

ONLINE AUCTION DECEPTION USING A SHILL BIDDING AGENT

JARROD TREVATHAN¹, WAYNE READ,¹ and RAJNI GOEL²

¹*James Cook University, Australia, jarrod.trevathan / wayne.read@jcu.edu.au*

²*Howard University, USA, rgoel@howard.edu*

Abstract:

Shill bidding is the act of using bids in an online auction to drive up the final price for the seller, thereby defrauding legitimate bidders. While ‘shilling’ is recognized as a problem and shill bidding is strictly forbidden in online auctions, presently there is little to no established means of defense against shills. This paper presents a software bidding agent that follows a shill 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. Its ability to inflate the price has been tested in a simulated marketplace and experimental results are presented. Furthermore, the agent is used to assist in developing algorithms to detect the presence of shill bidding in online auctions.

Keywords: Fraud, price inflation, software bidding agents, graph algorithms.

1. Introduction

Online auction fraud is rampant and every year it 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. The seller can either register as a bidder under a false identity, or be in collusion with one or more of the bidders. 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 it forces them to bid against false bids hence paying significantly more in order to win.

Automated agents for conducting electronic commerce are becoming more common to the extent that an idea has been touted in which all human input in auctions will eventually be performed by autonomous bidding agents. However, such an environment would definitely spawn undesirable behavior (e.g., cheating, stealing payments, etc. [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.

There are many auction types (e.g., Vickrey, Dutch, etc.), with the most popular being the English auction. In an English auction, bidders outbid each other in an attempt to win an item by being the highest bidder. English auctions are employed in online auctions such as those offered by eBay^a and Ubid^b, and these auctions are particularly susceptible to shill bidding practices. Shill bidding is strictly forbidden by commercial online auctioneers, and is a prosecutable offence [6]. However, the online environment makes shilling easy as bidders are anonymous, creating a need to develop schemes to detect the presence of shilling during an online English auction.

To develop a method to detect shill bidding behavior, we constructed a program 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. The shill bidding agent operates in a similar manner, however, it is working to inflate the auction's price and lose (thereby forcing the winner to pay more), rather than to trying to win the auction for the lowest price.

This paper presents a software bidding agent that follows a shill bidding strategy. To our knowledge no literature exists for a type of bidding agent we refer to as a *malicious bidding agent*. 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. We also show how the agent has been used to aid in the development and testing of shill detection techniques. 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 organized as follows: Section 2 describes general shill behavior. Section 3 presents a shill bidding agent that shills in a single auction. Section 4 evaluates the agent's ability to shill, contrasting safe and risky approaches. Section 5 discusses methods for

^a <http://www.ebay.com>

^b <http://www.ubid.com>

detecting shill bidding in online auctions and how the agent has been used to develop shill detection techniques. Section 6 provides some concluding remarks and avenues for future work.

2. Shill Behaviour

There is often much confusion regarding what constitutes shill behavior. Bidding behavior 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, making 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 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 un-inflated 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.

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 behavior and characteristics:

- A shill tends to bid exclusively in auctions only held by one particular seller, however, this alone is not sufficient to incriminate a bidder. It may be the case that the seller is the only supplier of an item the bidder is after, or that the bidder really trusts the seller (based on previous dealings).
- A shill tends to have a high bid frequency. An aggressive shill will continually outbid legitimate bids to inflate the final price. A shill typically will bid until the seller's expected payoff for shilling has been reached. Or until the shill risks winning the auction (e.g., near the termination time or during slow bidding).
- A shill has few or no winnings for the auctions participated in.
- 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 closing time of the auction.

- 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.
- It is common behavior for legitimate bidders trying to conserve their money to also only bid the minimal amount. In situations where a bidder's valuation is higher than the current bid, a bidder is more likely to outbid by more than the minimum amount.
- A shill's goal is to try and stimulate bidding. As a result, a shill will tend to bid more near the beginning of an auction. This means a shill can influence the entire auction process compared to a subset of it. Furthermore, bidding towards the end of an auction is risky as the shill could accidentally win.

The most extreme shill bidding strategy is referred to here as *aggressive shilling*. An aggressive shill continually outbids everyone thereby driving up the price as much as possible. This strategy often results in the shill entering many bids.

In contrast, a shill might only introduce an initial bid into an auction where there have been no prior bids with the intent to stimulate bidding. This kind of behavior is a common practice in both traditional and online auctions. However, most people typically do not consider it fraudulent. Nevertheless it is still shilling, as it is an attempt to influence the price by introducing spurious bids.

This is referred to as *benign shilling* in the sense that the shill does not continue to further inflate the price throughout the remainder of the auction. A benign shill will typically make a "one-off" bid at, or near the very beginning of the auction.

Regardless of the strategy employed, a shill will still be a bidder that often trades with a specific seller but has not won any auctions.

Another factor that affects a shill's strategy is the value of the current bid in relation to the reserve price. For example, once bidding has reached the reserve price it becomes more risky to continue shilling. This is conditional on whether the reserve is a realistic valuation of the item that all bidders share.

3. A Shill Bidding Agent

Numerous bidding agents have been proposed in literature [1, 3, 4, 5]. The Trading Agent Competition (TAC)^c [11] 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

^c <http://www.sics.se/tac/>

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.

Similar to a virus or worm, a malicious bidding agent is a bidding agent that behaves with the 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. 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 our own special purpose auction server. The auction server allows the agents to be tested in a controlled (and legal) manner (see Section 5.3 for further details).

In the following sections we develop and test 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 behavior.

3.1. 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 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 $H = \{h_1, h_2, \dots, h_n\}$, $|H| = n$, 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.).

3.2. 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.

```

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

B - Directive D3
boolean D3(risk limit  $\theta$ ) {
    calculate d
    calculate  $t_S$ 
    if ( $t_c \leq t_S$ ) then
        return true
    else
        return false
}

```

Figure 1 - Pseudocode for Shill Bidding Agent Directives 1 and 3

 D_1 - 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_2 - 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 periodically polls the Auctioneer to check if a new rival bid has been submitted. The polling interval length 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.

 D_3 - 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, θ , $0 \leq \theta \leq 1$. The agent is prohibited from bidding if the auction is more than $(\theta \times 100)\%$. 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 θ increase the risk that a shill might win an auction. $D_3(\theta)$ is a function that takes the risk limit θ , and returns **true** or **false** regarding whether the agent should continue to bid. (See Figure 1B.)

D_4 - Bid until the target price has been reached.

Many auctions specify a *reserve price* at the outset, which is the minimal winning amount that the seller will accept for the item. If the reserve price is not attained, then the auction is cancelled (i.e., the highest bidder is not required to pay, and the seller keeps the item). 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 .

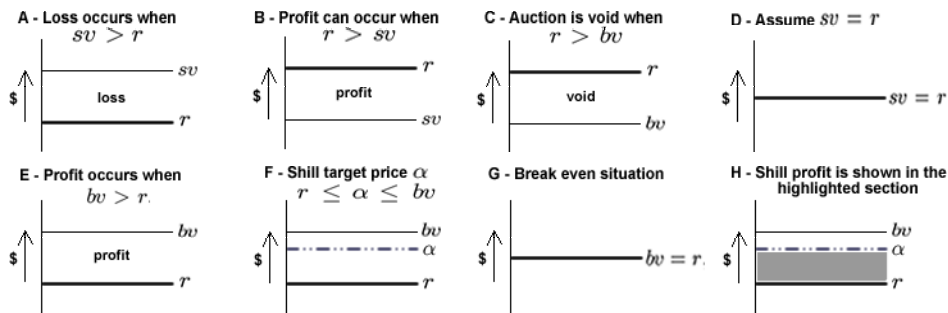


Figure 2 - Pricing and Strategies

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 touch on the complex topic of *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 α . The higher α is, the more profit the shill can attain. However, this increases the chances of failure (i.e., winning the auction). Ideally, the choice of α should be $r \leq \alpha \leq bv$ (see Figure 2F). When choosing α , 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).




Figure 3 - Final price bid distribution over a series of relatively concurrent auctions

Dumas *et al* [2] performed an analysis of datasets from eBay and Yahoo^d 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.

```

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
}

```

Figure 4 - Pseudocode for Shill Bidding Agent Directives 4 and 5

$D_4(p, \alpha)$ is a function that takes the current price p , the shill's target price α , and returns **true** or **false** regarding whether the agent should continue to bid. The agent will only bid when the current price p , is less than or equal to α . (See Figure 4A.)

^d www.yahoo.com/auction

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 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 (i.e., it returns **false** if the number of bids is below the threshold, or **true** otherwise).

First, we must determine each bid's normalised time in terms of the current time, t_c . This is represented as δh_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 μ , $0 \leq \mu \leq 1$, 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., $\delta h_i \geq (1 - \mu)$). κ denotes the number of bids submitted in the last $(\mu \times 100)\%$ of elapsed time from t_c . Increasing μ , increases the level of risk for the shill (as the agent is influenced by bids further in the past). κ is calculated as follows:

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

where $0 \leq \kappa \leq H$.

Next, κ 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 δt_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 behavior.

When no bids have been submitted for an auction, the shill agent will attempt to stimulate bidding by submitting the first 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 μ , and returns **true** or **false** regarding whether the agent should continue to bid. (See Figure 4B.)

- How the Agent Operates -

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

```

shill agent( $a_{id}$ ,  $\alpha$ ,  $\theta$ ,  $\mu$ ) {
  do {
    obtain price quote( $a_{id}$ )
  }
}

```

```

if (  $D_3(\theta)$  AND  $D_4(p, \alpha)$  AND  $D_5(\mu)$ 
    submit bid( $a_{id}$ ,  $D_1(p)$  )
} 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.

4. Performance

The agent was implemented 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 α . The agent was considered *unsuccessful* if it won the auction, or if it failed to inflate the final price to α prior to ceasing bidding.

The shill agent was pitted against four Zero Intelligence (ZI) agents (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 5A.

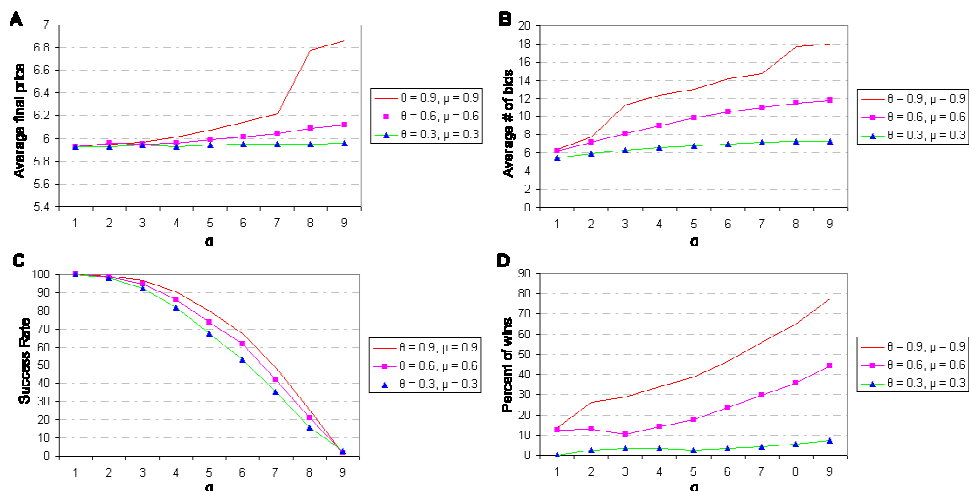


Figure 5 - Graphs illustrating the agent's performance with increasing risk factors θ and μ . A shows the increase in average final price with increases in α . Likewise B shows the increase in average number of bids per auction with increases in α . C illustrates how the agent's success rate decreases with α and D shows the increase in failures due to winning the auction.

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 α , the more bids a shill will have to submit. This can be seen in Figure 5B.

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, α , θ and μ . More risk adverse shills tended to fail by not meeting α , whereas riskier shills tended to fail by winning the auction. Figure 5C shows how the agent's success rate decreases with increases in α . Figure 5D 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 α , 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 θ and μ to help determine whether it is 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 shill to slowly inflate the price, rather than enter an initial large amount.

Claim 5 *The shill agent was less effective against sniping agents.* The shill 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 shill vs. all sniping agents), the shill's ability to influence the price decreased to 0. Thus the shill failure rate was 100% due to α not being met (unless normal bidding inadvertently reached α). Bid volume was also significantly lower.

5. Detecting Shill Bidding

This section describes methods to detect shill bidding, our proposed detection approaches, and how the shill bidding agent described in this paper has been used to aid in the development and testing of our proposed shill detection techniques.

At present there is limited coverage on how to detect shill bidding. eBay has been involved in many legal disputes where bidders/sellers have been accused of shilling (see [6]). eBay's policy clearly outlines undesirable bidder behavior and the penalties for shill bidding. However, eBay does not state exactly what factors they use to determine whether shilling has occurred, nor how to detect which bidders are shills. Furthermore, there is no means of recourse for an innocent seller who has incorrectly been accused of facilitating shilling in his/her auctions.

Shah et al [12] use data mining techniques to produce evidence of shilling. Their work used data from approximately 12,000 commercial auctions looking for associations between bidders and sellers. Bidders (or groups of bidders) that participated frequently in auctions held by particular sellers were deemed suspect. However, the authors' state that their analysis is very limited in that it only looks for simple associations. They suggest that a much more thorough analysis must be performed using complex associations which consider a wider range of shill behavior. There are also companies who offer data mining techniques to detect fraud in online auctions, however, like eBay, these companies have not made their techniques public.

There are also other characteristics possessed by shills. However, the use of these as a means of detection is dubious. For example, when a shill accidentally wins, the seller usually re-sells the item in a subsequent auction. The most obvious solution for detecting shilling is to observe if the same item is resold in a subsequent auction by the seller. This requires assigning a unique *item id* to each item sold by the seller. If the same *item id* is detected twice (i.e., the item has been resold by the seller), this provides some indication that shilling is taking place.

The problem with this approach is that it is only relevant for situations where a shill accidentally wins the item. That is, shilling is not detected in the case where a legitimate bidder has won and paid an inflated price. Furthermore, the seller might attempt to re-list the

item under a new *item id*. This then allows them to re-sell the item without raising suspicion. For these reasons, tracking resold items is not a sound method for detecting shilling.

Another commonly used detection method exploits the property that shills tend to be within close geographical proximity to each other. There are two main reasons for this. Firstly, collusion among bidders typically occurs amongst friends who live/operate near each other. This is because the costs of communication and coordinating the shilling process are less than over long distances. Secondly, if the seller has several aliases, these aliases will typically be registered in the same location.

However, geographical proximity alone is an unreliable indicator of shill behavior. First of all, the decreasing costs of communication hardware have made coordinating shill bidding over long distances much more affordable. Secondly, geographical information cannot be deduced from the raw auction data (as can the other characteristics) and involves cross-referencing users with the registration database, which raises anonymity concerns.

One proposed shill detection method examines the source IP address of a bidder. This is to prevent a shill from bidding using the same computer by logging in under different aliases. However, this is unreliable if more than one legitimate bidder must use the computer to bid (e.g., a computer in an Internet Cafe or a public library). Furthermore, this raises definite anonymity and privacy issues with regard to the bidders.

Other detection techniques suggest that seller feed back ratings and other statistics provided by eBay's feedback system can be used to detect fraud. The feedback system allows buyers and sellers to report on their dealings with each other, which becomes an indicator of an individual's honesty and reliability. This technique looks for anomalies in the feedback data to determine whether a seller/bidder fits a fraudster's profile. However, the problem with this approach is that it is specific to eBay, and is based solely on the integrity of the feedback system. Literature has well documented the short falls of eBay's feedback system. Among these shortfalls include that it is easy to generate false feed back, and that 99% of sellers seem to have feedback ratings greater than 97%, which suggests that the feedback rating is largely worthless as an indicator of a seller's integrity. Furthermore, the feedback system doesn't really encompass shill bidding, but rather just a seller's ability to dutifully deliver the item, and a bidder's reliability in making payment.

Existing shill detection methods are not public. It can be argued that it is dangerous to give a shill bidder knowledge of the inner workings of the detection method as this might allow them to subvert the system. However, the security (or ability) of the system should not rest with keeping it secret. In cryptography it is usually the case that confidence in a proposed security technique is only acquired over time if it withstands public scrutiny and attacks against its security. Likewise, if an adversary is allowed to attack the shill detection methods, then the confidence in a scheme will grow and ultimately better detection techniques can evolve.

5.1. Detecting a Shill Bidder Using Historical Auction Data

There are multiple scenarios that a seller and/or shill bidder can engage in to inflate an auction's price. This section provides an overview of our proposed approach to detect **one** shill who is working for **one** specific seller. [10] outlines more elaborate scenarios.

Our proposed solution observes bidding patterns over a series of auctions for a particular seller, looking for the shilling behaviour outlined in Section 2. The goal is to obtain statistics regarding a bidder's conduct, and deduce a measure called a *Shill Score*, that indicates the likelihood that s/he is engaging in shill behaviour.

The Shill Score algorithm targets core shilling strategies. A shill that deviates too far from these characteristics is less effective, and won't significantly alter the auction outcome. This approach acts as both a detection mechanism and a deterrent to shill bidders. To avoid detection, a shill must behave like a normal bidder, which essentially stops them shilling.

The Shill Score algorithm basically works as follows (see Trevathan and Read [9] for further details): A bidder i , is examined over m auctions held by the same seller for the behaviour outlined in Section 2. Each characteristic of shill behaviour is given a rating, which is combined to form the bidder's Shill Score. The algorithm gives a bidder a value between 0 and 10. The closer the Shill Score is to 10, the more likely that the bidder is a shill. The algorithm's goal is to determine which bidder is most inclined to be the shill out of a group of n bidders. The shill behavioural ratings are calculated as follows:

- α Rating - Percentage of auctions (by a particular seller) bidder i has participated in.
- β Rating - Percentage of bids bidder i has made out of all the auctions participated in.
- γ Rating - Normalised function based on the auctions bidder i has won out of the auctions participated in.
- δ Rating - Normalised inter bid time for bidder i out of the auctions participated in.
- ε Rating - Normalised inter bid increment for bidder i out of the auctions participated in.
- ζ Rating - Normalised time bidder i commences bidding in an auction.

Each rating is between 0 and 1, where the higher the value, the more suspicious the bidder. A bidder's Shill Score (denoted as SS) is calculated as the weighted average of these ratings:

$$SS = \frac{\omega_1\alpha + \omega_2\beta + \omega_3\gamma + \omega_4\delta + \omega_5\varepsilon + \omega_6\zeta}{\omega_1 + \omega_2 + \omega_3 + \omega_4 + \omega_5 + \omega_6}$$

where ω_i , $1 \leq i \leq 6$, is the weight associated with each rating.

[9] provides details and justifications for selection of weight values. If a bidder wins an auction, then his/her α , β , δ , ε and ζ ratings are 0 for the particular auction (as winning is inconsistent with a shill's goals). Note that the higher each individual rating, the higher the Shill Score (γ , δ and ε are inverted).

A bidder can observe other bidders' Shill Scores for an auction held by a particular seller. If the bidder feels that one (or more) of the other participating bidders' Shill Scores are too

high (i.e., there is potential that shilling is occurring), the bidder can then decide whether or not they want to risk participating in the auction (or other auctions held by the seller).

5.2. Using a Shill Bidding Agent to Test Shill Detection Techniques

The Research Auction Server⁶ at James Cook University, is an online server for conducting research into security issues regarding online auctions (see [7]). Both real and simulated auctions are performed to test the effectiveness of the detection methods. Some of these tests involve using human bidders who are bidding in both fake and real auctions (i.e., using fake and real money). One of the bidders is tasked with being a shill (note that this is illegal to do on commercial auction servers). The ‘human’ shill has been surprisingly effective in inflating the auction price. The majority of bidders are oblivious to the fact they are being shilled, and many of them actively engage in competitive rallies with the shill bidder causing the price to increase quickly. *(All victim bidders were reimbursed after the tests.)*

As it is laborious and time consuming to organise auctions involving human bidders and for the human shill bidder to actively observe auctions, the shill bidding agent described in this paper was implemented. The agent removed the need for a human to constantly observe an auction for opportunities to shill. Essentially a seller’s auctions became self-inflating. The power of the shill bidding agent also became more apparent when sophisticated colluding strategies were implemented. Essentially this reduced the costs and effort of coordinating a collusive shill bidding attack by a seller. The seller need not have friends bid in his/her auction, but only has to set up fake accounts (using false registration data). The colluding shill bidding agents then take care of the rest.

As described in Section 4, the shill agent was also pitted against other automated bidding agents. This allows tens of thousands of auction simulations to be performed that would not be practical with human bidders. These simulations can be used to quickly refine the shill detection techniques using a large amount of auction/bidding data.

6. Conclusions

This paper presents a malicious software bidding agent that follows a shill 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 shill bidding agent has been implemented on RAS to aid us in developing shill detection techniques. The agent’s ability to shill was tested using a simulated auction market

⁶ <http://auction.maths.jcu.edu.au>

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 shill bidder.

In future work we plan to investigate data mining techniques to try and determine if any complex associations exist between shill bidders that are not being taken into account by the Shill Score and the collusion detection mechanisms. Furthermore, the detection techniques are being expanded encompass cases involving multiple seller collusion. That is, multiple colluding sellers who control pools of colluding shill bidders.

References

1. D. Cliff, Minimal-intelligence Agents for Bargaining Behaviors in Market-based Environments, *Hewlett Packard Labs, Technical Report HPL-97-91*, (1997).
2. M. Dumas, L. Aldred and G. Governatori, A Probabilistic Approach to Automated Bidding in Alternative Auctions, In *Proceedings of the eleventh international conference on World Wide Web*, (ACM Press, 2002), pp. 99-108.
3. S. Gjerstad and J. Dickhaut, Price Formation in Double Auctions, *Games and Economic Behavior*, 22, (1998), pp. 1-29.
4. D. Gode and S. Sunder, Allocative Efficiency of Markets with Zero Intelligence Traders: Market as a Partial Substitute for Individual Rationality, *Journal of Political Economy*, 101, (1993), pp. 119-137.
5. J. Rust, J. Miller and R. Palmer, Behavior of Trading Automata in a Computerized Double Auction Market, In *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, (1992).
6. J. Schwartz and J. Dobrzynski, 3 men are charged with fraud in 1100 art auctions on eBay, *The New York Times*, (2002).
7. J. Trevathan and W. Read, RAS: a system for supporting research in online auctions, *ACM Crossroads*, 12.4, (2006), pp. 23-30.
8. J. Trevathan and W. Read, Undesirable and Fraudulent Behavior in Online Auctions, In *Security and Cryptography Conference (SECRYPT)*, (INSTICC Press, 2006), pp. 450-458.
9. J. Trevathan and W. Read, Detecting Shill Bidding in Online English Auctions, *Handbook of Research on Social and Organizational Liabilities in Information Security*, Chapter 27, (Information Science Reference, 2008), pp. 446-470.
10. J. Trevathan and W. Read, Investigating Shill Bidding Behaviours Involving Colluding Bidders", *Journal of Computers*, vol. 2, issue 10, (Academy Publisher, 2007), pp. 63-75.
11. M. Wellman and P. Wurman, The 2001 Trading Agent Competition, *Electronic Markets*, 13.1, (2003), pp. 4-12.
12. H. Shah, N. Joshi and P. Wurman, Mining eBay: Bidding Strategies and Shill Detection. *WEBKDD 2002 - MiningWeb Data for Discovering Usage Patterns and Profiles*, Lecture Notes in Computer Science. (Springer-Verlag, 2003), pp. 17-34.