        if ( $\delta t_c < 0.5$ )
            return true
        else if ( $\kappa > 1$ ) then
            return true
    return false
}

A - Directive D4
boolean D4(current price  $p$ , target price  $\alpha$ ) {
    if ( $p \leq \alpha$ ) then
        return true
    else
        return false
}

```

Fig. 4. Pseudocode for Shill Bidding Agent Directives 4 and 5

bid. This is a common practice in auctions. It is intended that psychologically the presence of an initial bid raises the item's worth, as competitors see that it is in demand.  $D_5(\mu)$  is a function that takes the risk limit  $\mu$ , and returns true or false regarding whether the agent should continue to bid. (See Figure 4B.)

The aforementioned directives govern the agent's operation depending on the state of the auction. The following pseudocode illustrates the agent's behaviour:

```

shill agent( $aid, \alpha, \theta, \mu$ ) {
    do {
        obtain price quote( $aid$ )
        if ( $D_3(\theta)$  AND  $D_4(p, \alpha)$  AND  $D_5(\mu)$ )
            submit bid( $aid, D_1(p)$ )
    } while ( $D_3(\theta)$  AND  $D_4(p, \alpha)$ )
}

```

The agent initially requests a price quote. If no bids have been submitted, then  $D_5$  returns true and the agent submits a bid for the amount returned by  $D_1$ . The agent then repeatedly requests price quotes to ensure that it is able to bid quickly if there is a rival bid (i.e.,  $D_2$ ). When a rival bid is submitted, the agent will bid only if the remaining directives,  $D_3$ ,  $D_4$ , and  $D_5$  are satisfied. The agent executes in this manner (requesting and evaluating price quotes), until either  $D_3$  or  $D_4$  becomes false.

#### IV. PERFORMANCE

The agent was implemented on RAS and has been tested with other types of bidding agents in a simulated auction market. The shill agent was assessed on its ability to inflate an auction's final price. The agent was considered *successful* if the winning price of a rival bid equals or exceeds  $\alpha$ . The agent was considered *unsuccessful* if it won the auction, or if it failed to inflate the final price to  $\alpha$  prior to ceasing bidding.

The shill agent was pitted against four Zero Intelligence agents (ZI) (see [4]). ZI agents are designed to simulate an ordinary bidder in an auction. Each ZI agent is assigned a random amount between \$0.05 and \$10.00 (according to a

uniform distribution), which it submits as a proxy bid at a random time during the auction. It was assumed that there was no reserve price.

Tests indicated that large numbers of ZI agents makes the shill less effective at influencing the final price (as there is no need to stimulate bidding). The average final price and bid volume for an auction **without shilling** were \$5.90 and 3.95 bids respectively. The standard deviation for closing prices was \$1.95.

**Claim 1** *The shill agent raised the average winning price compared to auctions it didn't participate in.* The average price was 1% – 25% higher in auctions where shilling occurred (depending on the shill's risk profile). Where  $\alpha < bv$ , the shill can influence the price the most. After this, the average final price becomes affected by the shill's winning bids. This can be seen in Figure 5 A.

**Claim 2** *The bidding volume increased in auctions where the shill agent was present.* Introducing a shill agent increased the average bid volume by up to 420%. Much of this can be attributed to the agent incrementally outbidding proxy bids. In general, the higher  $\alpha$ , the more bids a shill will have to submit. This can be seen in Figure 5 B.

**Claim 3** *Riskier shill agents acquire more profit, but at the expense of an increased number of failures.* We conducted numerous tests that altered the shill's risk parameters,  $\alpha$ ,  $\theta$  and  $\mu$ . More risk adverse shills tended to fail by not meeting  $\alpha$ , whereas riskier shills tended to fail by winning the auction. Figure 5 C shows how the agent's success rate decreases with increases in  $\alpha$ . Figure 5 D shows how the percentage of these failures that are due to winning the auction.

**Claim 4** *A shill agent that places a single proxy bid for the target price  $\alpha$ , is less effective.* Tests showed that shills employing this strategy tend to fail more by winning the auction. An agent employing the shilling strategy outlined in this paper uses  $\theta$  and  $\mu$  to help determine whether it is




Fig. 5. Graph illustrating the agent's performance with increasing risk factors  $\theta$  and  $\mu$ . The horizontal axis represents the target price  $\alpha$ . **A** shows the increase in average final price with increases in  $\alpha$ . Likewise **B** shows the increase in average number of bids per auction with increases in  $\alpha$ . **C** illustrates how the agent's success rate decreases with  $\alpha$  and **D** shows the increase in failures due to winning the auction.

safe to continue. Psychologically, the use of a larger number of small bids is also more likely to lure bidders into placing higher bids. Furthermore, it is less suspicious for the skill to slowly inflate the price, rather than enter an initial large amount.

**Claim 5** *The skill agent was less effective against sniping agents.* The skill agent was pitted against varying numbers of ZI agents following a sniping strategy (see [5]). As the number of sniping agents reached saturation (i.e., one skill vs all sniping agents), the skill's ability to influence the price decreased to 0. Thus the skill failure rate was 100% due to  $\alpha$  not being met (unless normal bidding inadvertently reached  $\alpha$ ). Bid volume was also significantly lower.

## V. CONCLUSIONS

This paper presents a malicious software bidding agent that follows a skill bidding strategy. The agent incrementally increases an auction's price forcing legitimate bidders to submit higher bids in order to win the item. The agent ceases bidding when the desired profit from shilling has been attained, or in the case that it is too risky to continue bidding without winning the auction.

The skill bidding agent has been implemented on RAS to aid us in developing skill detection techniques. The agent's ability to shill was tested using a simulated auction market involving other bidding agents. The agent raised the overall bidding volume as well as increased the average final price across all auctions. Shills with a higher profit risk factor managed to acquire a larger amount of profit. However, this came at the expense of an increased number of failed auctions. Tests conducted showed that the shilling strategy employed by the

agent, is superior to placing a single proxy bid for the target price. Engaging in sniping reduces the effectiveness of a skill bidder.

In an extended version of this paper, we have developed an adaptive skill bidding agent. When used on a series of auctions for substitutable items, the agent is able to revise  $\alpha$  and other risk factors based on its past performance. In future work we plan to investigate skill agents which engage in collusive behaviour to avoid detection.

## REFERENCES

- [1] Cliff, D., Minimal-intelligence agents for bargaining behaviours in market-based environments, *Hewlett Packard Labs, Technical Report HPL-97-91*, 1997.
- [2] Dumas, M., Aldred, L. and Governatori, G., A probabilistic approach to automated bidding in alternative auctions, In *Proceedings of the eleventh international conference on World Wide Web*, 99–108. ACM Press, 2002.
- [3] Gjerstad, S. and Dickhaut, J., Price formation in double auctions, *Games and Economic Behavior*, 22, 1–29, 1998.
- [4] Gode, D. and Sunder, S., Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality, *Journal of Political Economy*, 101, 119–137, 1993.
- [5] Rust, J., Miller, J. and Palmer, R., Behaviour of trading automata in a computerized double auction market, In *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, 1992.
- [6] Schwartz, J. and Dobrzynski, J., 3 men are charged with fraud in 1 100 art auctions on eBay, *The New York Times*, 2002.
- [7] Trevathan, J. and Read, W., RAS: a system for supporting research in online auctions, *ACM Crossroads*, 12.4, 23–30, 2006.
- [8] Trevathan, J. and Read, W., Undesirable and fraudulent behaviour in online auctions, In *Security and Cryptography Conference (SECURITY)*, 450–458, 2006.
- [9] Trevathan, J. and Read, W., Detecting shill bidding in online English auctions, *Technical Report*, 2006. Available at: <http://auction.maths.jcu.edu.au/research/shill.pdf>
- [10] Wellman, M. and Wurman, P., The 2001 trading agent competition, *Electronic Markets*, 13.1, 4–12, 2003.