

Investigating Shill Bidding Behaviour Involving Colluding Bidders

Jarrod Trevathan and Wayne Read

School of Mathematics, Physics and Information Technology, James Cook University, Australia

Email: {jarrod.trevathan, wayne.read}@jcu.edu.au

Abstract—Shill bidding is where spurious bids are introduced into an auction to drive up the final price for the seller, thereby defrauding legitimate bidders. Trevathan and Read presented an algorithm to detect the presence of shill bidding in online auctions. The algorithm observes bidding patterns over a series of auctions, and gives each bidder a *shill score* to indicate the likelihood that they are engaging in shill behaviour. While the algorithm is able to accurately identify those with suspicious behaviour, it is designed for the instance where there is only one shill bidder. However, there are situations where there may be two or more shill bidders working in collusion with each other. Colluding shill bidders are able to engage in more sophisticated strategies that are harder to detect. This paper proposes a method for detecting colluding shill bidders, which is referred to as the collusion score. The collusion score, either detects a colluding group, or forces the colluders to act individually like a single shill, in which case they are detected by the shill score algorithm. The collusion score has been tested on simulated auction data and is able to successfully identify colluding shill bidders.

Index Terms—Auction fraud, shill bidding, reputation system, colluding bidders, bidding agents, graph theory

I. INTRODUCTION

Online auctions are a popular means to exchange items. However, to the unsuspecting, the auction process is fraught with peril. Auction fraud is an ever-increasing problem that is magnified in the online environment. For example, a seller might accept payment and not deliver an item, or misrepresents the item to be of greater value. Furthermore, auction participants might engage in practices such as siphoning, sniping, bid rigging or shilling, that are designed to influence the auction in a manner that disadvantages others (see [6]).

Shill bidding is the act of introducing fake bids into an auction on the seller's behalf, in order to artificially inflate an 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 one of the shills accidentally wins, then the item is re-sold in a subsequent auction. Shill bidding is a problem as it forces legitimate bidders to pay significantly more.

The advent of online auctions such as eBay¹ and ubid² have made shill bidding much easier. This is due

to bidders not being physically present, which allows a shill to anonymously influence the bidding process. Furthermore, it is relatively simple for a seller to register under many aliases, and operate in rings with impunity. For example, in March 2001, a U.S. federal grand jury charged three men for their participation in a ring of fraudulent bidding in hundreds of art auctions on eBay (see Schwartz and Dobrzynski [3]). The men created more than 40 User IDs on eBay using false registration information.

Trevathan and Read [8] presented an algorithm (referred to as the Shill Score or SS algorithm) to detect the presence of shill bidding in online English auctions. The algorithm observes bidding patterns over a series of auctions, and gives each bidder a *shill score* to indicate the likelihood that they are engaging in shill behaviour. The algorithm is able to accurately identify those with suspicious behaviour in the case where there is only one shill bidder. However, as the example above illustrates, there are situations where multiple shills are in collusion with each other.

Collusive shill bidding refers to the instance where two or more shills work together. Collusion makes it more difficult to determine whether shilling is occurring in an auction, and which bidders are responsible for the shill bids. This is because colluding shills can distribute the work evenly among each other to collectively reduce their shill scores. For example, a group of shills can take turns at submitting bids, and/or alternate at participating in an auction. This has the effect of making an individual shill appear to be a regular bidder.

This paper proposes methods for detecting colluding shill bidders (extending upon our previous work [9]). Several *collusion ratings* are presented that indicate whether collusion is occurring, and which bidders are most likely to be in collusion with each other. The collusion ratings look for typical collusive behaviour among bidders, and form a bidder's *collusion score*. The group of shills is faced with a choice of either, a) acting as a group and being detected by the collusion score, or b) acting individually (like a single shill), and being detected by the SS algorithm. The collusion score has been tested on simulated auction data to gauge its effectiveness in detecting colluding shills.

This paper is organised as follows: Background on general shill behaviour is discussed in Section II. Section III describes shill detection methods and provides an

This paper is based on "Detecting Collusive Shill Bidding," by J. Trevathan and W. Read, which appeared in the Proceedings of the 4th International Conference on Information Technology - New Generations (ITNG), Las Vegas, USA, April 2007. © 2007 IEEE.

¹<http://www.ebay.com>

²<http://www.ubid.com>

overview of the SS algorithm. Section IV extends the SS algorithm to detect collusive shill behaviour, and identify which bidders are in collusion with each other. Section V describes how the proposed collusion score performs on simulated auction data with colluding shills. Section VI provides concluding remarks.

II. SHILL BEHAVIOUR

This section provides an insight into general shill behaviour. It describes a shill's mindset, characteristics and strategies, and presents an example of shill behaviour in an auction. There are other possible behaviours that can be considered as shilling. Furthermore, there is often much confusion over what constitutes shill behaviour. However, we believe that this section outlines the optimal or 'best' strategies for a shill to take in terms of obtaining the highest shilling profit, whilst avoiding winning an auction.

A. Shill Mindset

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.

Another of a shill's main goals 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 won, 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.

B. Shill Characteristics and Strategies

The following outlines typical shill characteristics assuming we are using an auction that terminates at a set expiration time (refer to [6], [8] for a more detailed explanation):

- 1) A shill usually bids exclusively in auctions only held by one particular seller.
- 2) A shill tends to have a high bid frequency. An aggressive shill will continually outbid legitimate bids inflating the final price until the seller's expected pay-off for shilling has been reached, or until the shill risks winning the auction (e.g., near the termination time or during slow bidding).
- 3) A shill has few or no winnings.

TABLE I.
AN EXAMPLE AUCTION WITH ONE SHILL

Bid #	b_{id}	Price	Time
15	b_1	\$ 33	20:03
14	b_2	\$ 32	12:44
13	b_1	\$ 31	12:42
12	b_2	\$ 26	5:05
11	b_1	\$ 25	5:02
10	b_2	\$ 21	2:47
9	b_3	\$ 20	2:45
8	b_2	\$ 15	1:07
7	b_1	\$ 14	1:05
6	b_2	\$ 9	0:47
5	b_3	\$ 8	0:45
4	b_2	\$ 6	0:20
3	b_3	\$ 5	0:19
2	b_2	\$ 2	0:06
1	b_1	\$ 1	0:05

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 closing time of the auction.

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 likely that legitimate bidders will be deterred from bidding.

6) 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.

We refer to the most extreme shill bidding strategy 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 has been no prior bids with the intent to stimulate bidding. This behaviour is a common practice in 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. We refer to this 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 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. A shill's strategy is also affected by 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.

C. Shill Example

Table I illustrates an example auction with three bidders. Each bidder is denoted as b_1 , b_2 and b_3 respectively. Bidders b_1 and b_3 are legitimate, whereas b_2 is a shill.

b_2 engages in aggressive shill behaviour by outbidding a legitimate bid by the minimal amount required to stay ahead and within a small time period of the last bid. b_2 's bids force the other bidders to enter higher bids in order to win. If b_2 was not participating in this auction, b_1 would have only needed to pay \$21 in order to win. Instead b_2 caused b_1 to pay \$33, thus the shill has inflated the price by \$12.

In this example, b_2 exhibits the typical shill behaviour described above. This is evidenced by: 1. High frequency of bids (i.e., b_2 has submitted more bids than both the other bidders); 2. Has not won the auction despite the high number of bids; 3. Quick to bid after a legitimate bidder; and 4. Only bids the minimal amount to stay in front.

III. SHILL DETECTION

Until recently, there was limited coverage on exactly how to detect shill bidding. A basic attempt was proposed by Shah *et al* [4] using data mining techniques. Rubin *et al* [2] use three broad metrics for measuring what they consider 'shill behaviour', to construct a reputation system. Cheng and Xu [1] recently showed how shill bidders can potentially be identified by the manner in which bidders behave in two concurrently running auctions for identical items. Trevathan and Read [8] propose a solution (i.e., the SS algorithm), that observes bidding patterns over a series of auctions for a particular seller, looking for the shilling behaviour outlined in Section II. 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 SS 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 SS algorithm basically works as follows (see Trevathan and Read [8] for further details): A bidder i , is examined over m auctions held by the same seller for the behaviour outlined in Section II. Each characteristic of shill behaviour is given a rating, which is combined to form the bidder's shill score. The shill score 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:

$$SS = \frac{\omega_1\alpha + \omega_2\beta + \omega_3\gamma + \omega_4\delta + \omega_5\epsilon + \omega_6\zeta}{\omega_1 + \omega_2 + \omega_3 + \omega_4 + \omega_5 + \omega_6} \times 10$$

where ω_i , $1 \leq i \leq 6$, is the weight associated with each rating. [8] 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).

IV. COLLUDING SHILLS

The SS algorithm considers only the basic scenario with **one shill**. In the case outlined in Section I, there were three sellers which used 40 different aliases (40 shills in effect). The sellers understood that there was less chance that they would get caught if they utilised multiple names to take alternating turns at shilling. This makes it more difficult for authorities to determine which bidders are shills, as collusive behaviour allows shills to appear as more regular bidders.

In terms of the SS algorithm, introducing multiple shills makes the task of detection much harder. This is because colluding shills can engage in more sophisticated strategies in an attempt to thwart the detection algorithm. This section discusses these strategies and presents methods by which the SS algorithm can be extended to detect colluding shills. We restrict our attention to the case where there is only **one** seller which controls multiple shills. That is, we are not investigating strategies involving multiple sellers.

In some cases, geographical proximity can be an indication of collusion if there are several shills within a close area that participate in the auction. For example, in the shill case, two of the men were from California and the other was from Colorado. However, this is not a reliable indicator of shilling and may raise privacy concerns, as it requires examining the registration database for such relationships. Checking suspect colluders' geographical proximity would generally only occur once the SS-collusion algorithm has been run, thus ensuring there is strong evidence for doing so.

As stated previously, the main goal of shilling is to *drive up the price of an item*. In the situation where there is only one shill, the shill's secondary goal is to attempt to do this in such a manner that it *minimises his/her shill score*. When there is more than one shill, there are particular strategies that the group (of shills) can engage in to influence some factors contributing to their individual shill scores. Therefore the group's collective

TABLE II.
AN EXAMPLE AUCTION WITH TWO SHILLS ALTERNATING BIDS

Bid #	b_{id}	Price	Time
1	b_1	1	0:05
2	b_2	2	0:06
3	b_1	5	0:19
4	b_3	6	0:20
5	b_1	8	0:45
6	b_2	9	0:47
7	b_1	14	1:05
8	b_3	15	1:07
9	b_1	20	2:45
10	b_2	21	2:47
11	b_1	25	5:02
12	b_3	26	5:05
13	b_1	31	12:42
14	b_2	32	12:44
15	b_1	35	20:03

goal (secondary to shilling) is to *minimise its member's skill scores*.

Despite being able to use more complicated strategies, the group as a whole must still conform to certain behaviour in order to be effective as a skill. With regard to the SS algorithm, all that skills can do by colluding is to reduce their α , β and ζ ratings. The γ , δ and ϵ ratings are still indicative of skill bidding. For example, none of the colluding skills will ever win an auction. Furthermore, it is still in the group's interests to bid quickly, and by minimal amounts to influence the selling price. Therefore, inter bid times and increments will be consistent for skills.

There appear to be three possible strategies that can be employed by colluding skills. The following example illustrates the first strategy. Here there are two colluding skills (s_1 and s_2). Each skill takes alternating turns at bidding, i.e., s_i bids, then s_2 bids, then s_1 bids again, etc. This behaviour is shown in Table II. Here there are three bidders denoted b_1 , b_2 and b_3 respectively. b_1 is a legitimate bidder, but b_2 and b_3 are skills. b_2 and b_3 take alternating turns at outbidding b_1 .

This strategy has the effect of lowering b_2 and b_3 's β ratings (i.e., the number of individual skill bids in an auction). We refer to this strategy as the *alternating bid strategy*.

The second strategy is for colluding skills to take turns at shilling for a particular auction. For example, given two auctions, one skill will bid exclusively in the first auction, while the other skill bids only in the second auction. This strategy lowers the skills' α ratings (i.e., number of auctions participated in), but does not affect their β ratings. We refer to this strategy as the *alternating auction strategy*.

The third strategy is to use a combination of the alternating bid and alternating auction strategies. This can be used to alter the group's α and β ratings between the two extremes. We refer to this strategy as the *hybrid strategy*. An example of a hybrid strategy would be for skills s_1 and s_2 to alternately bid in auction₁, s_3 and s_4 alternately bid in auction₂, then s_1 and s_3 alternately bid in auction₃, etc. This continues until all combinations of bidders have been used, and then the process repeats. In reality, colluding skills would probably employ a hybrid strategy.

TABLE III.
AUCTION WITH TWO BIDDERS - ONE BIDDER (b_1), ONE SHILL (s_1)

n_j	b_1	s_1	bid sequence	b_1	s_1
3	2	1	b_1, s_1, b_1	$\lceil n_j/2 \rceil$	$\lfloor n_j/2 \rfloor$
4	2	2	s_1, b_1, s_1, b_1	$n_j/2$	$n_j/2$
5	3	2	b_1, s_1, b_1, s_1, b_1	$\lceil n_j/2 \rceil$	$\lfloor n_j/2 \rfloor$
6	3	3	$s_1, b_1, s_1, b_1, s_1, b_1$	$n_j/2$	$n_j/2$

The following notation is used throughout this paper:

Let $L = \{1, \dots, \ell\}$ be a set of bidder numbers; $|L| = \ell$

Let $M = \{1, \dots, m\}$ be a set of auction numbers; $|M| = m$

Let $\mathcal{B} = \{b_1, \dots, b_\ell\}$ be a set of bidders and $\mathcal{S} = \{s_1, \dots, s_{\ell'}\}$, $\mathcal{S} \subset \mathcal{B}$, be the set of skills, where $\ell' < \ell$.

A. Alternating Bid Strategy

If the alternating bid strategy is employed, it is in the group's best interests to evenly alternate. If a particular skill bids more than other skills, this will increase that individual's β rating and hence their skill score. This violates the group's collective goal of minimising all its member's scores. There are rules that govern the number of bids a group of skills can submit during an auction. We will examine some of the possible scenarios.

Table III illustrates a case where there are two bidders and one of them is a skill. b_1 must bid $\lceil n_j/2 \rceil$ times to win. The skill can submit at most $\lfloor n_j/2 \rfloor$ bids, otherwise they will win the auction (and fail as a skill).

Table IV illustrates a case where there are three bidders and one of them is a skill. Consider the worst case scenario for the skill in terms of the SS algorithm where they must bid after every legitimate bid. This situation results in the maximum β rating a bidder could possibly get. As in the previous example, the skill can submit at most $\lfloor n_j/2 \rfloor$ bids. In reality, a skill will not normally need to bid this many times, as there are would be situations where two or more legitimate bidders might outbid another bid before a skill has time to respond. However, we must remember that is the worst case scenario (i.e., the skill must bid after every legitimate bid). In contrast, the best case scenario would be where the skill wouldn't have to bid at all, because the legitimate bidders had pushed the price beyond the seller's expected pay-off for shilling.

Given the worst case scenario with n legitimate bidders and one skill, regardless of which bidder bids, collectively the legitimate bidders must bid $\lceil n_j/2 \rceil$ times to win. This is shown in Table IV, where bidders b_1 and b_2 combined must outbid the skill, s_1 .

In terms of colluding skills, increasing the number of skills decreases the number of times each skill must bid

TABLE IV.
AUCTION WITH THREE BIDDERS - TWO BIDDERS (b_1, b_2), ONE SHILL (s_1)

n_j	b_1	b_2	s_1	bid sequence	b_1	b_2	s_1
3	1	1	1	b_1, s_1, b_2	$\lceil n_j/2 \rceil$	$\lfloor n_j/2 \rfloor$	
4	1	1	2	s_1, b_1, s_1, b_2	$n_j/2$	$n_j/2$	
5	2	1	2	b_1, s_1, b_2, s_1, b_1	$\lceil n_j/2 \rceil$	$\lfloor n_j/2 \rfloor$	
6	2	1	3	$s_1, b_1, s_1, b_2, s_1, b_1$	$n_j/2$	$n_j/2$	

TABLE V.
AUCTION WITH THREE BIDDERS - ONE BIDDER (b_1), TWO SKILLS
(s_1, s_2)

n_j	b_1	s_1	s_2	bid sequence	b_1	s_1	s_2
4	2	1	1	s_1, b_1, s_2, b_1	$n_j/2$		$n_j/2$
5	3	1	1	b_1, s_1, b_1, s_2, b_1	$\lceil n_j/2 \rceil$		$\lceil n_j/2 \rceil$
6	3	2	1	$s_1, b_1, s_2, b_1, s_1, b_1$	$n_j/2$		$n_j/2$
7	4	2	1	$b_1, s_1, b_1, s_2, b_1, s_1, b_1$	$\lceil n_j/2 \rceil$		$\lceil n_j/2 \rceil$
8	4	2	2	$b_1, s_1, b_1, s_2, b_1, s_1, b_1, s_2$	$n_j/2$		$n_j/2$

(thereby decreasing an individual skill's β rating). This alleviates the worst case scenario for shilling. Given ℓ' skills (where $\ell' < \ell$), the group can split the blame by dividing the number of bids submitted amongst themselves. An individual skill can submit as little as $\lfloor n_j/2 \rfloor / \ell'$ bids in an auction. This is shown in Table V with one bidder and two skills.

In the case of *maximum collusion* (i.e., all bidders are skills except for one, $\ell' = \ell - 1$), to win the legitimate bidder is forced into submitting the most bids for an individual in the auction (i.e., $\lceil n_j/2 \rceil$ bids). If the legitimate bidder wins, then his/her β rating will be 0 for the auction (i.e., the SS algorithm does not penalise winning bidders as winning is inconsistent with a skill's goals). However if the bidder does not win, his/her β rating will be significantly higher than all the other bidders (i.e., the skills), as a result of having to compete with the skill bids. The later scenario becomes common as more legitimate bidders are added. In this situation a bidder's γ rating (number of wins), will be a significant factor in determining whether they really are a skill. This example shows how colluding skills can lower their individual skill scores while shifting attention to innocent bidders.

The major question maximum collusion poses is, "how do we detect and deter collusion among skill bidders to ensure that the SS algorithm does not incriminate a legitimate bidder?" We propose a solution to the problem, which we refer to as the *collusion graph*.

- Collusion Graph -

The collusion graph indicates which bidders are likely to be in collusion with each other. There are two different forms of the collusion graph based on whether skills use the alternating bid, or the alternating auction strategy. Two new collusion ratings are introduced, η and θ , that measures a bidder's conduct in terms of collusive behaviour. These ratings **are not** used in the calculation of the skill score. Instead, the ratings supplement the skill score, and serve to bring possible colluders to the attention of the SS algorithm. We will discuss each rating in turn.

To detect colluding groups employing the alternating bid strategy, bidders are represented as a graph $G = (V, E)$. V is the set of bidders, and E is an edge between two bidders, indicating that they have both participated in the same auction. The goal of the collusion graph is to find a subset, $C \subset V$, which contains the bidders that are most likely to be in collusion with each other.

A bidder initially has no edges connecting it to other bidders (i.e., the set E is empty). Given two bidders,

$v_i, v_j \in V, i \neq j$, that participate in the same auction, an edge $e_{i,j}$ is added to E that connects these two bidders together. If bidders v_i, v_j participate in more than one auction simultaneously, weights are added to the edge $e_{i,j}$ to indicate the number of auctions they were present in together.

The idea is that G will form a graph connecting colluding bidders together. The higher the edge weighting between two bidders, the greater the likelihood that the two bidders are in collusion with each other. Figure 1 gives an example of a collusion graph. Here there are two colluding bidders. In general, skills will have the most number of edges (i.e., highest degree), and higher edge weightings than legitimate bidders.




Figure 1. Example Collusion Graph

Given a node with degree k , each edge weight is denoted as $w_j, 1 \leq j \leq \ell$. The base collusion rating, η'_i , for a bidder i is calculated as the sum of the edge weights incident to the node:

$$\eta'_i = \sum_j^k w_j$$

To calculate bidder i 's normalised collusion rating η_i , we first need to find the maximum and minimum base collusion ratings for the graph. These are denoted as η^{max} and η^{min} respectively. The normalised collusion rating is calculated as:

$$\eta_i = \frac{\eta'_i - \eta^{min}}{\eta^{max} - \eta^{min}}$$

where $0 \leq \eta_i \leq 1$.

Colluding skills will have similar η collusion ratings. This value is initially used to bring suspected colluders to the attention of the SS algorithm, (but is not used in the calculation of the skill score). Given a bidder with a high collusion rating, we can then observe his/her γ, δ and ϵ ratings, which will provide a strong indication whether the bidder is engaging in collusive behaviour. Consistent features for the group of skills will be their γ, δ and ϵ ratings.




Figure 2. Potential Colluding Bidders

Figure 2 shows the previous example of the collusion graph after the algorithm has been run. The bold nodes

indicate the suspect bidders. The η ratings for these bidders is 1. The next highest collusion rating for a bidder is 0.27, which is significantly lower than the three suspected colluders. This example shows that the collusion graph is able to immediately single out suspect bidders. These bidders can then be further investigated based on their γ , δ and ϵ ratings.




Figure 3. Example Collusion Graph

Figure 3 illustrates an example collusion graph with two skills and three bidders. In this example there are three auctions. Each bidder participates in exactly one auction, whereas the skills participate in all three auctions. The bidding sequences for each auction are $(s_1, b_1, s_2, b_1, s_1, b_1)$, (s_2, b_2, s_1, b_2) and $(s_2, b_3, s_1, b_3, s_2, b_3)$ respectively. The right side of Figure 3 shows that after the algorithm has been run, the colluding skills strongly emerge as the most suspicious.

The collusion graph can be represented as an adjacency matrix. Each entry in the matrix lists the edge weighting for a pair of bidders. The collusion graph from Figure 3 gives the following adjacency matrix:

	s_1	s_2	b_1	b_2	b_3
s_1	—	3	1	1	1
s_2	3	—	1	1	1
b_1	1	1	—	0	0
b_2	1	1	0	—	0
b_3	1	1	0	0	—

Base collusion ratings for each bidder can be obtained by summing the entries in a bidder's column.

Colluding skills using the alternating bid strategy will have approximately the same α rating. The collusion graph reflects this in the edge weights between bidders. Colluding skills will also have participated with a similar number of legitimate bidders. The collusion graph reflects this by the degree of a bidder's node. Colluding skills will have approximately the same η rating.

However, in this form the η rating does not significantly bind groups of bidders together. There are a lot of innocent bidders that are incorporated into the collusion graph by chance. These bidders will tend to have the lower η ratings. Less significant collusion structures can be pruned from the collusion graph, by using the observation that colluding skills using the alternating bid strategy will have approximately the same β rating.

Given two suspect colluding bidders, we can further compare them based on their β ratings. We refer to this as a *binding factor*. The binding factor, $\phi_{i,j}^\beta$, gives bidders i and j a value between 0 and 1, based on how similar their β ratings are. A binding value of 1 indicates a strong binding (high likelihood of collusion), whereas a binding value of 0 indicates a weak binding.

Let β_i and β_j denote the β ratings for bidders i and j respectively. If $\beta_i = \beta_j$, then the binding factor, $\phi_{i,j}^\beta$, is

1. If this is not the case, then the binding factor is given a decreasing value based on the percentage difference between the two values. The binding factor for two suspect bidders is calculated as:

$$\phi_{i,j}^\beta = \begin{cases} 1 & \text{if } \beta_i = \beta_j \\ \beta_i/\beta_j & \text{else if } \beta_i < \beta_j \\ \beta_j/\beta_i & \text{otherwise} \end{cases}$$

where $0 \leq \phi_{i,j}^\beta \leq 1$.

Consider the maximum collusion scenario held over m auctions. The legitimate bidder will have high α and β ratings (i.e., high participation rate and number of bids). The skills will have the same α ratings as the bidder, but low β ratings. The collusion graph will initially incriminate the legitimate bidder with the skill's group (due to the high participation rate). However, the legitimate bidder will have a high β rating, as they must submit more bids than the group of skills to win (i.e., $\lceil n_j/2 \rceil$ bids in the extreme scenario).

In this case, the binding factor, $\phi_{i,j}^\beta$ can be used to compare the bidder to any of the skills. $\phi_{i,j}^\beta$ will be low for the bidder, (i.e., s/he does not bind to any of the skills). Whereas $\phi_{i,j}^\beta$ will be high for the colluding group. Therefore $\phi_{i,j}^\beta$ serves to exonerate the legitimate bidder. Furthermore, if the bidder has actually managed to win any auctions, his/her γ rating will be low and thus provide a strong indication of integrity.

In order to determine which bidders are the most likely to be in collusion, a *collusion score* (denoted as CS^η) is used. CS^η combines the η collusion rating, binding factor and other aspects of the skill score ratings, to form an index which measures the extent of collusive behaviour a bidder exhibits in terms of the alternating bid strategy.

The set of bidders \mathcal{B} , is partitioned into disjoint sets of suspected colluding bidders. This is done based on how similar bidders' η ratings are to other bidders. Let $\mathcal{C} = \{C_1, C_2, \dots, C_k\}$ denote the set of all sets of suspicious bidders. Initially, the first bidder, i is assigned to C_1 . Bidder j ($i \neq j$) is added to C_1 if $\eta_i = \eta_j \pm \lambda$, where λ is an error factor. This process is repeated for all remaining bidders until every bidder is assigned to a set in \mathcal{C} . A bidder that is not suspicious will be assigned to a singleton set (i.e., s/he is the only bidder in the set). Note that this is probably not the 'best' method for grouping bidders, as some bidders whose η ratings fall within the range of $\eta_i \pm \lambda$ might miss out on being grouped with more similar bidders later in the process (i.e., those with a closer η rating). A more optimal grouping method remains the focus of future work.

A further problem with this approach to grouping is that two bidders may have similar η ratings, but have not actually participated in any auctions together. In this case, grouping the two bidders together would be incorrect as the grouping is in conflict with what behaviour the collusion graph is trying to highlight. Instead, when two bidders are checked for similarity with their η ratings, a secondary check is then performed to see whether they have actually participated in any auctions together. If they

have, they are placed in the same group, otherwise they are placed in separate groups.

Once two (or more) bidders have been singled out based on the similarity of their collusion ratings, CS^η checks the bidder's binding ratings. ϕ_σ^β is the average binding factor for a bidder with the other suspect colluding group members. For a singleton set $\phi_\sigma^\beta = 0$. A potential skill will have a low win rate, high bid rate, and low bid increment. If $\eta_i > 0.5$, CS^η is calculated as the average of the bidder's γ , δ , ϵ and η ratings, and the average binding factor ϕ_σ^β . A bidder with a skill score of zero is excluded. Bidder i 's CS^η is calculated as:

$$CS_i^\eta = \frac{\omega_1\alpha + \omega_2\gamma + \omega_3\delta + \omega_4\epsilon + \omega_5\zeta + \omega_6\eta + \omega_7\phi_\sigma^\beta}{\omega_1 + \omega_2 + \omega_3 + \omega_4 + \omega_5 + \omega_6 + \omega_7} \times 10$$

where $0 \leq CS^\eta \leq 10$, and ω_i , $1 \leq i \leq 7$, is the weight associated with each rating.

The ζ rating is included in calculating CS^η , however, there are several points to note about its use. For an individual auction, the skill that bids first will have a higher ζ rating than the later skills. As a result, over m auctions, each skill must also alternate at bidding first. This ensures that over time each of the group members will have a similar ζ rating. A binding factor could be used for two suspect bidders based on their ζ ratings. However, this is not an overly reliable collusion indicator. Nevertheless, the ζ rating still provides some useful information. For a skill bidder engaging in the alternating bid strategy, the ζ rating will range from medium to high. Empirical evidence suggests that $\zeta = 0.75$ on average for the alternating auction strategy. In a group of bidders with similar γ , δ and ϵ ratings, the ζ rating can often be the determining factor that incriminates a skill bidder.

B. Alternating Auction Strategy

The alternating auction strategy is optimal for a colluding group, if the number of skills, ℓ' , is equal to the number of auctions held, m . In the skill case from Section I, there were 1,100 auctions, however, the skills only had 40 aliases. In this case, given m auctions, the group can evenly reduce its member's α ratings if each skill participates in m/ℓ' auctions. However, there is a certain amount of effort required to create new aliases and/or recruit new skills. This reduces this strategy's effectiveness, especially when m is large.

- Dual Collusion Graph -

The alternating auction strategy requires a different approach. The idea is to look for bidders that have never competed against each other in an auction. In this situation, the "dual" of the collusion graph is used. In the dual collusion graph, the nodes that have edges in the regular graph are removed, and edges are added between nodes that did not have edges in the regular graph. Edge weights from the regular graph have no meaning, as either two bidders have participated in at least one auction together, or they haven't. All edge weights from the regular graph are discarded. All edges in the dual graph have a weight of 1.




Figure 4. Example Dual Collusion Graph

Figure 4 shows an example of the dual collusion graph involving two bidders and two skills. There are two auctions. The skills participate in one auction each, whereas the regular bidders participate in both auctions. The bidding sequences for each auction are (s_1, b_1, s_1, b_2) and (s_2, b_1, s_2, b_2) respectively. The adjacency matrix for this collusion graph is:

	s_1	s_2	b_1	b_2
s_1	—	0	1	1
s_2	0	—	1	1
b_1	1	1	—	2
b_2	1	1	2	—

The dual collusion graph shows an association between s_1 and s_2 . The adjacency matrix for the dual collusion graph is:

	s_1	s_2	b_1	b_2
s_1	—	1	0	0
s_2	1	—	0	0
b_1	0	0	—	0
b_2	0	0	0	—

θ is a collusion rating that indicates whether a bidder may be engaging in the alternating auction strategy. This is obtained from the dual collusion graph in a similar manner to the η rating. Given a node with degree k , each edge weight is denoted as w_j , $1 \leq j \leq \ell$. The base collusion rating, θ'_i , for a bidder i is calculated as the sum of the edge weights for the node's degree:

$$\theta'_i = \sum_j^k w_j$$

To calculate a bidder's normalised collusion θ rating we first need to find the maximum and minimum base collusion ratings for the graph. These are denoted as θ^{max} and θ^{min} respectively. The normalised collusion rating is calculated as:

$$\theta_i = \frac{\theta'_i - \theta^{min}}{\theta^{max} - \theta^{min}}$$

where $0 \leq \theta_i \leq 1$.

Similar to the argument for the alternating bid strategy, it is in the group's best interest to evenly take turns at participating in each auction to collectively keep their α ratings low. Given two suspect colluding bidders, we can further compare them based on their α ratings as they should have very similar values (assuming that they have evenly alternated in their participation). The binding factor, $\phi_{i,j}^\alpha$, gives bidders i and j a value between 0 and 1, based on how similar their α ratings are. A binding value of 1 indicates a strong binding (high likelihood of collusion), whereas a binding value of 0 indicates a weak binding.

TABLE VI.
SKILL SCORE AND COLLUSION BEHAVIOUR PROFILES

Rating		Single Skill	Alternating Bid	Alternating Auction	Hybrid
α	participation rate	high	high	low	low-medium
β	bid rate	high	low	high	low-medium
γ	loss rate	high	high	high	high
δ	bid frequency	high	high	high	high
ϵ	size of bid	high	high	high	high
ζ	early participation	high	medium	high	medium-high
η	alternating bid		high		medium-high
θ	alternating auction			high	low-medium
ϕ_{σ}^{β}	binding factor		high		medium-high
ϕ_{σ}^{α}	binding factor			high	medium-high

Let α_i and α_j denote the α ratings for bidders i and j respectively. If $\alpha_i = \alpha_j$, then the binding factor, $\phi_{i,j}^{\alpha}$, is 1. If this is not the case, then the binding factor is given a decreasing value based on the percentage difference between the two values. The binding factor for two suspect bidders is calculated as:

$$\phi_{i,j}^{\alpha} = \begin{cases} 1 & \text{if } \alpha_i = \alpha_j \\ \alpha_i/\alpha_j & \text{else if } \alpha_i < \alpha_j \\ \alpha_j/\alpha_i & \text{otherwise} \end{cases}$$

where $0 \leq \phi_{i,j}^{\alpha} \leq 1$.

In order to determine which bidders are the most likely to be in collusion, a *collusion score* (denoted as CS^{θ}) is used. CS^{θ} combines the θ collusion rating, binding factor and other aspects of the skill score ratings, to form an index which measures the extent of collusive behaviour a bidder exhibits in terms of the alternating auction strategy.

The set of bidders \mathcal{B} , is partitioned into disjoint sets of suspected colluding bidders. This is done in a similar manner as the alternating bid strategy. Once two (or more) bidders have been singled out based on the similarity of their collusion ratings, CS^{θ} checks the bidder's binding ratings. ϕ_{σ}^{α} is the average binding factor for a bidder with the other suspect colluding group members. For a singleton set $\phi_{\sigma}^{\alpha} = 0$. A potential skill will have a high bid rate, low bid increment, and generally commences bidding early in an auction. If $\theta_i > 0.5$, CS^{θ} is calculated as the average of the bidder's δ , ϵ , ζ , and θ ratings, and the average binding factor ϕ_{σ}^{α} . Again, a bidder with a skill score of zero is excluded. Bidder i 's CS^{θ} is calculated as:

$$CS_i^{\theta} = \frac{\omega_1\beta + \omega_2\gamma + \omega_3\delta + \omega_4\epsilon + \omega_5\zeta + \omega_6\theta + \omega_7\phi_{\sigma}^{\alpha}}{\omega_1 + \omega_2 + \omega_3 + \omega_4 + \omega_5 + \omega_6 + \omega_7} \times 10$$

where $0 \leq CS^{\theta} \leq 10$, and ω_i , $1 \leq i \leq 7$, is the weight associated with each rating.

Note that the γ rating is excluded as it is a function of the number of auctions participated in, versus the number of winnings. This is unreliable if the colluding group has engaged in an alternating auction strategy, as the α rating will be low, thus affecting the γ rating. The alternating auction strategy has no ramifications for the ζ rating (i.e., the time a bidder commences bidding). That is, for each auction the skill will behave in a typical skill-like manner (i.e., bidding early). This means that the ζ rating is still indicative of a skill.

C. Hybrid Strategy

A hybrid strategy combines both the alternating bid and alternating auction strategies. By doing so, the colluding group can lower each member's α and β ratings. The optimal hybrid strategy requires the coalition to choose $\binom{\ell'}{p} = m$, where ℓ' is the total number of skills, p is the number of skills to be used each auction, and m is the total number of auctions.

Table VI shows the skill score ratings for each shilling strategy. Of particular interest here is the comparison between a single skill's profile and a group of colluding skills using a hybrid strategy. At most the group can only minimally influence their skill scores. However, if they do so they will be detected by the collusion graph. The only way to avoid detection is for one of the group members to have a skill score higher than the others (thus conforming to the profile of a single skill).

A hybrid strategy can be detected using a combination of the η and θ ratings. The ideal situation for the group is to keep the α and β ratings even for all colluders. However, the binding factors $\phi_{i,j}^{\beta}$ and $\phi_{i,j}^{\alpha}$, prevent the group from evenly distributing their work. As doing so will result in the colluding group having high binding factors. This forces some skills to bid more, and/or participate in more auctions than other skills (thus raising their individual α and β ratings). An extreme hybrid strategy would entail using each skill once only. However, there are practical limitations on doing this.

To detect skills employing a hybrid strategy, we introduce a third collusion score (denoted as CS^h). CS^h examines the η ratings, and the binding factors for bidders that fit the hybrid profile shown in Table VI. θ is not used, as this value is typically low for the hybrid strategy. The δ and ϵ ratings will still be indicative of a skill bidder. However the α , β and γ ratings can not be relied on. At best the hybrid strategy only achieves a moderate reduction in the group's ζ ratings. Again, a bidder with a skill score of zero is excluded. Bidder i 's CS^h is calculated as:

$$CS_i^h = \frac{\omega_1\gamma + \omega_2\delta + \omega_3\epsilon + \omega_4\zeta + \omega_5\eta + \omega_6\phi_{\sigma}^{\beta} + \omega_7\phi_{\sigma}^{\alpha}}{\omega_1 + \omega_2 + \omega_3 + \omega_4 + \omega_5 + \omega_6 + \omega_7} \times 10$$

where $0 \leq CS^h \leq 10$, and ω_i , $1 \leq i \leq 7$, is the weight associated with each rating.

Partitioning the set of bidders \mathcal{B} , into disjoint sets of suspected colluding bidders is a much more difficult task with the hybrid strategy. Let $Group_{\eta}$ and $Group_{\theta}$ denote

particular groupings for the η and θ ratings respectively. One possible approach is as follows: If bidders i and j (where $i \neq j$) both belong to $Group_\eta$ and also to $Group_\theta$, then these two particular bidders can be grouped together under the hybrid strategy. However, empirical evidence shows that this is often not the case, and bidders that belong to the same $Group_\eta$ often do not belong to the same $Group_\theta$. There is some evidence showing that several bidders that belong to $Group_\eta$ do in fact map to several bidders in group $Group_\theta$ who are shills (but weren't originally in $Group_\eta$). In future work we hope to some how use this information to further develop an accurate grouping method for the hybrid strategy.

D. Applying the SS Algorithm

In this section we briefly outline how the SS Algorithm can be applied to aid in detecting colluding shill bidders.

Consider first the alternating bid strategy. In this approach we treat all member shills as one shill and then apply the SS algorithm to it. Consider the following bidding sequence with two shills (s_1, s_2), and four bidders (b_1, b_2, b_3, b_4):

$$s_1, b_1, s_2, b_2, s_1, b_3, s_2, b_4$$

Since s_1 and s_2 have a suspicious bidding pattern, these two bidders are considered as being the same bidder which is denoted as s_s . The bidding sequence then becomes:

$$s_s, b_1, s_s, b_2, s_s, b_3, s_s, b_4$$

From here the SS algorithm can be run on this bidding sequence to determine whether the group s_s exhibits characteristic shill bidding behaviour.

However, there are problems with such a naive approach. Firstly, how do you determine which bidders are suspicious? The following instance is easy:

$$s_1, b_1, s_2, b_2, s_1, b_3, s_2, b_4$$

But what if the shill bidding sequence doesn't exactly match this? That is, the shill was outbid before s/he had time to respond, or if the shills have stopped bidding (i.e., the shill target price has been reached, or it is too close to the auction's end)? Consider the following bid sequence with four shill bidders:

$$s_1, b_1, s_2, b_2, b_1, s_1, b_2, b_4$$

In this case the colluding shills have not made every second bid, therefore combining them into a single group is difficult. The approach simply cannot be used by selecting the first bidder and then grouping them with every second bidder. Furthermore, the grouping task becomes more difficult when the number of colluding shills is large. Consider the following bid sequence:

$$s_1, b_1, s_2, b_2, s_3, b_3, s_4, b_4$$

Here each shill has only bid once. This gives the appearance of seemingly innocent behaviour.

This approach does have merit when combined with the alternating bidding collusion rating CS^η . CS^η can single out a suspect group and then the SS algorithm is used on the entire group.

This approach can also be applied to the alternate auction strategy. Consider the following bidding sequences for auctions a_1, a_2, a_3, a_4 :

$$\begin{aligned} a_1 &- s_1, b_1, s_1, b_2, s_1, b_1, s_1, b_3 \\ a_2 &- s_2, b_1, s_2, b_2, s_2, b_2, s_2, b_3 \\ a_3 &- s_1, b_1, s_1, b_2, s_1, b_2 \\ a_4 &- s_2, b_1, s_2, b_2, s_2, b_3, s_2, b_3, s_2, b_2 \end{aligned}$$

By combining the identified group members together as if they were a single bidder s_s , the bidding sequence becomes:

$$\begin{aligned} a_1 &- s_s, b_1, s_s, b_2, s_s, b_1, s_s, b_3 \\ a_2 &- s_s, b_1, s_s, b_2, s_s, b_2, s_s, b_3 \\ a_3 &- s_s, b_1, s_s, b_2, s_s, b_2 \\ a_4 &- s_s, b_1, s_s, b_2, s_s, b_3, s_s, b_3, s_s, b_2 \end{aligned}$$

As in the alternating bid strategy, the alternating auction collusion score CS^θ , can be used as the basis for the grouping.

V. PERFORMANCE

This section describes how the collusion algorithm performs on simulated auction data generated using autonomous bidding agents. The Research Auction Server (RAS) at James Cook University is an online auction server used for conducting research into privacy and security issues in online auctions (see [6]). The SS algorithm and collusion graph have been implemented on and tested using RAS. Elements of the experimental setup include:

Zero Intelligence (ZI) Agent - An agent designed to simulate an ordinary bidder in an auction. It is assigned a random amount, which it tries to submit as a proxy bid at a random time throughout the auction. The agent's price is generated randomly according to a uniform distribution.

Simple Shill Bidding Agent - An agent designed to insert fake bids into the auction in order to inflate the price for the seller. 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 (e.g., during slow bidding, or near the auction's end). (See [7] for implementation specific details regarding the simple shill agent.) In general, the riskier the shill agent, the more profit it can attain, but the higher the likelihood that it will win the auction. A risky shill agent will display more characteristics that will increase its shill score than a risk adverse shill.

Auction - A software simulated auction. The auction has a start and end time. The bidding agents can submit bids during this time, where the auction outcome depends on the agents' actions. All auctions are English auctions with proxy bidding.

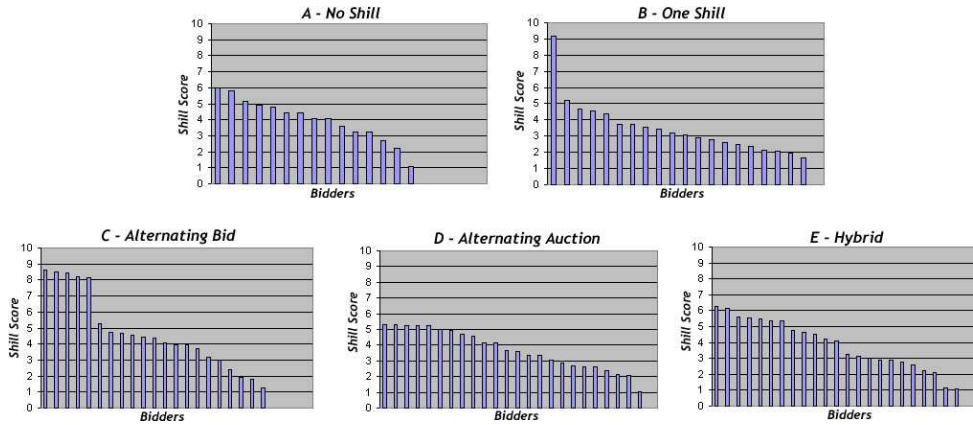


Figure 5. Skill Scores for Simulated Auctions with Differing Collusion Strategies

userid	α	β	γ	δ	ϵ	ζ	Shill Score	η	Group _{η}	ψ_{σ}	Collusion Score _{η}
SHILL_g	0.45	0.11	0.89	0.91	0.98	0.82	6.43	0.98	16	0.93	8.90
SHILL_d	0.45	0.12	0.89	0.89	0.95	0.85	6.41	0.94	16	0.93	8.85
SHILL_h	0.45	0.10	0.89	0.89	0.98	0.77	6.35	1.00	16	0.90	8.76
SHILL_c	0.45	0.13	0.89	0.95	0.96	0.87	6.50	0.86	15	0.93	8.74
SHILL_e	0.50	0.12	0.90	0.87	0.96	0.76	6.55	0.98	16	0.91	8.69
SHILL_j	0.45	0.11	0.89	0.88	0.96	0.70	6.28	0.94	16	0.95	8.52
SHILL_b	0.50	0.14	0.90	0.88	0.97	0.68	6.52	0.87	15	0.93	8.31
SHILL_l	0.45	0.11	0.89	0.84	0.95	0.62	6.15	0.94	16	0.95	8.27
SHILL_a	0.50	0.14	0.90	0.87	0.95	0.66	6.48	0.92	14	0.77	8.13
SHILL_f	0.45	0.11	0.89	0.85	0.95	0.67	6.22	0.92	14	0.77	8.11
agent'h	0.35	0.20	0.86	0.90	0.56	0.62	5.46	0.69	5	0.00	0.00
agent'r	0.10	0.71	0.50	0.82	0.86	0.69	4.34	0.17	12	0.00	0.00
agent's	0.25	0.26	0.00	0.75	0.70	0.61	3.13	0.46	2	0.78	0.00
agent't	0.15	0.13	0.67	0.28	0.83	0.27	3.49	0.24	13	0.00	0.00
agent'q	0.30	0.23	0.83	0.84	0.49	0.48	4.98	0.50	11	0.00	0.00
agent'o	0.05	0.20	0.00	0.97	0.99	0.39	2.53	0.22	10	0.00	0.00
agent'f	0.10	0.08	0.50	0.94	0.79	0.29	3.46	0.22	4	0.00	0.00
agent'q	0.15	0.23	0.00	0.85	0.67	0.15	2.34	0.42	3	0.71	0.00
agent'e	0.20	0.21	0.00	0.89	0.73	0.60	3.03	0.46	2	0.78	0.00
agent'd	0.20	0.29	0.00	0.45	0.75	0.25	2.39	0.40	3	0.69	0.00
agent'b	0.15	0.29	0.00	0.76	0.70	0.33	2.50	0.49	2	0.73	0.00
agent'c	0.25	0.18	0.00	0.77	0.69	0.58	3.05	0.41	3	0.66	0.00
agent'i	0.05	0.20	0.00	0.00	0.93	0.00	1.23	0.00	6	0.00	0.00
agent'i	0.15	0.12	0.00	0.71	0.82	0.27	2.36	0.40	3	0.48	0.00
agent'n	0.15	0.45	0.00	0.73	0.68	0.54	2.81	0.32	9	0.00	0.00
agent'a	0.10	0.53	0.00	0.80	0.66	0.56	2.73	0.10	1	0.00	0.00
agent'm	0.10	0.39	0.00	0.96	0.61	0.71	2.83	0.21	8	0.00	0.00
agent'l	0.25	0.31	0.80	0.66	0.40	0.60	4.64	0.38	3	0.65	0.00
agent'k	0.20	0.26	0.00	0.56	0.74	0.37	2.58	0.32	7	0.00	0.00
agent'o	0.15	0.17	0.00	0.91	0.77	0.51	2.75	0.46	2	0.68	0.00




Figure 6. Example collusion scores for the alternating bid strategy

Four main test types were conducted. Each test consisted of ten auctions with 20 ZI bidders. All ZI bidders were given the opportunity to participate in any auction. However not all ZI agents can participate in every auction. This is because a ZI agent is unable to participate in a particular auction if the current bid exceeds its proxy bid at the stage in the auction when it is first permitted to bid. The number of shills used varied with the type of test being conducted. The shill agent's target price was set to \$6.50 and programmed to bid in a moderately risk-adverse manner (see [7]). All tests were

repeated, and standard results are given.

Claim 1 Colluding shills were able to reduce their collective shill score in relation to a single shill. As expected, the colluding shills have lower α and β ratings than a single shill, and as a result lower shill scores. Figure 5 A shows a series of auctions without shilling. The single shill's shill score was 9 (see Figure 5 B), whereas five colluding shills using the alternating bid strategy managed to reduce their shill scores to 8 (see Figure 5 C). Five colluding shills employing

the alternating auction strategy achieved skill scores around 5 (see Figure 5 D). This clearly shows that the alternating auction strategy is more effective against the SS algorithm. However, as previously stated, the alternating auction strategy is more expensive in terms of the number of skills required, when compared to other strategies. Five colluding skills employing a hybrid strategy achieved a skill scores in the range of 5 to 6 (see Figure 5 E). In general, the larger the colluding group (relative to the number of legitimate bidders), the lower the group's skill scores.

Claim 2 *The collusion graph accurately identified bidders that engaged in the alternating bid strategy.* This test involved ten colluding skills employing the alternating bid strategy. Figure 6 graphically shows each bidder's collusion score. The ten skills consequently have the highest ratings. Figure 6's bottom section shows the skill statistics for the top ten rating bidders in terms of collusion score. The skills were divided into three distinct groups, which none of the regular bidders were members. The suspicious group members substantially bind with ϕ_σ^β ranging between 0.77 and 0.95. It is interesting to note that all the legitimate bidders scored 0 (for CS^η). Results were consistent when this test was repeated.

Weightings for CS^η were set based on empirical observation. In this example the weights are $\omega_1 = 1$ (α), $\omega_2 = 3$ (γ), $\omega_3 = 2$ (δ), $\omega_4 = 2$ (ϵ), $\omega_5 = 5$ (ζ), $\omega_6 = 5$ (η) and $\omega_7 = 3$ (ϕ_σ^α). While α (number of auctions participated in) is a significant factor, it was given a lower rating for this particular test. This was due to the large number of skills compared to the small number of auctions. Even though every skill had the opportunity to participate in each auction, not every skill did as ten skill bids were not always needed (i.e., ten being the number of skills). As a result, the α rating for each skill was not significantly high.

The γ rating (number of winnings) was given an increased weighting, as even though each skill may not have got the chance to participate in every auction, for the auctions they did participate in, they will have lost. ϵ and δ are still strongly indicative of skill bidding behaviour (i.e., fast bidding by small bid increments), and were weighted appropriately.

The ζ rating (time commenced bidding) was given a higher weighting. Even though ζ will be reduced for the alternating bidding strategy (as skills alternate and the commencement time will average out), this was a significant factor that exonerated legitimate bidders with high δ and ϵ ratings from skills. (That is legitimate bidders who bid in few auctions, and when they did bid it was quickly and by a small amount, which resulted in them attaining high δ and ϵ ratings.)

As can be seen from the results, skills engaging in the alternating bidding strategy rate extremely high in terms of the η rating. As such, the η rating was also given a significant weighting. The binding factor ϕ_σ^α , also performed strongly for the skills, but is less significant

than the η rating for determining collusive behaviour.

Since the grouping was not accurate (i.e., the skills were split between several groups), we individually singled out the skills and run the SS algorithm on the group as a whole. The skill score for the group (when run over a series of tests) was consistent with the results of *Claim 1* for a single skill. That is, the group's collective skill score over a series of tests ranged between 8 and 10.

Claim 3 *The dual collusion graph highlighted bidders that exhibited behaviour indicative of the alternating auction strategy.* This test used ten colluding skills, where each skill participated exclusively in one of the ten auctions. Figure 7 graphically shows each bidder's collusion score. In this example, 15 bidders attain a collusion score greater than 0. Out of these bidders, 11 rate higher than 7. The skills had the top ten highest consecutive ratings. Figure 7's bottom section shows the skill statistics for all bidders in terms of the collusion score. The skills were allocated to two groups. However, several legitimate bidders were also caught up in some of these groups. This illustrates that detecting skills is not an exact science. Innocent bidders can be falsely associated with a colluding group. Likewise a skill might miss out on being assigned to a suspect group, hence avoiding detection. Suspicious group members bind strongly with ϕ_σ^α ranging between 0.88 and 1.

Weightings for CS^θ were set based on empirical observation. In this example the weights are $\omega_1 = 2$ (β), $\omega_2 = 1$ (γ), $\omega_3 = 4$ (δ), $\omega_4 = 5$ (ϵ), $\omega_5 = 9$ (ζ), $\omega_6 = 6$ (θ) and $\omega_7 = 5$ (ϕ_σ^β). The β ratings for this test weren't overly high due to the skill agents being relatively risk adverse. ω_1 can be increased with riskier skills. ω_2 was given the lowest weighting due to the nature of the γ rating. The γ rating is dependent on the number of auctions participated in proportional to the number of winnings. In this example, each skill only participated in one auction each, and as such each γ rating never exceeded 0.50. ω_2 can be increased with a larger number of auctions. δ and ϵ are still indicative of skill behaviour (regardless of the strategy employed) and were weighted appropriately. The ζ rating was given the highest weighting ($\omega_5 = 9$) as this is a truly defining characteristic that separates a legitimate bidder who has participated in only one auction (and lost) that has high δ and ϵ ratings. θ and ϕ_σ^β are rated relatively high to pull out bidders that display collusive behaviour indicative of the alternating auction strategy

Similar to the alternating bidding strategy, the grouping was not accurate (i.e., the skills were split between several groups). Again we individually singled out the skills and run the SS algorithm on the group as a whole. The skill score for the group (when run over a series of tests) was consistent with the results of *Claim 1* for a single skill. That is, the group's collective skill score over a series of tests ranged between 8 and 10.

Claim 4 *The collusion score was able to successfully*




Figure 7. Example collusion scores for the alternating auction strategy

single out suspect bidders employing a hybrid strategy. To validate this claim, various tests were conducted involving numerous hybrid strategies. For hybrid strategies that favoured either extreme (i.e., alternating bid or alternating auction) rated highly in the respective collusion scores. Figure 8 shows the shill and collusion scores for a hybrid strategy using $\binom{5}{2}$ combinations of skills over 10 auctions. The five skills were the highest rating in terms of CS^h . However, similar to CS^θ , some innocent bidders are incorrectly grouped with the skills.

Weightings for CS^h were set based on empirical observation. In this example the weights are $\omega_1 = 2$, $\omega_2 = 5$, $\omega_3 = 5$, $\omega_4 = 10$, $\omega_5 = 1$, $\omega_6 = 1$ and $\omega_7 = 1$. The γ rating was given a relatively low weighting due to the small number of auctions in the test. The δ and ϵ ratings were given increased weightings to single out any bidders that bid quickly and by small amounts. The ζ rating was given an extremely high weighting because tests showed that skills employing a hybrid strategy always rated highly in this area (i.e., bid early). The η rating and binding factors were given a low rating as these are less reliable indicators due to problems with groupings.

VI. CONCLUSION

This paper proposes methods for detecting colluding shill bidders. We have identified three main strategies a

colluding group of skills can engage in to collectively reduce their individual skill scores. Colluding skills can be detected by examining a group of bidders for typical collusive behaviour in terms of the identified strategies. We introduce the collusion graph, which indicates relationships between groups of bidders that exhibit collusive behaviour in the form of the alternating bid strategy. The dual collusion graph is used to indicate relationships between groups of bidders that exhibit collusive behaviour in the form of the alternating auction strategy. Once a suspicious group is discovered, they are then further examined using a binding factor that indicates how similar the group members are based on aspects of their skill score ratings (i.e., α , β , γ , δ , ϵ and ζ). The collusion score combines all of these ratings and gives each bidder a score based on the likelihood that they are engaging in collusive shilling behaviour.

The ability of the collusion score to detect collusive shill bidding was tested on simulated auction data. Experimental results indicate that colluding skills are able to decrease their skill scores, compared to a single skill. The collusion algorithm was able to successfully identify skills that engaged in the alternating bid strategy. The algorithm highlighted major colluding groups using the alternating auction strategy. When skills used a hybrid strategy, the algorithm partitioned potential colluding groups. Some

Userid	α	β	γ	δ	ϵ	ζ	Shill Score	n	Group _n	ϕ^{β}_{σ}	Collusion Score _n	θ	Group _{\theta}	ϕ^{α}_{σ}	Collusion Score _{\theta}	Collusion Score _h
SHILL_b	0.20	0.30	0.50	0.94	0.97	1.00	4.87	0.44	13	0.56	0.00	0.41	4	0.89	0.00	9.05
SHILL_e	0.20	0.28	0.50	0.72	0.97	1.00	4.65	0.40	5	0.98	0.00	0.53	6	0.63	7.53	8.69
SHILL_d	0.20	0.50	0.50	0.92	0.90	1.00	4.97	0.48	16	0.00	0.00	0.35	4	0.89	0.00	8.66
SHILL_a	0.20	0.35	0.50	0.73	0.95	1.00	4.71	0.36	7	0.26	0.00	0.59	6	0.63	7.67	8.38
SHILL_c	0.20	0.42	0.50	0.58	0.96	1.00	4.65	0.24	15	0.00	0.00	0.65	6	0.63	7.65	7.96
agent^p	0.10	0.07	0.00	0.90	0.50	0.87	2.54	0.32	12	0.67	0.00	0.53	6	0.81	6.65	7.10
agent^f	0.10	0.40	0.00	0.93	0.57	0.92	2.98	0.20	6	0.00	0.00	0.71	5	0.00	6.22	6.91
agent^d	0.40	0.26	0.75	0.48	0.70	0.73	5.31	0.92	4	0.72	7.34	0.00	1	0.75	0.00	6.83
agent^c	0.20	0.47	0.00	0.93	0.92	0.64	3.51	0.60	3	0.00	4.81	0.24	3	0.67	0.00	6.81
agent^h	0.30	0.28	0.67	0.79	0.65	0.63	4.87	0.76	8	0.60	6.63	0.12	7	0.69	0.00	6.76
agent^e	0.20	0.27	0.50	0.45	0.74	0.72	3.94	0.40	5	0.98	0.00	0.41	4	0.89	0.00	6.65
agent^l	0.10	0.29	0.00	0.83	0.94	0.55	2.78	0.68	11	0.43	5.28	0.12	7	0.31	0.00	6.34
agent^s	0.10	0.22	0.00	0.67	0.98	0.57	2.62	0.32	14	0.00	0.00	0.59	6	0.81	6.47	6.14
agent^n	0.10	0.21	0.00	0.90	0.48	0.53	2.34	0.32	12	0.33	0.00	0.53	6	0.81	5.77	5.58
agent^j	0.10	0.11	0.00	0.63	0.22	0.83	2.03	0.24	9	0.00	0.00	0.65	6	0.81	6.00	5.55
agent^i	0.40	0.18	0.75	0.59	0.52	0.43	4.91	0.88	4	0.72	6.47	0.00	1	0.75	0.00	5.49
agent^g	0.30	0.21	0.00	0.70	0.63	0.47	3.05	0.72	8	0.68	5.20	0.18	7	0.69	0.00	5.40
agent^a	0.20	0.11	0.00	0.98	0.66	0.34	2.72	1.00	1	0.00	4.84	0.00	1	0.50	0.00	5.23
agent^q	0.10	0.09	0.00	0.75	0.59	0.43	2.10	0.36	7	0.26	0.00	0.59	6	0.81	5.50	5.10
agent^r	0.30	0.17	0.67	0.72	0.49	0.33	4.29	0.48	13	0.56	0.00	0.35	4	0.67	0.00	5.01
agent^m	0.30	0.12	0.00	0.73	0.64	0.18	2.74	0.72	11	0.43	4.20	0.24	3	0.67	0.00	4.23
agent^t	0.40	0.13	0.75	0.18	0.58	0.05	4.19	0.76	8	0.52	4.66	0.12	7	0.58	0.00	3.08
agent^o	0.10	0.07	0.00	0.00	0.38	0.00	0.82	0.32	12	0.67	0.00	0.53	6	0.81	2.90	1.59
agent^b	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	0.00	0.00	0.00	0	1.00	0.00	0.40
agent^k	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	0.00	0.00	0.00	0	1.00	0.00	0.40

Figure 8. Example collusion scores and skill information for the hybrid strategy

innocent bidders may be incriminated and some skills may not be detected. However, this work forms the basis for more sophisticated skill detection techniques that minimise the extent of false incrimination and non-detection. Other future work involves examining the effects of multiple seller collusion on skill detection, and using data mining techniques to improve upon existing work conducted.

ACKNOWLEDGMENT

The Authors would like to thank Ickjai Lee for his comments and suggestions on this paper.

REFERENCES

[1] Y. Cheng and H. Xu, "A Formal Approach to Detecting Shilling Behaviours in Concurrent Online Auctions," in *Proceedings of the 8th International Conference on Enterprise Information Systems*, 2006.
 [2] S. Rubin, M. Christodorescu, V. Ganapathy, J. Giffin, L. Kouger and H. Wang, "An Auctioning Reputation System Based on Anomaly Detection," in the *Proceedings of the 12th ACM Conference on Computer and Communications Security (CCS)*, pages 270–279, Alexandria Virginia, 2005.
 [3] J. Schwartz and J. Dobrzynski, "3 men are charged with fraud in 1,100 art auctions on eBay," in *The New York Times*, 2002.
 [4] H. Shah, N. Joshi and P. Wurman "Mining for Bidding Strategies on eBay," in *SIGKDD'2002 Workshop on Web Mining for Usage Patterns and User Profiles*, 2002.

[5] J. Trevathan and W. Read. "RAS: a system for supporting research in online auctions," *ACM Crossroads*, ed. 12.4, pages 23–30, 2006.
 [6] J. Trevathan and W. Read "Fraudulent and Undesirable Behaviour in Online Auctions," in the *Proceedings of the International Conference on Security and Cryptography (SECRYPT)*, pages 450–458, 2006.
 [7] J. Trevathan and W. Read, "A Simple Shill Bidding Agent," in the *Proceedings of the 4th International Conference on Information Technology - New Generations*, pages 933–937, 2007.
 [8] J. Trevathan and W. Read "Detecting Shill Bidding in Online English Auctions," *Social and Human Elements of Information Security: Emerging Trends and Countermeasures* (in press), 2007 Available at <http://auction.maths.jcu.edu.au/research/shill.pdf>
 [9] J. Trevathan and W. Read, "Detecting Collusive Shill Bidding," in the *Proceedings of the 4th International Conference on Information Technology - New Generations*, pages 799–808, 2007.

Jarrod Trevathan received his PhD from James Cook University in 2007. He is a lecturer/researcher at James Cook University, an analyst/programmer for Osmotion Pty Ltd, and an Associate Editor for ACM Crossroads. His research interests include privacy, security and fraud prevention methods in e-Commerce applications.

Wayne Read is an Associate Professor and the Head of the School of Mathematics, Physics and Information Technology at James Cook University. His research interests include mathematical modelling, series solutions and e-Commerce.