Automated Inference of Shilling Behavior in Online Auction Systems

BY

FEI DONG
B.S., Harbin Institute of Technology (China), 2006

PH.D. THESIS

Submitted as partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Computer Science
in the Graduate college of the
University of Illinois at Chicago, 2012

Chicago, Illinois

Defense Committee:

Sol Shatz, Chair and Advisor
Peter Nelson
Barbara Di Eugenio
Andy Johnson
Haiping Xu, University of Massachusetts Dartmouth

This dissertation is dedicated to:

my parents, Guizhi Zhao, Jianhua Dong and my husband, Juzheng Li.

ACKNOWLEDGMENTS

My deepest gratitude is to professor Sol M. Shatz for his mentoring, encouragement, support, understanding and patience during my graduate studies at UIC. He taught me not only how to express ideas but also think critically. He gave me the freedom to explore independently, and at the same time the guidance when I am lost. I have been fortunate to have him as my advisor.

My co-advisor Dr. Haiping Xu has been always there to listen and give feedback in the last five years. I am deeply grateful to him for many discussions that enriched my thoughts and helped me sort out the technical details. I am also thankful to him for carefully reading and commenting countless revisions.

I would like to thank the members of my dissertation committee, Dr. Peter Nelson, Dr. Barbara Di Eugen, Dr. Andy Johnson for their insightful comments and constructive criticisms, which lead my dissertation to a higher standard.

The most special thank you goes to my family, whom this dissertation is dedicated to. I want to thank my parents for their never fading support, belief, love and everything they have done for me. They made all of this possible. I also owe a huge thank you to my best friend and husband, Juzheng Li. His unwavering love, unconditional support and endless encouragement have sustained me throughout all of my graduate studies.

FD

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

IC3	**Internet Crime Complaint Center**
BPS	**Bids Per Seller ratio**
ATM	**Agent-based Trust Management**
SDFS	**Shill-Deterrent Fee Schedule**
IPV	**Independent Private Value**
SCBS	**Shill Counteracting Bidding Strategy**
N model	**Average number of bids model**
M model	**Average minimum starting bid model**
P model	**Bidders' Profile model**
SS	**Shill Score**
LTL	**Linear Temporal Logic**
DAM	**Dynamic Auction Model**
LAMSTAR	**Large Memory Storage and Retrieval**
D-S	**Dempster-Shafer**
BMA	**Basic Mass Assignment**
TLB	**Time of Last Bid**
CBA	**Concurrent Bid Activity**
ACB	**Abnormal Concurrent Bids**
AF	**Average Feedback**
BIA	**Bidding Increment Activity**
WPB	**Wins Per Bid**
NOW	**Number of Wins**
NOB	**Number of Bids**
AS	**Affinity for Sellers**
NB	**Number of Bids**
SP	**Starting Price**
NLARU	**No Longer A Registered User**

SUMMARY

Shill bidding is one of the most prevalent forms of auction frauds that violate the integrity of online auctions. A shill is a person who pretends to be a legitimate buyer and feigns enthusiasm for an auctioned item by bidding up the auction price. Shill bidding is very popular in online auctions due to a speculative risk for a huge potential gain and it is often covered up, making victims suffer without notice. Shill bidding produces undesirable effects not only on the auction participants but also on the auction mechanism itself as a resource allocation market. In the worst case scenario, shill bidding could lead to auction market failure.

This dissertation attempts to address these problems from three major aspects:

1) We provide a state-of-the-art review of the online auction fraud, especially shill bidding. We classified online auction frauds, summarized shill bidding indicators, reviewed shill bidding research literatures in both economic and technology sides, and then pointed out potential research challenges.

2) We discuss reasoning under uncertainty for shill detection in online auctions using Dempster-Shafer theory. Based on the conceptual framework of Dempster-Shafer theory, a unique practical shill detection approach has been proposed. This method in essence takes evidence from different levels, i.e., auction-level and bid-level, into consideration. The knowledge from auction properties and bidding behaviors are represented and quantified. Using Dempster's rule of combination, we combined evidence that enforces each other and resolved the conflicts between different pieces of evidence. The case study shows that our proposed approach is accurate and practical for real world deployment.

3) We explore price comparison as a possible approach to identifying shill bidding in online auctions. We studied whether auction users can infer shill-bidding behavior from the difference between actual auction price and expected auction price. By employing chi-square test of independence and a logistic regression model, we examined the contrary predictions made by extensions to existing auction theory in an attempt to explain how bidders can infer shill bidding from the difference between final selling price and expected price. The results show that the relationship of final auction price and expected auction price could be considered as a reliable indicator of shill bidding. We also found that a lower-than or equal-to expected final auction price is quite persuasive in concluding

SUMMARY (continued)

unlikely shill bidding and a higher-than-expected final auction price suggests possible shill bidding. We believe that the rules derived from the hypotheses tests in this dissertation are helpful and applicable for both auction houses and auction users. Honest bidders can protect themselves from being cheated and reduce the risk of monetary loss in online auctions. In addition, auction houses can adopt the rules to complement existing shill detection techniques and enhance the confidence of auction users towards auction houses.

I. INTRODUCTION

A. Online Auction

The Internet revolution and advances in information and communication technology have laid the groundwork for online auction systems to become a new profitable business platform. To cite just one well-known example, eBay, the world's largest auction website, announced $11.7 billion revenue for the year of 2011. Clearly many buyers and sellers are attracted to online auctions. According to one source [NCL08], at least 31% of Americans who have Internet access regularly participate in online auctions, accounting for a sizeable total of 35 million people.

1. Online auction classification

There are dozens of traditional auction types; however, the types of auctions on the Internet are limited. The primary auction types are the English auction, Dutch auction, first price sealed-bid auction and second price sealed-bid auction (also known as Vickrey auction) [V61, MW82, MM87]. Due to its widely known auction rules and its efficiency as a resource allocation mechanism, the English auction has become the most popular type of online auction among both service providers and consumers. It is a typical open, ascending-price auction in which bidders compete with each other by placing higher bids. When a predefined time expires or the highest bid reaches the pre-determined buyout price (a price at which the seller would like to sell the item during the auction), the highest bidder wins the auction. The winner pays the highest price, namely, the last bid. According to a preliminary study [CX08], the English auction and its variants account for 88% of Internet auctions while Dutch auctions account for 1%; others such as sealed-bid auctions and double auctions comprise the remaining 11%.

Unlike the English auction, the Dutch auction is a type of descending-price auction [B92]. This type of auction requests a high price at the beginning, and then the price is lowered gradually until a participant is willing to accept the price, or a predetermined minimum price is reached. The winning participant pays the last asked price. Dutch auction is also used in online auctions where multiple identical items are sold simultaneously to one or more winning bidders. It is equivalent to a multi-unit English auction in the economics literature [MM05].

While both English and Dutch auctions are open auctions in which participants know each bidder's bidding price, sealed-bid auctions are auctions in which each bidder bids just once and the bid price is kept as a secret during the auction. The first price sealed-bid auction is an auction in which all bidders submit their bids at the same time, and all participants are ignorant of others' bids. The winner is the one with the highest bid, and pays that bid.

The second price sealed-bid auction (Vickrey) is also a sealed-bid auction, like first price sealed-bid auction. The only difference is that, in second price sealed-bid auction, the winner pays the second highest bid rather than the winner's own bid. It has been proved that the second price sealed-bid auction is a mechanism strategically equivalent to the English auction, but it gives bidders an incentive to bid their true values, making this type of auction important for auction theory [L99]. From the auction theory point of view, eBay's automated proxy bidding is similar, but not identical, to the Vickrey auction. On eBay, the winning buyer as determined by the automatic proxy bidding system does not pay their own highest bid, but instead pays the second highest bid plus a predefined minimum increment. This is different from a Vickrey auction because eBay's bidding proxy can submit bids multiple times in an auction on behalf of the buyer, and the buyer can also change the maximum bid accordingly. However, in a Vickrey auction, every bidder can only submit a bid once and the bid is not changeable.

In English, Dutch and sealed-bid auctions, there is usually only one seller in an auction. However, in a double auction, there might be many sellers, where sellers and buyers offer and submit bids in any order [WWW98]. Then bids are ranked from highest to lowest, and offers are ranked from lowest to highest to generate a supply and demand profile. When offers and bids are matched (bids move down and offers move up), the required quantities of goods are exchanged. Double auctions are commonly used in futures markets.

One more way to classify auctions is according to their durations. Some auction sites have a fixed time schedule, in which when time runs out, the auction is closed. For instance, eBay offers users the option of 1, 3, 5, 7 and 10 day auction durations. A one day auction on eBay lasts exactly 24 hours. Other auction sites such as uBid and Yahoo!Auction (US Yahoo!Auction has declared its retirement since June 16, 2007) have a time schedule similar to that of some traditional auctions – when the time limit runs out, time is extended as long as someone outbids the others.

2. Bidding strategy

A bidding strategy is a plan of bids designed to achieve a particular goal, such as winning an item at a low price or exposing highest bidders. In auctions, bidders may choose a single bidding strategy or use mixed bidding strategies. In fact, there may be as many bidding strategies as bidders. In Table I we list some representative bidding strategies that are commonly used in auction markets [SJS03].

TABLE I.

FIVE COMMON BIDDING STRATEGIES

STRATEGY	DESCRIPTION
Skeptic	Bid multiple times but bid as low as possible each time.
Proxy bidding	Specify a maximum bid initially and then authorize the proxy to bid automatically as many times as necessary up to the maximum.
Sniping	Bid in the last seconds, leaving no time for anyone else to outbid.
Unmasking	Bid several times in a short period of time with the purpose of exposing the maximum bid or the highest bidders.
Evaluator	Bid just once at an early time with a high value.

B. Online auction fraud

Although the number of sellers and buyers attracted by online auctions is growing rapidly, this contemporary business medium faces an important challenge – auction fraud [A02, CSK05, GS08]. Both sellers and buyers can participate in auction fraud for their own benefit. Data released by the U.S. Federal Bureau of Investigation's Internet Crime Complaint Center (IC3) reveals that 93,771 auction complaints were received in 2006, representing 45 percent of all Internet fraud complaints [M07]. Auction complaints remain the largest source of Internet-related complaints, consistently ranking at the top of the list for many years [FTC03].

According to IC3, there are several ways online auction fraud can occur: misrepresentation of a product for sale, non-delivery of merchandise or services sold, triangulation (fraudsters purchase items using a stolen credit card, selling the items to uninitiated buyers thereby retaining the cash and transferring the risk of seizure to the end recipient), fee stacking (charging extra money after an auction is over), selling black market goods, multiple bidding (buyers inflate prices using aliases, which frustrates competitors, then at the last moment the high bids are withdrawn to secure a low bid), and finally shill bidding (sellers or their associates place bids on their own auctions for fraudulent purpose).

To understand online auction fraud, it is convenient to first classify the various types of online auction fraud according to the three time periods in which the fraudulent behavior can take place: pre-auction, in-auction and post-auction (Figure 1). Misrepresentation of items, selling of black market goods and triangulation usually occur before the auctions start, so we classify them as pre-auction fraud; and non-delivery of goods and fee stacking occur after auctions close, so we consider them as post-auction fraud. In-auction fraud is the main focus of this dissertation and so the discussion of in-auction fraud is postponed to the following sections.

Because both pre-auction and post-auction frauds involve offline behaviors that can often easily be noticed by buyers and sellers, investigation of such frauds relies more on real-world evidence than on online prevention and detection mechanisms. However, in-auction fraud happens while transactions are in progress, thus it may occur without leaving direct physical evidence, and worst of all may not even be noticed by the victims. In addition, while pre-auction fraud and post-auction fraud have already attracted researchers' and policy makers' attention, in-auction fraud has attracted much less attention due to its complexity in detection [GT08, BP02]. In order to reduce the loss to victims and to protect online business participants, in-auction fraud deserves more attention and effort from mechanism designers and information technology researchers.

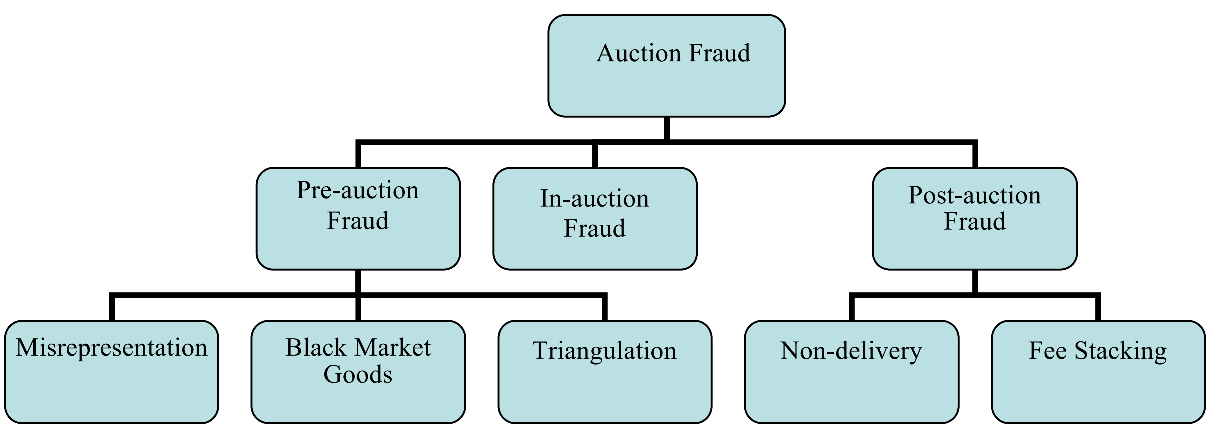

Figure 1. Auction fraud categorization

Many consumer guidance websites and newspapers have provided auction fraud detection tips, such as checking if a seller and a suspicious bidder are from the same geographic region; and searching a shill suspect's bidding history to determine if a seller and the shill suspect have a partnership. However, even if a suspected shill and a seller are located far from each other, they still could be partners. With modern Internet communication applications, sellers and shills can communicate with each other very easily as if they were sitting next to each other. In addition, whereas there may be a large number of historical records for long-time sellers, it is difficult and time-consuming for a bidder to discover the partnership between sellers and shills. Unfortunately, in-auction fraud is so sophisticated and tricky that such tips are extremely difficult for consumers to effectively apply.

Many researchers from economics, business, system science and computer science have realized the severity of this problem [G04, GS06, NB08]. They are working to combat auction fraud and have delivered some preliminary analysis and results. In this chapter, we aim to provide an overview of the state-of-the-art Internet in-auction fraud prediction, prevention and detection techniques, and to highlight challenging research issues in this interesting new area. The surveyed researches are from both the economics and the computer science perspectives.

C. The Nature of Online In-Auction Fraud

Current and emerging state-of-the-art auction platforms rely heavily on information technology, and hence any weakness in information systems could be utilized by malicious users to maximize their own profits. Although in-auction fraud has an important influence on new technology-based sectors of the economy such as e-commerce, there is a lack of mechanisms to fight against in-auction fraud. This is because in-auction fraud has not been a major issue in traditional auctions, and therefore, it was barely considered in traditional literature.

Due to the nature of Internet applications such as a high degree of anonymity, incomplete legal constraints and lower barriers to entry and exit, it is difficult to combat in-auction fraud. The Internet characteristics inevitably result in information asymmetry between sellers and buyers in online auctions, where one party of the transaction has more or better information than the other party. In an auction, sellers usually have better knowledge about the item than buyers; therefore shills who are associated with the sellers also have better knowledge than other buyers. The fraudulent participants intentionally make use of this information asymmetry to obtain benefit and deliberately hide their fraudulent or opportunistic behaviors so as to avoid being detected and caught. Since most of the fraudsters sign up on auction sites with fake identity information, when fraudulent behaviors are exposed, it is hard for investigators to find out the real identity under the meaningless net IDs. Afterwards the fraudsters could start over again using new net IDs so as to conceal their identities and reduce the chance of being punished.

Online auction houses do not always exhibit a positive commitment to actively solving the auction fraud problem. Some researchers claimed that the policies of existing auction houses typically would not discourage users from cheating [KW03]. Chua et al. pointed out that with the current oversimplified evaluation practice that relies on transaction feedbacks, the online auction platforms provide leverage for a con-man to manipulate the system with the intent of deceiving [CW04]. So far, though a few online auction fraud detection approaches have been proposed, online auction systems that support smashing auction fraud are still very few [XSB08, PXG07].

Furthermore, unlike traditional auctions, online auction rules usually impose fixed time durations for auctions, and online auctions typically last longer than traditional auctions. For example, eBay offers users the option of 1, 3, 5, 7 and 10 day auction durations. These optional auction durations on the one hand satisfy the sellers' need to attract more bidders, but on the other hand they leave malicious buyers and sellers more time to cheat. The longer the auction lasts, the more chances for fraudsters to make their fraudulent behaviors look normal. With enough time, cheaters could potentially behave like normal bidders with few clues left for fraud investigation. Thus, detection of auction fraud using statistical evidence related to time would be compromised by the normalized behavior. For example, if an auction lasts one day, the fraudsters may place their first several fraudulent bids within the first few hours after the auction begins. However, if an auction lasts many days, the cheaters may not be in much of a hurry to place shill bids. A bidder might be suspected if this bidder places a high bid for a common item in one auction, but places no bid in another auction of an identical item at a lower price, running concurrently on the same site. Nevertheless, given enough time, the seller or seller's accomplice can wait until there are no sellers selling the same item for less, and then place a shill bid with less chance of being caught. The statistical investigation would therefore be highly affected.

Last but not least, although many existing fraud patterns have been suggested, they can merely be used as clues for investigation rather than evidence. Each pattern in itself may have more than one rational explanation, including fraudulent explanation and innocent explanation. A detailed discussion of the uncertainties involved in such patterns is discussed in Chapter IV.

II. SHILL BIDDING

In this Chapter, we introduce shill bidding in detail. First, the impact of shill bidding is presented. Then, different types of shill bidding, which is also a type of in-auction fraud, is introduced. Next, the indicators of shill bidding are summarized. At the end of this chapter, we discuss the challenges exist in shill bidding research.

A. Impact of Shill Bidding

Shill bidding has become more and more severe. It is probably the most prevalent form of online auction fraud. A shill is a person posing as a legitimate buyer who feigns enthusiasm for the item on auction by bidding up the price, thus serving as an accomplice to the seller. The role of a shill can be played by an associate of the seller, such as a friend or family member, or by the seller himself posing as a legitimate buyer under a fake online auction ID. With the growth of auction business, shill bidding has become not ignorable. Mass media including Consumer Affairs, the New York Times and USA Today have covered this kind of auction fraud frequently in the last few years [GA08, F07]. According to recent criminal charges, some felons placed shill bids in thousands of auctions, driving up the price from several dollars to a few thousand dollars [CA04]. Although the punishment for auction fraud could be severe (e.g., several years in prison with fines), shill bidding is still very popular due to a speculative risk for a huge potential gain. We notice that even some sellers in eBay's Power Seller program [ebay09] with 100% positive feedback confessed in online forums that they would have made much less money if they did not use shill bidding.

Shill bidding severely violates commonly perceived notions of fairness in auction markets. Wang et al. showed that private-value English auctions with shill bidding could result in a higher expected seller profit than other types of auctions, violating the classical revenue equivalence theory [WHW01]. Kauffman and Wood examined the effects of shill bidding on final bid price in rare coin auctions and showed that some bidders might view shill bids as signals that an item is worth more, thus they might pay more than other bidders who cannot see such signals [KW05].

In a worst-case scenario, shill bidding could result in an insufficient market or even market failure, which has not been considered in traditional auction literature. Shill bidding could lead to a vicious spiral of auctions. Shill bidding could result in the "shiller's curse" as indicated in [WHW02]. When buyers suspect the existence of shill bidding, they may shield their bid and wait for the seller to keep the item and sell it for less in the next round of the auction. If no bidder outbids the shill, the seller will probably try to sell the item again later. If the shill phenomenon exists in an auction market persistently, buyers would fear the existence of shills whenever they want to buy something through online auctions, and such buyers would likely bid at a price that is much lower than their valuation, considering the existence of potential shill bidding. The above process can be repeated, and then both sellers' and buyers' trust toward the auction market deteriorates. Consequently, the chaos may lead to inefficient or failed markets.

B. Key Concepts related to Auction and Shill Bidding

The current online auction literature reflects a consensus that regards buyers as bidders, even though the actual entity that places bids up to a certain limit could be an automatic bidding system, such as the proxy bidding system on eBay or the bid butler on uBid. Nonetheless, sellers and auctioneers are different according to the literature. In some auction models, researchers regard auctioneers the same as sellers, considering an auction as an online transaction between auctioneers (sellers) and buyers. Therefore, they often use the two terms "sellers" and "auctioneers" interchangeably [GM87, K04]. However, as online auction fraud becomes increasingly popular, the role of the auctioneer can no longer be overlooked. Wang et al. made a distinction between "auctioneer" and "seller." They argued that auctioneers, who are agents that conduct auctions as a third party with self-interest, play an essential and critical role in auctions, especially in in-auction fraud prevention [WHW01]. Therefore, in this dissertation, we define sellers as people who own the item for auction, while auctioneers are the entities who coordinate auctions. In online auctions, an auctioneer is simply an auction house, such as eBay. Table II provides a comparison between auctioneer and seller.

Before discussing details of shill bidding, it is useful to first understand some basic concepts related to the fee structure of auction houses. Auction houses charge user fees for the services they offer. For instance, the total cost on eBay includes an insertion fee and a final value fee, where the insertion fee is charged whenever a seller lists an item for sale, and the final value fee is charged if the item is sold. The insertion fee is calculated as a function of the starting price or the reserve price, whichever is higher. The starting price is an initial price of an auctioned item, while a reserve price is the minimum price at which the seller is willing to sell the item. Since a high starting price may drive away potential bidders, a seller typically sets a low starting price to attract bidders and a high reserve price to insure that the seller can “reserve” (keep) the item in case the winning bid is below a satisfactory value, i.e., the reserve price.

TABLE II.

ONLINE AUCTIONEER VS. ONLINE SELLER

ROLE	DESCRIPTION	MAJOR RESPONSIBILITIES	GOALS	EXAMPLE
Auctioneer	An entity that provides auction services to online auction users.	1. Provide a transaction platform and services to both sellers and bidders. 2. Make arrangements for auctions. 3. Place advertisements for auctioned items.	Provide a trust-worthy auction environment to customers and earn commissions.	eBay
Seller	An entity that offers items for sale. The seller may or may not be the owner of the item.	1. Run the auction by posting item descriptions and pictures, and taking bids. 2. Receive payments and provide the auctioned item(s) to the winner. 3. Pay commission fees to the auctioneer.	Obtain a high sale price and reduce commission costs.	A person who hosts an auction on eBay.

In-auction fraud may include shill bidding, bid shading, false bidding, multiple bidding, and bidding rings. Shill bidding happens in open auctions in which, by definition, bidders are allowed to compete with each other to bid multiple times. It could not happen in a sealed-bid auction due to the nature of concealed bids. Bid shading and false bidding are specific terms for cheating behaviors for first-price auctions and second-price auctions, respectively.

From the literature, it is clear that there are many different types of in-auction fraud, which are related to the role of the auction participants, namely the sellers and buyers. We now subdivide in-auction fraud into sellers’ fraud and buyers’ fraud, and explain them in detail. We adapt the taxonomy presented in [JZB07] and elaborate the concept of shill bidding, one type of seller-based fraud, into three concrete forms of shilling behavior that can occur in online auctions. As Figure 2 shows, sellers’ fraud includes competitive shilling, reserve price shilling, buy-back shilling and false bidding, while buyers’ fraud consists of bid shading, multiple bidding and bidding rings.

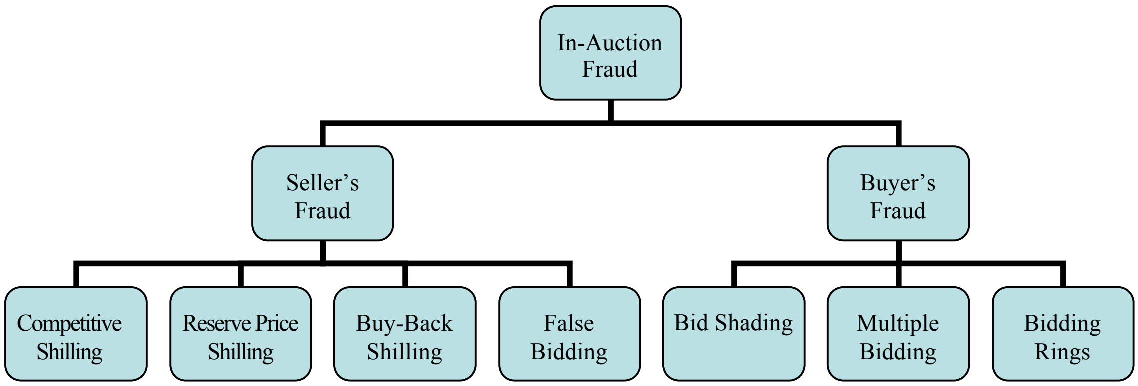

Figure 2. In-auction fraud categorization

Shill Bidding includes any activity in which a seller or an associate of a seller bids on the seller's own item in an auction. Shill bidding can be performed either by the seller or by individuals associated with the seller (including friends and family members), who may have a level of access to the seller's item information that is not available to the general community [ebay08]. We distinguish three types of shilling behaviors based on the seller's or shill's motivation to cheat either the other bidders, as in types (1) and (3) below, or the auction house itself, as in type (2).

(1) *Competitive shilling* is a bidding behavior that artificially drives up the bidding price of the auctioned item with no intention of actually buying. The purpose is to make a legitimate winner pay more than this person would otherwise pay, so that the seller can gain more profit [KW05]. The competitive shilling behavior can occur in both live and online auctions as long as the collusion between the seller and the shills remain unknown to the auction house and the other bidders. Intuitively, one can surmise that the shill bidders engaged by sellers pretend to be legitimate competing bidders, and use the shill bid to lure legitimate high-value buyers not only to bid up to their valuations but also to exceed them. This behavior cheats other bidders by inducing them to pay more for the item than they would have without the shill bids. To better understand competitive shilling, consider the following example situation. A seller hosts an auction of an item, say an unlocked cell phone. Initially, nobody places any bid at the auction. So, to attract some bidders to the auction, the seller, acting under another alias, places a competitive shill bid for the purpose of stimulating other bids. After someone outbids the shill bid, the seller places another shill bid for the purpose of driving up the bidding price. The seller hopes other legitimate bidders will outbid all further shill bids. All bids placed on behalf of the seller can be regarded as competitive shill bids and the bidding behavior is called competitive shilling.

(2) *Reserve price shilling*, first defined by Kauffman and Wood [KW03], is a bidding behavior motivated by the desire to avoid payment of the reserve price fee. By accepting the reserve price service, a seller is agreeing to pay the auction house a fee for the service. To minimize the payment of auction house fees but still reserve the item below a certain price, some high volume sellers will not set an "official" reserve price but instead engage shills to place bids on their auctions. For example, a seller may wish to sell an item at $200. If the seller chooses eBay's optional service to set a reserve price of $200, this seller would automatically incur a reserve price fee of $2.00 according to eBay's fee structure of 2012. To avoid payment of the reserve price fee but still reserve the item if the final bid for the item is under $200, the seller might first list the item at $9.99, paying the auction house a low insertion fee of $0.25. Then, either an associate of the seller, or the seller himself, places a bid at the price of $200 in hopes that a legitimate bidder will make a purchase at $201 or more. In fact, eBay in its rules against shill bidding also gives an example of reserve price shilling [ebay08]. Note that sellers do take a monetary risk when employing reserve price shilling: if nobody makes a purchase at more than the hidden reserve price, the seller must still pay both an insertion fee and a final value fee.

(3) *Buy-back shilling*, which to our knowledge has not been previously identified in the literature, is a bidding behavior employed by sellers, or other shills as agents of the seller, when the legitimate bidders do not bid an acceptably high price. The seller or shills would rather buy back the item and sell it again than sell the item now at a low price that does not reach their expectation. In this situation, shills behave as a normal bidder with the goal of buying the item at a bargain price and completing the transaction. Such activity cheats the other bidders by depriving them of purchasing an item at a bargain price. For example, as in the previous example a seller may wish to sell an item at $200 but initially sets the starting price at $9.99, producing an insertion fee of $0.25. When the auction is close to termination, if the highest bid has only reached $15, the seller may place a shill bid at $16.00 in order to buy the item back, even though the seller must pay the auction house a final value fee of $1.40 in addition to the $0.25 insertion fee. Nonetheless, this cost to the seller is trivial compared to the profit-loss the seller would have incurred if the item was sold at $15. The profit of buy-back shilling is obvious; hence we speculate that the buy-back shilling behavior does actually exist in online auctions.

Any of the above three shill bidding types can be enacted in one of the different forms described below, depending on who acts as the shill:

- Acting alone: A seller, or owner of the item for sale, carries out shill bids by himself or herself. The seller is able to register several IDs in an auction house, e.g., eBay or uBid. Using different IDs and pretending to be different legal bidders in order to bid multiple times in the seller's own auction, the seller can inflate the final auction price and profit.
- Seller collusion: Several sellers help each other to place bids on each others' transactions for their mutual benefits.
- Accomplice: A seller hires or invites family members and friends to serve as shills who will place bids on the seller's item, but instructs them to avoid winning.

Other types of in-auction fraud as shown in Figure 2 are described as follows:

False bidding: In a second price sealed-bid auction, each bidder bids only once in the auction and the winner pays the second highest bid rather than the highest. An auctioneer can help a seller profitably cheat by examining the bids under the table after all buyers have submitted their bids. Knowing all bids, the auctioneer can submit an extra bid to make the second highest price very close to the current highest price such that the seller can gain more profit [RTK90]. For example, after all buyers have submitted their maximum bids, the auctioneer learns that the highest bid is $200.00 for the auctioned Razor cell phone and the second highest bid is $120.00. The auctioneer (who could possibly be working on behalf of the seller) can help the seller insert an extra bid, say $198.00, which is quite close to the highest bid $200.00, but not beyond the highest bid. After the auction ends, the seller receives $198.00 rather than $120.00, and the extra $78.00 in revenue is gained by false bidding. This type of auction fraud may appear in auctions held by eBay when bidders are using the auction site's automatic bidding proxy system. Every buyer using the bidding proxy has to submit a sealed maximum bid to the bidding proxy. The bidding proxy then bids repeatedly by setting increments until the bid exceeds the buyer's predetermined maximum bid. When the seller is able to obtain all existing maximum bids, this seller can then place a second highest bid as a shill bid, which is slightly lower than the highest maximum bid. By doing so, the seller's revenue increases.

In addition to sellers' in-auction fraud, buyers can also cheat in Internet auctions. Bid shading, multiple bidding and bidding rings are common cheating approaches used by buyers:

Bid shading: In first price sealed-bid auctions, the winner of the auction pays the highest bid. If a bidder could use an unfair method to know the highest bid before the bids are disclosed, then the bidder could insert a bid just above the highest bid. The fraudulent bidder would thereby increase the probability of winning while minimizing the payment to the seller [PS03]. Let's take an auction of the game console Wii as an example. Assume Bidder 6 is willing to pay up to $400 for the game console before submitting any bid. When the auction begins, all bidders except Bidder 6 submit their bids. These bid values are lower than Bidder 6's valuation, ranging from $200 to $300, as shown as in the left-side table of Figure 3. Since the auction is a sealed one, a bidder typically cannot see the highest bid. By knowing the highest bid in some unusual way, Bidder 6 may guarantee a win by placing a bid at $301.

Figure 3. An example of Bid Shading

Another slightly different bidding strategy that is not regarded as fraudulent also goes by the name of bid shading [Z07]. In this case, bidders place bids below their true valuation of the item in order to avoid overpaying for the auctioned item.

Multiple bidding: Multiple bidding, also known as bid shielding, is similar to shill bidding except that it is a fraudulent behavior of buyers rather than sellers. The buyers register several aliases and use them to place multiple bids for the same item. By driving up the price with multiple auction identities, the buyers discourage other potential competitors. After that, they retract all high bids, leaving the lowest winning bid on the auction. At the end, the winner gets the auctioned item at a much lower price. For instance, consider a scenario for a Motorola Razor cell phone auction in which Bidder 3 bids $134.90. Bidder 1, who may have bid previously on this item, now realizes that Bidder 3 is a potential competitor. In order to try and force Bidder 3 out of the competition, Bidder 1 places 3 bids consecutively, namely $135.00, $270.00 and $280.00. The last two bids are obviously much higher than the previous bids; thus, when risk-neutral bidders see this situation, they will quit the auction instead of paying beyond the valuation. Therefore, Bidder 1 can secure the winning position. But, the cheating behavior comes about when Bidder 1 retracts the two high bids at the last minute of the auction, leaving only the $135.00 bid, which is the lowest cost to win the auction. This cheating method works only at auction websites that allow retracting bids. While almost all current auction websites' rules generally disallow retracting bids, they still allow retracting bids under exceptional circumstances such as a typographical error in entering the bid [eba08].As we observe, bids retractions occur often in many active auction websites.

Bidding Rings: Bidding rings is also a term related to bidders' fraud. It refers to collusive auction fraud behaviors conducted by several bidders. Several fraudulent bidders form a ring, and the ring members have an agreement not to bid against each other, either by avoiding bidding on the auction or by placing phony (phantom) bids to not compete with each other. The result is that the winner can win the auctioned item at a very low price.

C. Shill Bidding Indicators

From the perspective of sellers' fraud, the following bidder characteristics may indicate the existence of shilling behaviors. We assume that when a buyer is interested in an item, the buyer is able to know key sale information related to the item, such as auctioning prices of the same item in other auctions on the same website, the auctions' end time, etc.

1) In concurrent auctions, when bids are placed on an auction with a higher price rather than the ones with lower prices, the bidder might be a shill. For example, three sellers are auctioning the same brand new Motorola Razor cell phone concurrently. The current bidding prices for auctions held by seller 1, seller 2, and seller 3 are \$80.00, \$100.00, and \$150.00, respectively. To simplify matter, we assume that the Razor cell phones sold in the three auctions are identical, the three auctions end almost at the same time, and the reputation rankings of the three sellers have no effect on the bidders' decision. Now if a bidder places \$160.00 on the auction held by seller 3, knowing that a currently winning bid could be placed at a lower price with Seller 1 or Seller 2, the bidder should be highly suspected as a shill. This indicator can be supported by the fact that the legitimate buyers are usually looking for bargains at auction websites but unlike normal bidders, shill bidders do not care about price very much because they do not really intend to buy a cell phone and their real intention is to drive up the price.

2) Shills usually have higher number of Bids Per Seller ratio (hereafter BPS) than that of normal bidders. The auction house maintains records of the number of bids a bidder has placed for every seller that the bidder has

interacted with. Normal active bidders should have placed a number of bids on several different sellers' auctions, but shill bidders usually only deal with a very limited number of sellers, and most shill bidders only place bids on one or two sellers' auctions. A simple example is shown in Table III, where users *S1*, *S2* and *S3* are all professional cell phone sellers. *B1* and *B2* are bidders who have recently placed bids for cell phones. Bidder *B1* has only placed bids on auctions hosted by *S1*. Also, this bidder completed 6 transactions with *S1* with a total number of 39 bids. For *B1*, the BPS ratio for *S1* is 39; the BPS ratio for seller *S2* and *S3* are both 0. The bidder *B2* bids separately on each seller's auction. For *B2*, the BPS ratio for *S1*, *S2* and *S3* are 10, 6 and 8, respectively. The statistical result of BPS ratio shows that *B1* is possibly a shill because *B1* behaves "strangely" compared to normal bidders who bid on any auctions instead of a specific one, where *B1* only bid with one seller, *S1*.

TABLE III.
AN EXAMPLE OF SHILL INDICATOR

Seller / Bidder	S1		S2		S3	
	No. of auctions	No. of bids	No. of auctions	No. of bids	No. of auctions	No. of bids
B1	6	39	0	0	0	0
B2	1	10	1	6	1	8

3) Shills generally avoid winning auctions. Because the general purpose of sellers in deploying other people as shills or acting as shills themselves is to inflate the price but not to obtain the item, this indicator is straightforward. As we mentioned previously, the total cost of selling an item at an auction house includes an insertion fee and a final value fee. When listing an item at the auction house, the seller is charged an insertion fee. If the item sells successfully, the seller is also charged a final value fee. If the shill wins, the seller needs to pay both fees to the auction house, and the partnership of the shill and the seller together would not come out ahead because of the fees paid to the auction house. Therefore, shills bid with caution to avoid outbidding potential winners. There is an exception in buy-back shill. In this case, if at the end of the auction the price is much lower than the sellers' expected value, then the seller would bid aggressively to buy the item back with the hope to sell it at a high price in the next round of the auction.

4) Shills usually place bids at the very beginning of auctions if the auction duration is not long enough. By doing so, (1) sellers could set up a hidden reserve price earlier; 2) the early bids can serve as stimulating bids which attract other bidders; and (3) shills could leave legitimate bidders sufficient time to outbid the shill bids, thus reducing the chance of an unexpected, and undesired, winning of the auction.

5) Shills may bid with the minimum increment that is established by an auction house. Especially during the later stage of an auction, if a shill's bidding increment is too large, the bidding prices may drive away other buyers, resulting in the shill accidentally winning the auction.

6) Shills usually do not receive many feedbacks. This indicator stands because the easiest way for a seller to manage a shill is for the seller to act alone. Since a shill seldom wins auctions, the seller does not need to provide any feedback for the shill. Note that online auction houses typically allow users to provide feedback to either sellers or buyers. However, this evidence alone is insufficient for incriminating a shill, because it is possible that the bidder is a newly registered user and does not have any feedback yet.

Now from the auction perspective, the following characteristics of auctions themselves indicate that shills may be involved:

1) Auctions with shills have more bids on average than those without shills. Shills tend to outbid the legitimate bids frequently until the price reaches their expected value, or when the risk of winning the auction becomes high. The bids that the shills stimulated and placed contribute to the extra amount of bids in the auction.

2) The average minimum starting bid in an auction with shills is less. It means that the higher the starting bid (compared to book value), the less possibility that the auction involves a shill. Conversely, if the starting bid is much less than the book value, it is more likely it involves a shill. This indicator is explained and tested in [KW03].

Each of the above characteristics could be used as an *indicator* of the existence of shills, but not as evidence, because each characteristic may have some "innocent" explanations other than the existence of shilling behaviors. For example, one explanation for the first bidder characteristic given above could be as follows: some experienced buyers may prefer high rated sellers over low rated ones, even though highly rated sellers may sell the same item at a higher price. This preference for dealing with particular sellers is reasonable; common sense tells us that sellers with high reputations are trusted for quality of service. When a bidder has indicator 1, it is also possible that the bidder simply does not have enough information about the price of other identical items. As another example, consider indicator 2, as given previously. This behavior might well indicate the existence of shilling, but it might instead indicate something more innocent. For example, maybe the sellers are the unique sellers of a specific kind of hard-to-find item, and the buyers solely collect this kind of item in online auction houses. Finally, consider the following question: what if the bidder is aware of the existence of other similarly rated sellers who also sell the same item on the same site? It could still be possible that the bidder simply trusts some particular seller and that seller's goods have always met the buyer's requirements. Hence, each indicator itself cannot serve as obvious evidence of shills; a combination of several indicators may be more persuasive than a single indicator in detecting shills.

D. Research Challenges

In-auction fraud is quite different from pre- and post-auction fraud in that the latter two can be easily detected by the victims while in-auction fraud cannot. This difference makes solutions to pre-auction and post-auction fraud not adoptable for in-auction fraud. Research on combating in-auction fraud has only recently begun and several challenging research tasks are still open problems. Since this is a new research area, we identify, and briefly discuss, some research challenges that we feel are important in this area.

Development of effective reputation systems. Reputation systems are the easiest accessible tool that online auction participants can rely on to evaluate a seller or a buyer. However, as we analyzed previously in this dissertation, current reputation systems in major auction houses fail to provide users trustworthy and accurate information. Thus, it is important to propose effective reputation mechanisms to provide users with convincible ratings. By giving users an explicit indication of the genuineness of the rated user behind the reputation score, next-generation online reputation systems should be able to encourage trustworthy behaviors and significantly prevent in-auction fraud.

Real-time fraud detection. The most efficient way to reduce the loss resulting from in-auction fraud is to detect the fraudsters as early as possible. If the auction system can successfully detect the presence of auction fraud immediately after it happens, the auction house can cancel involved auctions so that the shills can be caught, and the victim can be protected from losing money and property. However, most existing fraud detection techniques cannot guarantee real-time detection of in-auction fraud. Efficient shill detection algorithms such as using a model checking based approach could be a very promising approach for real-time shill detection as demonstrated in [XC07, XBS09].

The lack of ground truth. Development of techniques and tools that aim to assist with detection of in-auction fraud, such as shilling behavior, is clearly not an easy task. An even more challenging task is the assessment of effectiveness of such techniques. How well do they really work in practice? This assessment is complicated by the subtle behaviors associated with such fraud and the lack of example data that includes actual, verified fraud behavior. So, how to obtain the ground truth becomes a problem. How does one demonstrate that a shill-detection technique is effective on a set of sample auction data, when the existence of actual shill behavior in the sample data is not known?

In theory, the most reliable way to obtain the ground truth would be to have access to confessions, or legal indictments, of sellers who employed shills or acted as shills themselves. However, this does not seem feasible. Since shill bidding is not considered as a very honest behavior, shills typically work hard to conceal their involvement in such activity. Yet, in the process of researching this area, we did discover several sellers in an online auction forum confessing that they had employed shills. However, we could not reliably connect the forum usernames to online auction house usernames since neither is a real identity. We also made some efforts to obtain auction data involving shills from auction houses such as eBay. Unfortunately, at this time, auction houses do not disclose such information.

In other domains, researchers often turn to “expert opinions” as an alternative to ground truth. For example, in the area of spam review detection, it is also difficult to get the ground truth regarding which reviews are actually spam or fake. Mukherjee et al. hired eight online review experts to assess the “spamicity” of 2400 English language reviews – to classify them as being “spam,” “borderline spam” or “non spam” [MLG12]. In the field of natural language processing, it is common to establish ground truth by manually annotating training data by field experts [CFL10].

In this dissertation, we address the ground truth problem by employing some commonly agreed principles that are indicative of shill behavior. Since the intention of shill bidding and primary ways to achieve shill bidding are understood, we establish a set of rules (defined in Section C of Chapter II and also presented in [DSX09, DSX10b]) to detect shills based on the features of shill bidding. So we consider bidders who violate a sufficient number of rules to be "effective" shills. Based on the commonly agreed principles, some collected auction data are manually investigated and labeled.

Capture of in-auction fraud evidence. As we have discussed in Section F of Chapter I, researchers have summarized several in-auction-fraud-bidding patterns. However, these patterns cannot serve as direct evidence of auction fraud because even if the bidders are confirmed as using these bidding patterns, they still can be innocent. There may be other reasonable, or non-shilling, explanations to the otherwise questionable behaviors. Further quantitative and qualitative analyses of auction fraud are critical for capturing in-auction fraud evidence. Once sufficient auction fraud evidence can be retrieved, the suspicious in-auction fraud could be verified automatically and accurately.

Adaptive anomaly detection. Similar to computer and network security, auction fraud detection faces the difficulty of becoming a battle between fraudulent online auction users and auction integrity researchers. Driven by the opportunity to achieve monetary profits, it can be expected that fraudulent auction participants will not stop fraud practices, but instead change their habitual bidding behaviors to circumvent existing anomaly detection systems. In order to make auction houses trustworthy, it is important to develop adaptive anomaly detection algorithms for capturing in-auction fraud. The adaptive anomaly detection algorithms must be adaptive to new conditions, and thus able to effectively detect and respond to new forms of fraud.

Fraud detection and verification using artificial intelligence techniques. Detection and verification of in-auction fraud requires human knowledge and reasoning capability. This provides motivation for exploring the challenging task of adapting artificial intelligence techniques (e.g., agent-based architectures and reasoning) to represent human knowledge for shill detection and verification. Success in this area would be valuable for achieving the goal of automated detection and verification of in-auction fraud.

III. RELATED WORK

Researchers from economics, system science, computer science and other fields have noticed the severity of the problem of auction-fraud due to shilling, and have made efforts to solve it. In this chapter, we review representative solutions to the problem of shill bidding in the following categories: trust management framework, prediction/prevention approaches, and detection approaches. In addition, some problems and tasks that remain as research challenges in the shill bidding analysis domain are identified. This material is based on the work in [DSX09b].

A. Trust Management Framework

Internet fraud has severely undermined the trust on which members of electronic application communities used to rely. Current existing electronic commerce applications such as online auction systems do not provide completely trustworthy services. There is a high demand for trusted online auction systems that provide trusted, secure and worry-free services.

Xu et al. recently presented an agent-based trust management (ATM) framework for online auctions [XSB08]. The ATM is defined as a multi-agent system [W01] that consists of a security agent, an analysis agent, a set of monitoring agents, auction agents, and bidding agents. Human bidders can specify flexible and complex bidding strategies in the interface of bidding agents so a bidding agent, on behalf of a human bidder, can communicate with auction agents to place bids automatically [RTK90]. Meanwhile, the security agent can dispatch monitoring agents to watch for bidding activities and detect suspicious users, and an analysis agent is responsible for analyzing users' bidding behaviors using live auction data and users' history information. Based on the analytical results, the security agent can re-evaluate a user's trust values in order to verify whether a suspect is a shill bidder. The proposed agent-based trust management module facilitates real-time trust re-evaluation by updating user roles and access permissions dynamically. As a result, the framework provides a solid foundation towards building a trustworthy networked system.

Similarly Yi et al. also applied software agent technology, as well as cryptographic technology, to automate and secure online auctions [YS00]. They presented a secure agent-mediated online auction framework. The framework consists of three components: online auctioneer, auction agent and online bidders. Before the auction starts, the auctioneer advertises an item for auction on the Internet, and all bidders who are interested in the item can register for the coming auction by showing their certificate to the auctioneer. When the auction starts, the online auctioneer generates and launches an auction agent. Following a route specified by the auctioneer, the agent traverses a list of online bidders B_1, B_2,…, B_n on the Internet by showing the auctioneer's certificate. The auction agent informs every bidder of the minimum increment along with the current highest bid for each round, collects bids from bidders and finally brings the bids back to the auctioneer. The procedure is repeated until the highest bid does not increase for three times.

In the agent-mediated online auction framework, public key and private key infrastructures are used to protect the auction from malicious bidders. Each bidder's bid message is first signed by the bidder and then encrypted with the public key of the auctioneer, so the auctioneer can check the authenticity of each bid. Since only the auctioneer can read the bid information and others do not know the private key of the auctioneer, no one can check the bids except the auctioneer. Thus, the framework can prevent malicious bidders from scanning or modifying another bidder's bid. In addition, before an auction agent leaves a bidder, the bidder is asked to update the non-repudiation database and send it to the next bidder. Once the auction agent arrives at the next bidder, the first thing the bidder needs to do is to reply the previous bidder with a signature on the auction agent. Therefore, once any malicious action is detected by the auctioneer, an investigation can be launched by checking the non-repudiation database to discover the malicious user. A significant aspect of cryptography-based online auction fraud detection methods is that the mechanisms of cryptography are able to provide non-repudiated evidence for investigation.

B. Prediction/Prevention Approaches

Preventive measures can be more effective in online auctions than reactive measures. Wily traders usually exploit loopholes left in procedural rules to "attack" honest users and challenge system and mechanism designers. If the auction procedural rules embedded in the software programs of online information technology applications are airtight, fraud activities can be easily prevented, avoided and eliminated.

Wang et al. designed a Shill-Deterrent Fee Schedule (SDFS) mechanism, which could reduce the extra profit brought by shill bidding in the context of independent private value (IPV) English auctions so as to deter opportunistic shills [WHW01]. Under the SDFS mechanism, the auctioneer charges the seller a listing fee and a commission fee. The seller sets only a single starting bid or a reserve price, without the option of setting both a low starting bid and a higher secret reserve price to lure in buyers (as currently allowable on eBay). The listing fee is a function of this initial reserve price, and the commission fee is calculated by the product of the commission rate and the difference between the winning bid and the reserve price. If the reserve price is too high, then the listing fee will be higher and the seller will probably lose the chance to sell the goods. If the reserve price is set too low in an effort to lower the listing fee, then the difference between the reserve price and the selling price will be high, with a correspondingly high commission fee. Therefore, SDFS encourages the sellers to set the reserve prices honestly. The commission rates vary from market to market and are mathematically determined by the online auction systems to guarantee no extra profit for shill bidding compared to honest sale. On the whole, SDFS is reasonable to inhibit shilling behavior.

Some items are not sold in the first round of an auction. This can occur for many reasons, including no bidders having placed bids on the auction, the final price of the auction not reaching the reserve price of the auction, or a seller engaged in a shill and accidentally won the auction. When an item is not sold the first time it goes up for auction, it will typically be offered for resale in a next round. Because there are a significant number of identical auctioned items in the same auction house, a great number of goods are sold after multiple rounds. Wang et al. analyzed shill bidding in multi-round online English auctions, and proved that there is no equilibrium without shill bidding in these auctions [WHW02]. They interpreted the finding as an incentive for shills and suggested a corrective pricing such as SDFS and a fair intermediary should be used to reduce the damage to the market.

Preventing in-auction fraud from happening is possibly the best solution, nonetheless, in some cases, when in-auction fraud cannot be prevented, approaches that can predict its occurrence can also reduce the risk to auction participants. Kauffman and Wood [KW03] examined how the fee structure on eBay may motivate shill bidding and first identified “reserve price shilling” based, in part, on their research into eBay auctions of rare coins in April 2001.They tested whether some questionable bidding behaviors are attributable to reserve price shilling. According to the test results, they built an empirical probit model to predict reserve price shilling based on the seller’s previous behavior before the auction begins.

In addition, researchers have tried to design bidding strategies using game theory in order to help honest users counteract scams [FT90]. Porter and Shoham proposed two equilibrium bidding strategies to counteract bid shading and false bids in sealed-bid auctions, namely first price sealed-bid auction and second price sealed-bid auction

[PS03]. An equilibrium bidding strategy is a Bayes-Nash equilibrium if the bidder's expected gain is maximized when the bidding strategies for all other bidders are fixed. Usually, the expected gain or utility function equals the product of the probability of winning and the difference between a winner's highest willing-to-pay price and the actual winning price. The probability of winning can be estimated by the probability of cheating and the probability that a bidder's highest willing-to-pay price is higher than that of a shill's. Therefore, when knowing the possibility of cheating, equilibrium can be derived to counteract a shill and maximize a bidder's expected gain.

Motivated by Porter and Shoham's work, Jenamani et al. derived an equilibrium bidding strategy for honest bidders to deal with shills in English auctions, and translated the equilibrium strategy into an algorithm called shill counteracting bidding strategy (SCBS) [BJZ05]. By bidding according to the algorithm, honest bidders could counteract shills in English auctions. Experiments are conducted to evaluate the proposed strategy and compare it with five other popular bidding strategies. The average expected utility of the agents with the proposed strategy is found to be the highest when the auction continues for a longer duration. In a later paper, Jenamani et al. showed that both theoretical and experimental results confirm that the equilibrium bidding strategy increases the bidders' expected utility; meanwhile, the authors also explained why English auction is popular over the Internet.

To date, research in fraud prevention and prediction is quite rare. Very few techniques have been proposed for combating in-auction fraud proactively. To prevent in-auction fraud, online auction mechanisms should be improved in order to deter fraudsters from committing fraud. In addition, the underlying information system design should be verified to make sure the properties of the system do not violate any auction mechanism and leave no opportunity for a participant to engage in fraud.

C. Fraud Detection Approaches

1. Using statistical methods

Current Internet auction systems rely solely on feedback based reputation systems to evaluate both buyers and sellers. Nevertheless, the existing traditional reputation system for auction houses has already shown its weakness in providing trusted information. Several researches have shown that the reliability of the reputation system of current auctions house, e.g., eBay, is debatable [CD00, CD03, RKZ00, RZ02]. First, the positive feedbacks are overwhelming but the negative feedbacks are deflated. Deceptive auction users take advantage of the weakness of current rating mechanisms in reputation systems by helping each other artificially build up a good reputation history regardless of their actual behaviors. Rubin et al. found 95% of eBay sellers have good reputation and 98% of their feedbacks are positive [RCG05]. Furthermore, existing reputation systems are easily manipulated. Malicious users could first accumulate a high feedback score by selling low value goods, and then deal high value goods with that good reputation. For example, a seller first sold pencils and gained a good rating. Now the same seller is selling used cars on the same auction site. Can we trust this seller? Probably not. Because the seller could cheat some used car buyers and then shift again to rebuild a reputation from pencil buyers. Moreover, the existing reputation system provides little information about sellers' degree of honesty. Users may find auction fraud information in feedbacks but when dealing with a seller with a long history, it is impractical to look at the feedbacks page by page. Unfortunately, the anti-fraud information has not been directly reflected in the reputation system so far. In all, the current reputation system can no longer satisfy people's need for evaluating trustworthiness in online transactions.

Rubin et al. [RCG05] proposed a new reputation system for auction sites to help users protect their interests by indicating auction fraud. The reputation score in the system is a 3-tuple <*N*, *M*, *P*>, where each variable is a number between 0 and 100 (100 indicates 100% confidence of anomaly, and 0 indicates no signs of fraud). The three variables come from three statistical models: average number of bids model (*N*), average minimum starting bid model (*M*), and bidders' profile model (*P*), respectively. The first model identifies sellers whose auctions, on average, attract more bids than auctions posted by other sellers. In this case, the abnormal situation could be produced either by fierce competition among buyers or by shilling behaviors. The first model does not provide an explanation of the cause for this abnormal situation. The second model, *M*, identifies sellers who have a large number of bids that cannot be explained by their low minimum starting bid (in the statistical model considered by the authors, each starting bid is associated with a number of bids it can attract). Although statistical results show a correlation between minimum starting bid and high volume of bids, it is still not reasonable for the anomalous auctions to attract an overly-high number of bids. Finally, the *P* model identifies anomalous sellers, whose auctions include a group of bidders who bid repeatedly and lose repeatedly as well. The last model explains that the high average number of bids is possibly caused by shill activities. This detection method is indeed a statistics based method.

Trevathan et al. designed an algorithm called *shill score* (SS) to detect the presence of shilling behaviors in online English auctions that have already completed [TR09]. The algorithm targets six very common shill strategies (like the shill indicators described in Section F). By examining each bidder's behavior over auctions hosted by the same seller, the algorithm gives ratings to each bidder based on how the bidder's behavior fits into each of the explicitly defined shill patterns. A bidder's final shill score is calculated in the form of an average of the weighted ratings. The higher the score, the more likely the user is a shill. However, the proposed approach failed in detecting collusive shills, i.e., multiple shills in collaboration with each other. The collusive strategies used by shill groups are much more complicated and sophisticated than the single shill strategies. Collusive shill could thwart the SS algorithm by normalizing the shill group's shill score. To address this problem, Trevathan et al. extended the SS algorithm and proposed a new algorithm [TR07] named the "collusion score" to detect collusive shills controlled by one seller. They analyzed three kinds of collusive bidding strategies that could be adopted by shill bidders: (1) alternating bid strategy, in which shills bid alternatively in the same auction; (2) alternating auction strategy, where different shills bid on different auctions and each shill bids exclusively in one auction; (3) hybrid strategy, which is the most complicated one that combines the first two strategies. Collusion graphs are utilized to examine shills in terms of the above three identified collusive shill strategies. Combining ratings of each examination, the collusion score is assigned to each bidder, indicating the likelihood that the bidder is engaging in collusive shilling behaviors. The situation where multiple sellers work in collaboration to do shilling is even more complicated than collusive shill controlled by one seller, thus, the proposed collusive score approach is not suitable for this case.

In work that is part of this dissertation, we identified and summarized a series of shill bidder properties and normal bidder properties [DSX09]. This work is discussed in Chapter II. Since each property involves uncertainty in incriminating auction shills, they employed the mathematical theory of evidence, called *Dempster-Shafer theory*, to combine evidence in order to reduce the uncertainty. Meanwhile, the conflict between evidence is also measured. A degree of belief is assigned as a shill score in order to quantify the likelihood of certain bidders being shills. Experiments showed that the approach has potential to be effective in reducing the number of false positives generated by any single piece of evidence.

The aforementioned reputation models of detecting in-auction fraud essentially make use of statistical methods. The model proposed by Rubin et al. can be used to test the statistical significance of how far the tested auctions are away from the benchmark auctions, and then to hypothesize that the statistical anomalies represent shills. The *SS* method statistically measures how the bidders' behaviors fit into the shill patterns, and calculates a score indicating the likelihood of shill. The Dempster-Shafer theory approach that we have proposed aims to reduce the uncertainty and resolve the conflicts between evidence that is used to incriminate shills. Note that most statistical methods have to analyze a large amount of data, where auction data must be carefully selected for comparison because unreliable benchmark auctions can decrease the statistical significance of differences, thus compromising the accuracy of the results.

2. Using data mining methods

Data mining (also called knowledge discovery) is a powerful computer-assisted process designed to analyze and extract useful information from historical data [HK01]. It allows users to analyze data from different dimensions or perspectives in order to uncover consistent patterns, anomalies and systematic correlations between data elements. The ultimate goal of data mining is to predict future behaviors and trends based on the discovered patterns and association rules. Several researchers have adopted data mining methods to detect shill associations and suspicious patterns.

Ford et al. presented a real-time self-adaptive classifier, called RT-SAC, which is capable of accurately classifying and smoothly adapting to new data [FXV12]. In addition, by utilizing a Prolog engine and a separate classification module, the classifier can be easily configured at runtime. The RT-SAC approach can greatly increase the efficiency and effectiveness for real-time shill detection in online auctions. Pandit et al. designed and implemented an online auction fraud detection system named NetProbe [CPF06, PCW07, ZZF08]. The key idea of the NetProbe is to infer properties of a user by properties of other related users. In particular, given a graph representing associations between auction users, the likelihood of a user as a fraudster is inferred by looking at the behavior of the user's immediate neighbors. The NetProbe system models auction data as a network graph in which sellers and bidders are represented by nodes, and transactions are represented by edges between sellers and bidders. Markov random field and belief propagation algorithms are utilized to unearth suspicious trading patterns created by fraudsters, and thus to detect possible fraudsters. In addition, to deal with the dynamic nature of online auction data, an incremental version of the NetProbe has also been proposed [PCW07]. The motivation behind incremental NetProbe is that the addition of new edges in the graph will not affect the whole graph, but only leads to minor changes to the immediate neighborhood of the edge. Properties of affected nodes are updated incrementally without wasteful re-computation. Experiments using large synthetic and eBay data sets demonstrated that the NetProbe was effective with high accuracy, and the incremental NetProbe had significant speedup in execution time with negligible loss of accuracy. It is worth noting that the NetProbe does not need to treat single shill and collusive shill separately; instead, it could detect compliances together with shill.

Shah et al. applied data mining techniques in detecting shill behaviors in eBay video game console auctions [SJS03]. They mined associations between buyers and sellers during a period of time and found that some users only bid in auctions hosted by one particular seller and seldom won, e.g., once or twice in all the transactions in which they were involved. They considered these as possible cases of shill bidding.

Data mining approaches, like reputation approaches, also require analyzing huge amounts of historical data, and therefore take a very long time to get results. Although incremental NetProbe can reduce the execution time to almost half of the original, it still cannot achieve real-time performance so far. As a tradeoff, data mining approaches do have the advantage of accuracy compared to other approaches.

3. Using formal methods

Formal methods use formal notations and logic, which are mathematically rigorous techniques, for the specification, development, and verification of software and hardware designs [CW96]. Model checking is one of the many automated formal methods that are used to check the correctness of a system. Specifications about a system are expressed in the form of logic formulas, and efficient algorithms are used to verify certain properties of a system by means of an exhaustive search of all possible states that it could enter during its execution. There are also researchers making use of formal methods to detect in-auction fraud.

Xu et al. introduced a formal model checking approach to detect shilling behaviors, especially the competitive shilling behaviors [XC07]. They first derive a formal auction model from real world auction data according to a predefined auction model template. The model is then verified using the SPIN model checker for behavioral properties, which are specified in pattern-based LTL (Linear Temporal Logic) formulas. The formulas are translated from some of the aforementioned shill strategies. Experiments showed that the proposed formal approach is feasible and efficient. Recently, this work has been extended for real-time detection of auction shills by defining a dynamic auction model (DAM) [XBS09]. Shilling behaviors in different stages of an auction, namely early stage, middle stage and final stage, are formally specified in LTL, and verified on the DAM using real-time model checking technique in order to discover shills suspects.

The primary advantages of the model checking solution are accuracy and the potential to support detection of shilling behaviors in real-time. The mathematical rigor of the model checking technique can explain the accuracy of the solution. Moreover, the authors used the model checker to verify bidders' behavior at the bid level rather than the auction level, so when a shill bids on an auction, the shill's behavior could not pass the check so can be detected immediately.

The proposed real-time model checking approach is very promising because it can possibly detect in-auction fraud before any payment and thus prevent most of the monetary loss for victims. Once the fraud is verified, the system can notify all participants immediately, and may also suspend or cancel the auctions.

Sections A, B, and C discussed three different types of solutions to the problem of in-auction fraud, namely trust management frameworks, prediction/prevention methods, and detection methods. In Table IV, these solution techniques are compared based on some performance issues such as time-efficiency and data requirements.

TABLE IV.

IN-AUCTION FRAUD SOLUTION TECHNIQUES

Solution Techniques		Real-Time	Historical Data
Trust Management Framework	ATM	Supported	Required
	Cryptography-based	Supported	Not required
Prediction or Prevention Approaches	Equilibrium Bidding Strategy	N/A	Not required
	Shill-Deterrent Fee Schedule	N/A	Not required
Detection Approaches	Statistical Methods	Not supported	Required
	Data Mining	Supported	Required
	Formal Methods (Model Checking)	Supported	Not required

D. Auction Data Processing

A substantial amount of work has been done in the study of auction data. Heijst et al. [HPW07] combined text mining and boosting algorithms to predict auction final prices. Ghani and Simmons [GS04] compared a regression model, a neural network and a decision tree, and they achieved the best result using the neural network when treating the price prediction problem as a series of binary classification problems. Lim et al. [LAH08] employed grey system theory to predict auction closing prices in a simulated auction environment. Different from the above approaches, we use a neural network approach, but for the special purpose of predicting the likelihood of auctions involving shilling activities. In particular, the expected auction price is learned from the Large Memory Storage and Retrieval (LAMSTAR) Neural Network [G07], where price prediction is based on features extracted from item descriptions, listings and bid properties.

E. Empirical Research

The body of empirical research on online auctions is growing. Roth and Ockenfels [RO02] had observed the prevalence of "bid sniping," which is a strategy that helps bidders counteract shill bidding by avoiding bidding until the final moment of an auction Lucking-Reiley [L99] and Bajari and Hortacsu [BH03], based on empirical investigations, concluded that a lower starting price can bring more bidders into an auction and may increase the final price of an auction. Kauffman and Wood [KW03] developed a probit model to study the factors that can be used to predict reserve price shilling. They found book value and starting bids are indicative of reserve price shilling. Unlike previous research, we propose to examine the relationship between shilling behavior and final auction prices using logistic regression, as presented in Chapter VI.

IV. REASONING UNDER UNCERTAINTY FOR SHILL DETECTION IN ONLINE AUCTIONS USING DEMPSTER-SHAFER THEORY

This chapter discusses a novel method of shill detection and the handling of the uncertainties associated with the evidence. This chapter is based on work in [DSX09, DSX10b].

A. Uncertainties of Shills in Online Auctions

Shill bidding usually occurs without leaving obvious direct physical evidence, thus it cannot be easily captured by the victims. Kauffman and Wood examined the effects of shill bidding on the final bidding price in rare coin auctions, and showed that some bidders might view shill bids as signals that an item was worth more, thus they would be likely to pay more than other bidders who could not see the signals [KW05].

Many shill bidding strategies or patterns have been identified in order to help investigate auction frauds. However, most of the previous findings involve uncertainties. For example, we might identify the following shill pattern in real auction data: “When a bidder tends to place bids in an auction with a higher current bidding price than the current price in a concurrent auction with an identical auctioned item, the bidder might be a shill” [XC07]. This is not a certain rule for shill detection because it is also possible that some experienced buyers may prefer highly rated sellers with better reputation for quality of service to lower-rated ones, even at the cost of a higher payment. Furthermore, to support automated detection of shill bidders, we need a consistent approach to representing and quantifying auction and bidder related knowledge. For instance, by calculating and analyzing a bidder’s winning ratio, we may have a better idea about the bidder’s actual bidding intention – if a bidder wins often, the bidder is not likely a shill because a shill typically avoid winning auctions.

To address these problems, in this chapter, we propose a decision support approach to certifying bidders’ behaviors immediately after each auction’s bidding cycle, but before the auction is officially closed. Similar to resolving a criminal case, we first collect evidence that supports a bidder being a shill as well as evidence that supports a bidder being an honest bidder. Since each piece of evidence involves uncertainties, it is appropriate to employ some formal reasoning technique [AAP03]. In this context, we propose to use belief functions in Dempster-Shafer (D-S) theory [S76, D68] to model the uncertainties associated with different pieces of evidence pertaining to varied bidding properties. This allows us to explicitly represent the uncertainties and combine knowledge from different sources of evidence to produce an aggregated assessment. Based on the assessment, a certification is issued to each bidder, which can assist both auction houses and auction participants in deciding the trustworthiness of bidders. This work extends our previous proposed framework of using Dempster-Shafer theory for shill detection [KW03]. The major extensions are as follows. First, auction-level properties and evidence are introduced, which complements the bid-level evidence for providing a macro-examination of auctions. Second, in the previous work, the focus was only on quantifying the degree of belief concerning if a bidder is a shill, rather than considering both cases – if a bidder is a shill or is not a shill. In this dissertation, each piece of evidence is determined to support either shilling behavior or normal bidding behavior based on the analysis of the corresponding quantified bidding property. Algorithms for calculating the belief of being not a shill as well as belief of being a shill are presented in this extended work. Therefore, the scope of the decision support system is significantly improved.

B. Introduction to Dempster-Shafer Theory of Evidence

Dempster-Shafer (D-S) theory of evidence is a mathematical theory that was developed by Dempster and Shafer in 1976 as a new approach for representing uncertainties and expressing conflict involved in a set of evidence [S76]. D-S theory has often been used to combine information (evidence) from different sources to calculate the probability of an event. Generally, D-S theory differs from traditional probability theory in that the former allows the explicit representation of ignorance and uncertainties in the evidence combination process. Furthermore, D-S theory allows assigning a probability to not only singletons but also a set of multiple alternative elements [D68]. These unique characteristics make D-S theory particularly attractive to designing and implementing complex systems. In this section we highlight some of the key concepts of D-S theory, including some examples from our domain of interest, shilling behavior in auctions.

The belief distribution of the D-S theory is based on a *universe of discourse* Θ (also called *frame of discernment*) that consists of a finite set of mutually exclusive atomic states in a problem domain [S76]. For example, in the auction shill detection domain, the frame of discernment for a bidder is $\Theta=\{shill, \sim shill\}$. The power set 2^{Θ}, which is the set of all possible subsets of Θ including the empty set, can be denoted as $2^{\Theta} = \{\emptyset, \{shill\}, \{\sim shill\}, \Theta \}$.

There are three important functions in D-S theory: Basic Mass Assignment function, Belief function, and Plausible function [D68]. The Basic Mass Assignment (BMA) function is m: $2^{\Theta} \rightarrow [0,1]$. It assigns a belief mass in the interval between 0 and 1 to each subset of the power set. The belief mass represents the impact of a piece of evidence to the subset of 2^{Θ}. The BMA function should verify the following two equations:

$$\sum_{A \in 2^{\Theta}} m(A) = 1 \qquad (1) \qquad\qquad m(\varnothing)=0 \qquad (2)$$

The empty set Ø represents a contradiction, which cannot be true in any state. Therefore, the BMA for Ø is assigned 0. The basic mass assignment $m(\Theta)$ can be interpreted as the measurement of conflict (in our application both states of shill and ~shill are present) and a mass is computed for the conflict. For the shill detection problem, Eq. (1) and Eq. (2) imply that $m(shill) + m(\sim shill) + m(\Theta) = 1$.

To obtain the overall belief of A, one must take the sum of beliefs on all subsets of A. As defined in Eq. (3), a *belief function* is defined as the mass sum of all Bs, which are subsets of A.

$$bel(A) = \sum_{B \subset A} m(B) \qquad (3)$$

D-S theory allows a belief of a subset of 2^{Θ} to be represented by intervals, bounded by belief and plausibility [D68] – for example, $bel(\{shill\}) \leq P(\{shill\}) \leq pl(\{shill\})$. The plausibility of A specifies the likelihood that it is not any other subset in 2^{Θ}. The quantity of plausibility of A is equal to one minus $bel(\sim A)$, that is $Pl(A)=1-bel(\sim A)$. For example, the degree of plausibility for shill is: $Pl(\text{shill}) = m(\{\text{shill}\}) + m(\{\text{shill}; \sim\text{shill}\})$. According to Eq. (3), it is easy to derive that the quantity of plausibility of A is equal to the sum of the masses of B, whose intersection with A is not empty, as shown in Eq. (4). For all $A \in \Theta$, $bel(A)$ forms a lower bound for A that could possibly happen, and $pl(A)$ forms an upper bound for A to happen, given by formula (5).

$$pl(A) = \sum_{B \mid B \cap A \neq \varnothing} m(B) \qquad (4)$$

$$bel(A) \leq P(A) \leq pl(A) \qquad (5)$$

Knowing any of the three functions m, bel, and pl, the other two can be deduced using Eq. (3) and Eq. (4) [S76].

Given independent belief functions over the same frame of discernment, we can combine the beliefs into a common agreement concerning a subset of 2^{Θ} and quantify the conflicts using Dempster's rule of combination

[S76]. Given two masses $m_1$and m_2, this combination computes a *joint mass* for the two pieces of evidence under the same frame of discernment. It is calculated as follows:

$$m_{1,2}(A) = (\sum_{Y_1 \cap Y_2 = A} m_1(Y_1)m_2(Y_2))/K \qquad (6)$$

where $K = 1 - \sum_{Y_1 \cap Y_2 = \varnothing} m_1(Y_1)m_2(Y_2)$

Note that K represents the renormalization factor, which is equal to one minus the amount of the conflicts between two masses pertaining to the subset A of the frame of discernment.

The combination rule is usually denoted as the orthogonal sum of belief values; in other words, the combination of belief from evidence a and belief from evidence b is denoted as $bel_{a,b}=bel_a \oplus bel_b$. Therefore, the global belief of A can be represented as $bel(A) = \oplus\ bel_i$, for all pieces of evidence that supports A.

To illustrate the concepts, consider a subset X of 2^Θ and evidence E_1 that yield a set of values represented by $m_{E1}(\{x\})$, $m_{E1}(\{\sim x\})$, and $m_{E1}(\{x,\sim x\})$. Suppose that evidence E_1 may provide, in general, some support that X is true, i.e., event x occurs, or some support that X is not true, i.e., event $\sim x$ occurs. In terms of the mass function, the BMAs for x and $\sim x$ are m_{E1} $(\{x\})$ and m_{E1} $(\{\sim x\})$, respectively. Lack of knowledge about whether x occurs or not is represented by m_{E1} $(\{x,\sim x\})$. The sum of the three values is one. i.e., $m_{E1}(\{x\}) + m_{E1}(\{\sim x\}) + m_{E1}(\{x,\sim x\}) = 1$. We can further assume that evidence E_1 is either reliable with probability 0.9 or unreliable with probability 0.1. Now, using our shilling behavior example, suppose that evidence E_1 supports that bidder i is a shill with 100% certainty. Considering E_1's reliability, E_1 gives 0.9 degree of belief for supporting that bidder i is a shill (i.e., $m_{E1}(\{x\}) = 0.9$), but zero degree of belief that bidder i is honest (i.e., $m_{E1}(\{\sim x\}) = 0$) because the evidence does not support bidder i is honest. The remaining degree of belief (0.1) is due to the uncertainty, i.e., $m_{E1}(\{x, \sim x\}) = 0.1$.

C. Shill Detection Under Uncertainty

1. An abstract model

Our proposed approach can be defined as an abstract model with 5-tuple <B, bel, P, M, R>, where

1). B = $\{b_1, b_2, \ldots b_n\}$ is a set of online auction bidders to be certified;

***2). bel*:** $B \rightarrow [0, 1]$ is a scoring function**.** There is a degree of belief for every online auction bidder, representing the system's belief that a bidder is a shill or not.

3). P = $\{p_1, p_2, \ldots, p_k\}$ is a set of bidders' properties, which can be considered as evidence either for shilling behaviors or normal bidding behaviors.

4). M = $\{m: P \rightarrow [0,1]\}$ is a set of mass assignment functions which quantify every piece of evidence into a mass that supports either *shill* or *~shill.*

5). R = $\{\theta, \varphi\}$ is the set of thresholds for making decisions on a bidder's certifications, where $\theta < \varphi$. The first element θ is the belief value threshold for determining if a bidder is a trusted bidder. If the value of $bel(shill_i)$ is below θ, the $bidder_i$ will be certified as a *Trusted Bidder*. The second element φ is the belief value threshold for determining shills, and it is larger than 0.5. If the value of $bel(shill_i)$ exceeds φ, the $bidder_i$ will be certified as a *Shill*. For any bidder, if the shilling score is between θ and φ, and *bel*(*shill*) is greater than or equal to *bel*(*~shill*), the certification of the bidder would be updated to *Suspect*.

Certifying a group of bidders B is to assign every bidder b_i in B a role to indicate the bidder's trustworthiness, i.e., deciding if a bidder is a *shill*, a *suspect*, or a *trusted bidder*. For any bidder's property $p_i \in P$, it can be utilized to either support a bidder is a legitimate bidder or a shill, depending on the nature of the property and the quantified value of the evidence. For example, if a bidder placed quite a few abnormal concurrent bids in a sellers' auction, it becomes evidence to support that the bidder is a shill. However, if a bidder places very few abnormal concurrent bids in online auctions, it should be considered as evidence to support the bidder is not a shill. Each property can only support a bidder for one state but not both. At the auction level, the decision boundary can be the average level of all auctions in the same category. If an auction's property value is significant (i.e., not close to the average value), the corresponding evidence can be considered as one to support that the auction involves shilling behavior or the auction is normal, depending on the value and the nature of the property. For example, suppose auctions in a certain category attracted 7.67 bids on average in the past 30 days. Now if an auction ends with 60 bids, we may consider that the auction involves shills. Since each property p_i for a bidder can only support one state, the rest of the belief from property p_i cannot commit to another state other than the universal set, i.e., the *frame of discernment.* Intuitively, the universal set, e.g., {*shill*, *~shill*}, can be interpreted as *uncertainty* about any state. The ability to represent and quantify uncertainties is a key advantage of Dempster-Shafer theory. The BMA for the evidence that corresponds to property p_i in supporting shilling behavior can be represented as in Eq. (7), Eq. (8), and Eq. (9).

$$m_{pi}(shill) = \alpha * f \tag{7}$$

$$m_{pi}(\sim shill) = 0 \tag{8}$$

$$m_{pi}(U) = 1-\alpha * f \tag{9}$$

where $0 < \alpha < 1$, and it is an adjusted value that can be understood as the strength of property p_i on determining if a bidder is a shill. The function f quantifies evidence for shill certification, where $0 < f \leq 1$.

The BMA for the evidence that corresponds to property p_j in supporting normal bidding behavior can be represented using Eq. (10), Eq. (11), and Eq. (12).

$$m_{pj}(shill) = 0 \tag{10}$$

$$m_{pj}(\sim shill) = \beta * g \tag{11}$$

$$m_{pj}(U) = 1-\beta * g \tag{12}$$

where $0<\beta<1$, and it is an adjusted value that is the strength of property p_j on determining if a bidder is not a shill. The function g quantifies evidence for supporting the bidder is honest, where $0<g \leq 1$.

2. The shill certification framework

An automated shill certification system can play a significant role in maintaining trust among online auction users. The major task of our proposed shill certification system is to identify shills and recognize honest bidders. Figure 4 depicts the shill certification framework based on D-S theory.

Auction bidders are certified mathematically using a data fusion method that combines information from different aspects of bidders' behaviors and auction-level features. The certification process classifies all bidders into categories reflecting the likelihood of an "actual" shill. Initially, every bidder in the auction house is categorized as a *Trusted Bidder*. When the bidding process of an auction ends, the auction enters a shill certification stage, and the auction does not officially close until the certification procedure is complete. After an auction is officially closed with valid certifications for all bidders, the seller and the winner of the auction can proceed for further activities such as payment, shipping, and mutual feedback.

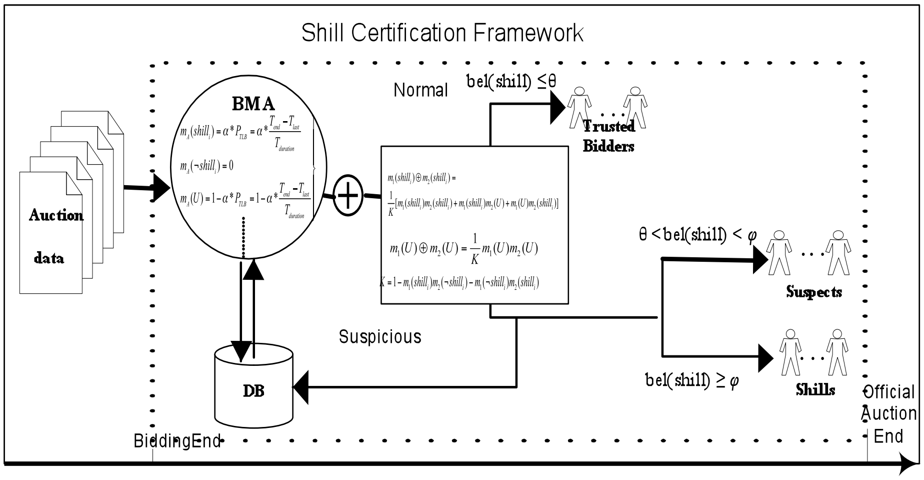

Figure 4. Shill certification framework

In the certification process, the monitored auction data along with historical statistical data stored in a database is the input to the basic mass assignment (BMA) module (as shown in Figure 4). Each bidder's behavior is checked and quantified based on the formulas that will be defined in Section D.1. If the BMA module does not have sufficient belief to support a bidder is a shill, the bidder is classified as a *Trusted Bidder* immediately. Note that this piece of information is stored into the historical database for future use. On the other hand, if a bidder's bidding behavior shows any shill bidding properties, the pieces of evidence obtained from BMA module are combined together using Dempster's rule [CV05] (denoted as ⊕ in Figure 4). The results of the evidence combination process are belief values that indicate the likelihood of being a shill or a normal bidder. Needed information for computing the belief values, such as the other evidence for the same bidder in the same auction, can be fetched from the database. Once the belief values are calculated, the certification system updates each bidder's certification according to the certification assignment rules as shown in Figure 5.

$bel(shill) \geq \varphi$ => Shill
$\varphi > bel(shill) > \theta$ => Suspect
$bel(shill) \leq \theta$ => Trusted Bidder

Figure 5. Certification assignment rules

The threshold (φ) of certifying bidders as *Shill* should be fairly high to reduce the number of false positives. For the bidders that are certified as *Suspect*, the values of their *bel(shill)* must be lower than φ but still greater than the values of *bel*(~*shill*). This means that the evidence is not sufficient enough to support a bidder for being a shill, even though the bidder behaved more like a shill than an honest bidder. As a result, the bidder is assigned a certification of *Suspect*. When any additional independent evidence is available, the certification of *Suspect* shall be validated again. If a bidder's certification changes, the new certification is committed to the database. If a bidder's certification is labeled as *Shill*, the bidder is subject to further investigation and possible punishment, but the shill-handling step is outside the scope of this dissertation.

The required statistical data for computing the basic mass assignment includes information such as the bidder's total number of bids in a certain period of time, and the number of bids in the particular auction. Such information is stored in a historical database, and can be fetched when needed. The database is updated periodically when each certification process completes.

3. Shill related properties: bid-level and auction-level

There are two types of properties that can be used to provide evidence of shilling behaviors – those properties associated with a particular bidder, such as the time when he placed his last bid, and those properties associated with the auction itself, such as the total number of bids in the auction. The auction-level properties can be used as evidence to support that an auction involves shills; while the bid-level properties can be used as evidence to support that a bidder is a shill. Note that if an auction is suspected of involving shills by the evidence at the auction level, every bidder in the auction is considered as a shill suspect initially. In other words, if an auction-level property is used as evidence for shill, it is used as evidence supporting shill for every bidder in the auction. On the other hand, if the property is used as evidence for not shilling, all bidders in the auction get one more piece of evidence to support that they are honest. Therefore, when combining the evidence at the auction level with the evidence at the bid level, the auction-level evidence is just used as a piece of bid-level evidence.

In order to demonstrate the feasibility of our approach, we define several bidding properties and auction features that are used in our case study. Note that the list of possible properties we provide is not necessarily complete; providing such a complete list is beyond the scope of this dissertation.

4. Bid-level properties

Property TLB (Time of Last Bid). The time a bidder places his last bid in an auction can reflect the genuineness of the bidding purpose. Generally speaking, shill bidders typically avoid placing bids in a later stage of an auction in order to reduce the risk of winning. In other words, a bidder who places a bid in the late stage of an auction is more likely an honest bidder who aims to win the auction. This evidence can support either shill or normal bidder, depending on the relative time at which a bidder places his last bid. We quantify the relative time of such bids by Eq. (13).

$$TLB = \frac{T_{end} - T_{last}}{T_{duration}} \tag{13}$$

where T_{end} is the end time of the auction; T_{last} is the time when the monitored $bidder_i$ places his last bid; and $T_{duration}$ is the duration of the auction. Thus, in terms of this particular evidence, the likelihood of a bidder being a shill increases as P_{TLB} increases. The earlier such a (last) bid is placed in an auction, the more suspicious is the bidder. When the last bid is placed in the final stage of the auction (we define it as $[0.9T_{end}, T_{end}]$, following the definition in [XBS09]), this information can be considered as evidence to support that the bidder is honest.

Property CBA (Concurrent Bid Activity). Shill bidders are not bargain hunters, while most legitimate bidders are. Because shill bidders' purpose is different from that of legitimate bidders, shill bidders typically do not favor items with lower prices. They may place bids in an auction that has a higher current bidding price rather than in some concurrent auctions that have *lower* current bidding prices [XC07]. We consider bidders placing such abnormal bids as candidates of shills. We capture this indicator of shilling behavior as the percentage of abnormal concurrent bids (*ACB*) placed by a bidder with respect to a particular seller as given by Eq. (14).

$$CBA = \frac{ACB_{(i,j)}}{\sum_{j=0, j \neq i}^{n} ACB_{(i,j)}} \tag{14}$$

where $ACB_{(i,j)}$ is the number of abnormal concurrent bids that $bidder_i$ has placed in auctions hosted by $seller_j$.

Property AF (Average Feedback): A feedback score is an indicator of an online auction user's reputation, which can be used as a predictor for the user's future behavior. Generally speaking, a positive rating increases a user's feedback score, and a negative rating decreases the feedback score. Since a high feedback score is important for a bidder to gain trust from sellers and other bidders, auction users try to maintain good reputations, which may possibly have been accumulated over a long period of time. Therefore, users with good feedback histories normally would not take the risk of being a shill. On the other hand, because shills seldom win auctions, they do not accumulate much feedback. This source of evidence can be quantified by comparing a bidder's feedback score against the average feedback (*AF*) score of all users in the same category. When a bidder's feedback score is less than the average, the evidence to support that a bidder is a shill is calculated by Eq. (15.1). For the same reason, if a bidder's feedback score is greater than or equal to the average, property *AF* shall be considered as evidence for normal bidder. In this case, the mass for it should be calculated according to Eq. (15.2).

$$AF = 1 - \frac{FB_i}{FB_{avg}} \qquad (15.1) \qquad\qquad AF = 1 - \frac{FB_{avg}}{FB_i} \qquad (15.2)$$

where FB_{avg} is the average feedback score of all bidders in the same category of auctions and FB_i is the feedback score of $bidder_i$.

Property BIA (Bidding Increment Activity): The minimum increment of an auction is usually set up by an auction house before the auction begins. Within different price ranges, the auction house requires different minimum increments. Typically, the minimum increment increases as the price range level grows. If a bidder wants to outbid another buyer, the bidder must place a bid that at least equals the current price plus a minimum increment. Generally speaking, a bargain hunter usually tries to place bids with the minimum increment so as to win the item at a price as low as possible, while a shill often adds a large increment in an attempt to raise the price quickly. As we have observed, shills tend to place such a larger increment at an early stage of an auction, and before the price reaches the normal price range, where the early stage of an auction is defined as the first quarter of the auction time, following the definition in [XBS09]. Typically there are only a few bids placed, but a shill bidder may be eager to drive up the price as early as possible. This is because if a shill places quite a large bid at the final stage of an auction, the price would likely drive potential buyers away and may cause the shill to become the auction winner. This is the worst situation, which shills try to avoid. When quantifying the evidence, we compare the average increment a bidder placed with the minimum increment, and only the increments placed before the final stage of an auction are considered. Recall that the final stage of an auction is defined as [0.9T, T], following the definition in [XBS09].

$$\mathrm{BIA} = \frac{\sum_{l=0,t=0}^{s,0.9T} \frac{m * MIN_l}{\sum_{i=1}^{m} I_i}}{s} \qquad (16)$$

where I_i is a bidder's i^{th} increment, s is the total number of different minimum increments in s different price ranges, m is the total number of increments the bidder added in a price range, and MIN_l represents the minimum increment of the l^{th} price range defined in the auction rules. For example, an auction house sets the minimum increment rule as follows. When the current price is between \$0.01 and \$0.99, the required minimum increment is \$0.05; the minimum increment is \$0.25, when the current price is between \$1.00 and \$4.99; when the current price is between \$5.00 and \$24.99, the required minimum increment is \$0.50, and so on. Suppose the final price of an auction is \$24.99 and the average increments that a bidder placed in the three different prices ranges are \$0.50, \$0.9, and \$2.00, respectively. Then the value of *BIA* for the bidder is calculated as the following:

$$BIA = \frac{\frac{0.05}{0.5} + \frac{0.25}{0.9} + \frac{0.5}{2.00}}{3} = 0.21$$

This result shows that the bidder might be a shill. The likelihood of a bidder being a normal bidder increases as the value of *BIA* decreases. But, when a bidder's increment exceeds a pre-specified value, the *BIA* property would be used as evidence to support that the bidder is a shill.

Property WPB (Wins Per Bid). Normal bidders usually win several auctions during a period of time, while shill bidders do not win. This property can be measured by Wins Per Bid (*WPB*). We compare the bidder's *WPB* value for a specific bidder, $WPB_{(i,j)}$ as shown in Eq. (17.1), with the bidder's overall *WPB* value as given by Eq. (17.2).

$$WPB_{(i,j)} = \frac{NOW_{(i,j)}}{NOB_{(i,j)}} \quad (17.1) \qquad WPB = \frac{\sum_{j=1, j \neq i}^{j=n} NOW_{(i,j)}}{\sum_{j=1, j \neq i}^{j=n} NOB_{(i,j)}} \quad (17.2)$$

where $NOW_{(i,j)}$ is $bidder_i$'s number of wins in the auctions that were hosted by $seller_j$, and $NOB_{(i,j)}$ is the number of bids that $bidder_i$ has placed in $seller_j$'s auctions. If the bidder's $WPB_{(i,j)}$ is lower than *WPB*, this information will be used as evidence to support that the bidder is a shill. Otherwise, this is a source of evidence that supports the bidder is not a shill.

Property AS (Affinity for Sellers). Shills usually have a close affinity for a particular seller. A normal bidder may place bids in different sellers' auctions, while a shill tends to participate in a great number auctions conducted by a particular seller who may have collaboration with the shill. The degree of abnormality of a bidder's bid activity is quantified by the percent of participation for a seller's auctions, given by Eq. (18). The abnormality increases as the bidding frequency grows.

$$AS_{(i,j)} = \frac{AUC_{(i,j)}}{NOA_j}$$

(18)

where $AUC_{(i,j)}$ is the number of auctions hosted by $seller_j$ and participated in by $bidder_i$, and $NOA_{(i,j)}$ is the total number of auctions hosted by $seller_j$ in a certain period of time. If a bidder's $AS_{(i,j)}$ score is low, the *AS* property will be used as evidence to support that the bidder is a normal bidder. Otherwise, this is a source of evidence to support that the bidder is a shill.

5. Auction-level properties

The properties specified previously are often considered as bid-level properties. Next we specify shill-relevant properties from the auction level. These properties can either support auctions with shill bidding or support valid auctions.

Property NB (Number of Bids): Auctions with shills usually end with more bids on average than those without shills. When comparing auctions of similar duration and items for bid, shills tend to outbid the legitimate bids frequently until the price reaches their expected value, or when the risk of winning the auction becomes high. The bids that the shills stimulated and placed contribute to the extra amount of bids in the auction. If the number of bids in an auction is more than the average number of bids in the same category, there is a chance that the auction involves shills. In this case, Property *NB* can be quantified as in Eq. (19.1). Similarly, if the number of bids in an auction is smaller than the average, the possibility that the auction employed shill bids is reduced. In fact, the fewer bids an auction has, the more possibility the seller is honest and did not employ shills. In this case, the degree of the auction being normal is given by Eq. (19.2).

$$NB = 1 - \frac{NB_{avg}}{NB_k} \qquad (19.1) \qquad\qquad NB = 1 - \frac{NB_k}{NB_{avg}} \qquad (19.2)$$

where NB_k is the number of bids that are placed in the monitored auction and NB_{avg} is the average number of bids that were placed in an auction of the same product over the last 30 days.

Property SP (Starting Price): The starting price (*SP*) of an auction that involves shills is usually less than the average starting price of auctioned items in the same category. In other words, the higher the starting bid (compared to book value), the less possibility that the auction involves shills. Conversely, if the starting bid is much less than the book value, it is more likely that the auction involves a shill. This indicator is explained and tested in [XSB08]. In auction houses, the commission fee is partly based on the starting price. So by lowering the starting price of an auction, sellers can save money on the commission fee. If an auction's starting price is higher than average, this might indicate that the seller has no intention to engage in shilling. Because, otherwise, if the seller planned to use shill bidding, he would not have to start the auction at a higher price, which incurs a higher listing fee, i.e., a part of the commission fee. Therefore, the higher the starting price, the more possibility the auction is normal, as given in Eq. (20.1). In contrast, if an auction's starting price is lower than average, it is possible that the auction involves shills. The property to support a shill auction is given by Eq. (20.2)

$$SP = \frac{SP_{avg}}{SP_k} \qquad (20.1) \qquad\qquad SP = \frac{SP_K}{SP_{avg}} \qquad (20.2)$$

where SP_k is the starting price of the monitored auction, and SP_{avg} is the average starting price for the same product over the last 30 days.

Since any of the properties mentioned above involve uncertainties, we now propose to employ D-S theory to reduce the uncertainties and conflicts in incriminating shills.

D. Shill Certification

1. Basic mass assignment

As we mentioned earlier, the shill certification process employs a mathematical theory, D-S theory of belief functions, to represent the uncertainties of evidence pertaining to different hypotheses. Consider the states {*shill*},

{~*shill*}, and Θ. We now provide basic mass assignments (BMA) for evidence *TLB, CBA, WPB, BIA, AF, AS, NB, and SP,* which were described in Section C.

BMA for Evidence TLB: The basic mass as in (21.1) is assigned to ~*shill* only if a bidder places his last bid in the final stage of an auction. In this stage, a shill only places bids occasionally and very carefully in order to avoid winning the auction. When a bidder is detected to have a tendency to stop bidding earlier, the bidder should be identified as a shill candidate and the mass is assigned to *shill*, as given by Eq. (21.2).

$$\left.\begin{array}{l} m_{TLB}(shill) = 0 \\ m_{TLB}(\sim shill) = \beta_{TLB} * (1 - TLB) = \beta_{TLB} * (1 - \frac{T_{end} - T_{last}}{T_{duration}}) \\ m_{TLB}(\Theta) = 1 - \beta_{TLB} * (1 - \frac{T_{end} - T_{last}}{T_{duration}}) \end{array}\right\} \quad (21.1)$$

$$\left.\begin{array}{l} m_{TLB}(shill) = \alpha_{TLB} * TLB = \alpha_{TLB} * \frac{T_{end} - T_{last}}{T_{duration}} \\ m_{TLB}(\sim shill) = 0 \\ m_{TLB}(\Theta) = 1 - \alpha_{TLB} * \frac{T_{end} - T_{last}}{T_{duration}} \end{array}\right\} \quad (21.2)$$

BMA for Evidence CBA: When a bidder passes up many chances to place a lower bid for an item in a concurrent auction, the bidder is considered as a shill candidate. On the other hand, if a bidder places very few abnormal concurrent bids, *CBA* is used as evidence to support the state of ~*shill*. Eq. (22.1) and Eq. (22.2) provide the basic mass assignments for *shill* and ~*shill*, respectively.

$$\left.\begin{array}{l} m_{CBA}(shill) = \alpha_{CBA} * CBA = \alpha_{CBA} * \frac{ACB_{(i,j)}}{\sum_{i=0, i \neq j}^{n} ACB_{(i,j)}} \\ m_{CBA}(\sim shill) = 0 \\ m_{CBA}(\Theta) = 1 - \alpha_{CBA} * CBA = 1 - \alpha_{CBA} * \frac{ACB_{(i,j)}}{\sum_{i=0, i \neq j}^{n} ACB_{(i,j)}} \end{array}\right\} \quad (22.1)$$

$$\left.\begin{array}{l} m_{CBA}(shill) = 0 \\ m_{CBA}(\sim shill) = \beta_{CBA} \\ m_{CBA}(\Theta) = 1 - \beta_{CBA} \end{array}\right\} \quad (22.2)$$

BMA for Evidence WPB: When a bidder's winning ratio (i.e., Wins Per Bid) in a specific bidder's auction is lower than his normal winning ratio, the bidder is likely to be a shill bidder. The basic mass assignment for the state that the bidder is a shill is given by Eq. (23.1). If the bidder's winning ratio for a specific bidder is greater than or equal to the overall average value, the basic mass for ~*shill* can be calculated by Eq. (23.2).

$$\left.\begin{array}{l} m_{WPB}(shill) = \alpha_{WPB} * (1 - WPB_{(i,j)}) = \alpha_{WPB} * (1 - \dfrac{NOW_{(i,j)}}{NOB_{(i,j)}}) \\ m_{WPB}(\sim shill) = 0 \\ m_{WPB}(\Theta) = 1 - \alpha_{WPB} * (1 - \dfrac{NOW_{(i,j)}}{NOB_{(i,j)}}) \end{array}\right\}$$

(23.1)

$$\left.\begin{array}{l} m_{WPB}(shill) = 0 \\ m_{WPB}(\sim shill) = \beta_{WPB} * WPB_{(i,j)} = \beta_{WPB} * \dfrac{NOW_{(i,j)}}{NOB_{(i,j)}} \\ m_{WPB}(\Theta) = 1 - \beta_{WPB} * \dfrac{NOW_{(i,j)}}{NOB_{(i,j)}} \end{array}\right\} \tag{23.2}$$

BMA for Evidence BIA: When a bidder is detected to place bids with increments that are much greater than the minimum increment before the final stage of an auction, the bidder should be identified as a shill candidate. The BMA is given by Eq. (24.1). On the other hand, if a bidder places bids with only small increments, the evidence should support that the bidder is an honest bidder. Eq. (24.2) shows how to calculate the basic mass assignment for bidder *i* being honest.

$$\left.\begin{array}{l} m_{BIA}(shill) = \alpha_{BIA} * (1 - BIA) = \alpha_{BIA} * (1 - \dfrac{\sum\limits_{l=0,t=0}^{s,0.9T} \dfrac{m * MIN_l}{\sum\limits_{i=1}^{m} I_i}}{s}) \\ m_{BIA}(\sim shill) = 0 \\ m_{BIA}(\Theta) = 1 - \alpha_{BIA} * (1 - \dfrac{\sum\limits_{l=0,t=0}^{s,0.9T} \dfrac{m * MIN_l}{\sum\limits_{i=1}^{m} I_i}}{s}) \end{array}\right\} \tag{24.1}$$

$$\left.\begin{array}{l} m_{BIA}(shill) = 0 \\ m_{BIA}(\sim shill) = \beta_{BIA} * BIA = \beta_{BIA} * \dfrac{\sum_{l=0,t=0}^{s,0.9T} \dfrac{m * MIN_l}{\sum_{i=1}^{m} I_i}}{s} \\ m_{BIA}(\Theta) = 1 - \beta_{BIA} * \dfrac{\sum_{l=0,t=0}^{s,0.9T} \dfrac{m * MIN_l}{\sum_{i=1}^{m} I_i}}{s} \end{array}\right\}$$

(24.2)

BMA for Evidence AF: Users with good feedback histories are usually less likely shills. If a bidder's feedback score is greater than or equal to the average of all users' in the same category, evidence *AF* supports that the bidder is a normal bidder. In this case, the basic mass is assigned as in Eq. (25.1). Otherwise, if a bidder's feedback score is lower than the average feedback score of all users, *AF* should be counted as evidence to support shilling behavior. In this case, the basic mass is assigned using Eq. (25.2).

$$\left.\begin{array}{l} m_{AF}(shill) = 0 \\ m_{AF}(\sim shill) = \beta_{AF} * AF = \beta_{AF} * (1 - \dfrac{FB_{avg}}{FB_i}) \\ m_{AF}(\Theta) = 1 - \beta_{AF} * (1 - \dfrac{FB_{avg}}{FB_i}) \end{array}\right\} \quad (25.1)$$

$$\left.\begin{array}{l} m_{AF}(shill) = \alpha_{AF} * AF = \alpha_{AF} * (1 - \dfrac{FB_i}{FB_{avg}}) \\ m_{AF}(\sim shill) = 0 \\ m_{AF}(\Theta) = 1 - \alpha_{AF} * AF = 1 - \alpha_{AF} * (1 - \dfrac{FB_i}{FB_{avg}}) \end{array}\right\} \quad (25.2)$$

BMA for Evidence AS: A bidder may place bids in auctions that were hosted by the same seller; however, if the bidder placed bids in most of that seller's auction, the bidder shall be suspected as a shill working with the seller. When $AS_{(i,j)}$ is within the normal range, property *AS* suggests the bidder is normal and AS is used as non-shill evidence, as in Eq. (26.1). If $AS_{(i,j)}$ is greater than a threshold, say 50%, the property *AS* will be used as evidence to support that the bidder is a shill, as in Eq. (26.2)

$$\left.\begin{aligned} &m_{PS}(shill) = 0 \\ &m_{PS}(\sim shill) = \alpha_{PS} * (1 - PS) = \alpha_{PS} * (1 - \frac{AUC_{(i,j)}}{NOA_j}) \\ &m_{PS}(\Theta) = 1 - \alpha_{PS} * (1 - \frac{AUC_{(i,j)}}{NOA_j}) \end{aligned}\right\} \quad (26.1)$$

$$\left.\begin{aligned} &m_{PS}(shill) = \beta_{PS} * PS = \beta_{PS} * \frac{AUC_{(i,j)}}{NOA_j} \\ &m_{PS}(\sim shill) = 0 \\ &m_{PS}(\Theta) = 1 - \beta_{PS} * \frac{AUC_{(i,j)}}{NOA_j} \end{aligned}\right\}$$

(26.2)

BMA for Evidence NB: When the number of bids in an auction is greater than the average number of bids in auctions in the same category, that auction might be suspected of having shill bids. For the same reason, if an auction completes with number of bids less than the average, the auction might not be compromised by shills. When NB_{avg} is smaller than NB_k, property *NB* suggests this auction may involve shills. The possibility of shills can be calculated using Eq. (27.1). When NB_k is smaller than NB_{avg}, evidence *NB* suggests the auction is normal. Then we assign the value to *NB*, the non-shill evidence, as in Eq. (27.2).

$$\left.\begin{aligned} &m_{NB}(shill) = \alpha_{NB} * NB = \alpha_{NB} * (1 - \frac{NB_{avg}}{NB_k}) \\ &m_{NB}(\sim shill) = 0 \\ &m_{NB}(\Theta) = 1 - \alpha_{NB} * (1 - \frac{NB_{avg}}{NB_k}) \end{aligned}\right\} \quad (27.1)$$

$$\left.\begin{aligned} &m_{NB}(shill) = 0 \\ &m_{NB}(\sim shill) = \beta_{NB} * NB = \beta_{NB} * (1 - \frac{NB_k}{NB_{avg}}) \\ &m_{NB}(\Theta) = 1 - \beta_{NB} * NB = 1 - \beta_{NB} * (1 - \frac{NB_k}{NB_{avg}}) \end{aligned}\right\} \quad (27.2)$$

BMA for Evidence SP: If the starting price of an auction is lower than normal, it may indicate that a shill might try to drive up the price in the early stage, which results in a hidden higher starting price. For the same reason, if the starting price of an auction is higher than normal, the auction may have less chance of involving shill bids. Eq. (28.1) and Eq. (28.2) shows how to quantify evidence *SP*. When SP_k is less than SP_{avg}, property *SP* suggest we trust the auction. We assign *SP* a mass, as given by Eq. (28.1). Otherwise, we assign a mass to shill evidence *SP*, as given by Eq. (28.2).

$$\left.\begin{aligned} m_{SP}(shill) &= \alpha_{SP} * (1-SP) = \alpha_{SP} * (1-\frac{SP_k}{SP_{avg}}) \\ m_{SP}(\sim shill) &= 0 \\ m_{SP}(\Theta) &= 1-\alpha_{SP} * (1-\frac{SP_k}{SP_{avg}}) \end{aligned}\right\} \quad (28.1)$$

$$\left.\begin{aligned} m_{SP}(shill) &= 0 \\ m_{SP}(\sim shill) &= \beta_{SP} * (1-SP) = \beta_{SP} * (1-\frac{SP_{avg}}{SP_k}) \\ m_{SP}(\Theta) &= 1-\beta_{SP} * (1-\frac{SP_{avg}}{SP_k}) \end{aligned}\right\} \quad (28.2)$$

2. Evidence combination

Once the basic probability assignments are obtained, different pieces of evidence are combined in a consistent manner to provide a more complete assessment on shill bidding, and thus it reduces the uncertainties involved in individual evidence. The evidence fusion procedure can be carried out using Dempster's combination rule. The corresponding rules of combining evidence for *shill* and ~*shill* are listed as the following:

$$beief(shill_i) = m(shill_i) \quad (29)$$

$$m(shill_i) = m_1(shill_i) \oplus m_2(shill_i) \oplus \cdots \oplus m_n(shill_i) \quad (30)$$

$$belief(\sim shill_i) = m(\sim shill_i) \quad (31)$$

$$m(\sim shill_i) = m_1(\sim shill_i) \oplus m_2(\sim shill_i) \oplus \cdots \oplus m_n(\sim shill_i) \quad (32)$$

To combine multiple pieces of evidence for shill bidding behaviors, we can compute $bel(u_i)$ by combining any pair of evidence first, and then combining the result with the remaining third one, forth one, and so on, For instance, Eq. (33) – Eq. (37) combines evidence 1 and evidence 2.

$$m_1(X) \oplus m_2(X) = \frac{\sum_{E_1 \cap E_2 = X \neq \varnothing} m(E_1) * m(E_2)}{1 - \sum_{E_1 \cap E_2 = \varnothing} m(E_1) * m(E_2)} \quad (33)$$

$$m_1(\mathrm{shill_i}) \oplus m_2(\mathrm{shill_i}) = \frac{m_1(\mathrm{shill_i})m_2(\mathrm{shill_i}) + m_1(\mathrm{shill_i})m_2(\mathrm{u}) + m_1(\mathrm{u})m_2(\mathrm{shill_i})}{1-\mathrm{k}} \quad (34)$$

$$m_1(\sim \mathrm{shill_i}) \oplus m_2(\sim \mathrm{shill_i}) = \frac{m_1(\sim \mathrm{shill_i})m_2(\sim \mathrm{shill_i}) + m_1(\sim \mathrm{shill_i})m_2(\mathrm{u}) + m_1(\mathrm{u})m_2(\sim \mathrm{shill_i})}{1-\mathrm{k}} \quad (35)$$

$$m_1(U) \oplus m_2(U) = \frac{m_1(U)m_2(U)}{1-k}$$

(36) where $k = m_1(\mathrm{shill_i})m_2(\sim \mathrm{shill_i}) + m_1(\sim \mathrm{shill_i})m_2(\mathrm{shill_i})$

(37)

The factor k is a measure of the amount of conflict between the two masses. The value of $bel(u_i)$ indicates the degree of credibility of u_i.

V. CASE STUDY

The method proposed in this dissertation has been successfully examined using real online auction data from eBay. Before presenting the results and our analysis, a few issues related to implementing the experiments need to be clarified.

First, the parameters used in this approach need to be specified and fixed in order to provide a consistent result. The *alpha* value for each piece of evidence were set subjectively based on different levels of importance as we observed. For example, we considered the evidence of Concurrent Bid Activity and the evidence of Affinity for Seller to be more important than other evidence. Therefore, we assign them a higher weight. In addition, the evidence of Time to Last Bid is not likely to be very reliable by itself to determine a shill in comparison to other evidence. Yet the later a bidder places a bid in an auction, the less suspicious it is that the bidder is a shill since a late bid increases the risk of winning. Based on such observations, we set the *alpha* value for each piece of evidence according to the importance of the evidence in determining a shill. Initially, the *alpha* values for TLB, AS, CBA, WPB, BIA, and AF are assigned 0.6, 0.9, 0.9, 0.8, 0.8, and 0.7, respectively. However, with these *alpha* values we observed that shill scores computed over some training data were not highly distinguishable. We manually tested and adjusted the alpha values on a set of training data in the category of WII game system until the shill scores for shills and innocent bidders were distinguishable; in other words by slightly adjusting the evidence weights, the *alpha* values not only reflect the importance of evidence but also make shills more distinguishable from innocent bidders. Thus, for the data reported here, we set α_{TLB}, α_{AS}, α_{CBA}, α_{WPB}, α_{AF}, α_{BIA}, α_{NB}, α_{SP} as 0.6, 0.95, 0.95, 0.9, 0.8, 0.7, 0.8 and 0.8, respectively. It is possible that the "best" alpha values may vary depending on the category of item up for auction. In future work, one can experiment with some approach, such as designing a neural network, to learn the precise *alpha* values automatically from larger data sets. Meanwhile, for this research, we set $\beta = \alpha$ for each piece of evidence just for simplicity.

A. Data Collection

The data used in our case study was collected from a recent auction on eBay with the title "Microsoft Xbox 360 Complete System and 20G Hard Drive." A detailed bidding history of the auction is shown in Figure 6. To protect the privacy of bidders, only symbolic IDs instead of the bidders' IDs are shown (along with the bidder's reputation score, shown in parenthesis). The detailed description of the item is provided in Table V.

B. Data Processing

The record for our case study is shown in Figure 7. Information about the auction, such as the detailed bidding history, can be obtained directly from the auction web pages. Most of the statistical data can be obtained from posted auction information and eBay web services. For instance, via eBay web services, we collected the 30-day average final price of all Xbox auctions that has the same options and it is $138.94. We also obtained that the average starting price of this kind of auction is $40.64; the average number of bids per auction is 7.67; and the average feedback score for bidders bidding actively in this category is 101.98. The statistics is updated periodically, but for the same category of items, it is constant in a short period of time, e.g., 1 day. Besides, some of the detailed information about the seller and bidders, such as feedback score and links to finished auctions, can be captured from their posted profiles.

TABLE V.

MICROSOFT XBOX 360 PRO SYSTEM - GAME CONSOLE - 20 GB

Description	
Model:	Microsoft Xbox 360
Hard Drive Capacity:	20GB
Features	
Audio Output:	Surround Sound
Video Output:	ATI Xbox 360 - 256-bit - 2D/3D graphics acceleration
Max.	1920 x 1080
Connections	
1 x AV cable port, 3 x USB 2.0, 1 x Ethernet (RJ-45)	

XBOX 360 Auction Bidding History						
	Bidders:12	Bids:61		Time Ended:MAY-07-09 11:58:07 PDT		
e***e	US$167.50	May-07-09 11:57:54 PDT		6***o	US$60.00	May-06-09 19:11:53PDT
o***i	US$165.00	May-07-09 11:57:58 PDT		p***k	US$60.00	May-06-09 19:15:27PDT
e***e	US$160.00	May-07-09 11:57:27 PDT		o***i	US$58.00	May-06-09 19:11:32PDT
o***i	US$160.00	May-07-09 11:57:33 PDT		6***o	US$56.00	May-06-09 19:11:15PDT
o***i	US$155.00	May-07-09 11:57:05 PDT		a***l	US$55.00	May-06-09 19:10:36PDT
o***i	US$150.00	May-07-09 11:56:34 PDT		6***o	US$55.00	May-06-09 19:10:59PDT
o***i	US$146.00	May-07-09 11:56:10 PDT		o***i	US$52.00	May-06-09 19:08:53PDT
s***h	US$143.50	May-07-09 11:55:39 PDT		6***o	US$50.00	May-06-09 17:57:50PDT
o***i	US$142.00	May-07-09 11:54:47 PDT		o***i	US$50.00	May-06-09 19:08:38PDT
s***h	US$138.50	May-07-09 11:41:07PDT		o***i	US$45.00	May-06-09 19:08:53PDT
f***a	US$136.00	May-07-09 11:02:16 PDT		o***i	US$42.00	May-06-09 19:08:24PDT
s***h	US$133.50	May-07-09 10:47:42 PDT		o***i	US$40.00	May-06-09 19:08:11PDT
o***i	US$131.00	May-07-09 10:35:33 PDT		o***i	US$38.00	May-06-09 19:07:54PDT
s***h	US$128.50	May-07-09 09:50:43 PDT		p***p	US$35.00	May-06-09 19:07:42PDT
s***l	US$128.00	May-07-09 09:45:58 PDT		i***e	US$30.00	May-06-09 12:17:49PDT
s***h	US$123.50	May-07-09 09:50:02 PDT		p***p	US$30.00	May-06-09 13:28:46PDT
s***h	US$118.50	May-07-09 09:26:31 PDT		P***p	US$25.00	May-05-09 17:29:14PDT
s***l	US$116.00	May-07-09 06:45:20 PDT		n***0	US$20.00	May-05-09 18:11:46PDT
o***i	US$115.00	May-07-09 06:08:34 PDT		p***k	US$13.00	May-05-09 17:33:43PDT
o***i	US$110.00	May-07-09 06:08:05 PDT		p***k	US$12.00	May-05-09 17:32:07PDT
o***i	US$102.00	May-07-09 06:07:42 PDT		p***k	US$11.00	May-05-09 17:31:45PDT
s***l	US$100.00	May-07-09 20:06:30 PDT		p***k	US$10.00	May-05-09 13:39:28PDT
o***i	US$98.00	May-07-09 06:07:25 PDT		P***p	US$10.00	May-05-09 17:29:06PDT
o***i	US$95.00	May-07-09 06:07:09 PDT		p***p	US$5.00	May-04-09 19:49:57PDT
o***i	US$90.00	May-07-09 06:06:53 PDT		p***k	US$5.00	May-05-09 13:38:05PDT
o***i	US$88.00	May-07-09 06:06:36 PDT		p***k	US$3.00	May-05-09 13:37:51PDT
6***o	US$86.00	May-07-09 05:28:00 PDT		p***k	US$2.50	May-05-09 13:37:34PDT
p***p	US$75.00	May-06-09 21:36:33 PDT		v***i	US$1.00	May-04-09 18:33:24PDT
p***p	US$70.00	May-06-09 21:36:25 PDT		p***p	US$1.00	May-04-09 19:49:47PDT
p***p	US$67.00	May-06-09 21:36:12 PDT		p***p	US$0.50	May-04-09 19:49:36PDT
p***p	US$65.00	May-06-09 21:15:14PDT		Starting Price	US$0.01	May-04-09 11:58:07PDT

Figure 6. Bidding history

We investigated auctions hosted by a particular seller during the past month, and the bidding history of every bidder who participated in this auction. The historical information is statistically processed. We counted each bidder's number of wins, total number of bids, total number of auctions participated, bid activity with the same seller, feedback score, number of abnormal concurrent bids, and etc. The statistical results are shown in Table VI. The basic masses assigned for evidence specified in Section D.1 are shown in Table VII and Table VIII.

RECORD							
Microsoft Xbox 360 System-Game Console-20GB Hard Drive							
AUCTION INFO	Starting price	Num. of Bids	Num. of Bidders	Duration	Start time	End time	Winning Bid
	0.01	42	12	3day(259200 s)	lay-04-09 11:58:07 PI	May-07-09 11:58:07 PD	US $167.50
HISOTRICAL STATISTICS	Average price	Average starting price	Average number of Bids	Average Feedbacks	Increment_{MIN}		
	138.94	40.64	7.67	101.98	($0.01-$0.99)$0.05; ($1.00-$4.99)$0.25;($5.00-$24.99)$0.5;($25.00-$99.99)$1.00;($100-$249.99)$2.5;($250-$499.99)$5.00;		
SELLER INFO	Seller ID	Num. of auctions	Bids attracted in video Game	Bids/Auction	Num. of categories	Feedback(percentage)	Active Since
	bascoo1998	36	909	25.25	7	9(90.9%)	MAY-3-2009

Figure 7. History record

TABLE VI.
STATISTICAL DATA

Bidder	Bids in this auction	Bid Activity with the seller[1]	Items bid on	Times bid on this seller's auction[2]	Same category sellers (all sellers) placed bid[3]	Total bids	Wins	Feed-back score	Num. of ACB[4]	Time from last bid to the end of auction[5]	Increment$_{AVG}$			
											$1	$0.5	$0.25	$0.05
e***e	2	1%	88	1	1(30)	112	2	642	0	13	0	0	0	0
o***i	21	79%	10	6	2(3)	82	1	2	11	9	6	0	0	0
s***h	6	100%	1	1	1(1)	6	0	0	0	148	0	0	0	0
f***a	1	100%	1	1	1(1)	1	0	0	0	3351	0	0	0	0
s***l	3	30%	120	30	4(6)	283	5	27	71	7929	$40	0	0	0
6***o	5	100%	3	3	1(1)	11	1	3	0	23407	$6.2	0	0	0
p***p	11	31%	5	1	5(5)	35	1	8	0	51094	$8.33	$10	$4	$0.495
p***k	8	83%	5	2	2(2)	30	0	0	1	60610	0	$11	0	0
a***l	1	7%	7	1	2(6)	17	0	20	0	60451	$3	0	0	0
i***e	1	100%	1	1	1(1)	1	0	7	0	85218	$5	0	0	0
n***0	1	50%	7	3	2(5)	24	0	8	0	150381	0	$19.5	0	0
v***i	1	53%	3	1	1(3)	3	0	0	0	235483	0	0	0	$0.99

1. This shows the percentage of all bids from this bidder that went to this specific seller. 2. In the last 30 days, the number of auctions that were hosted by the seller the bidder participated in. The seller's total number of Xbox Game System auctions is 36. 3. This shows that the bidder placed bids for how many different games system sellers. The number in the parenthesis is the total number of sellers the bidder has placed bid for, no matter what categories of goods they sell. 4. ACB stands for abnormal concurrent bid; 5. time from last bid to the end of auction (in seconds) = The duration of the auction - duration of a bidder's last bid since the auction begins. eBay only provides a bidder's bid history for the last 30 days.

TABLE VII.
THE BASIC MASS ASSIGNMENTS FOR BID-LEVEL EVIDENCE TLB, AS, AND ACB

Bidder	m_{TLB}(shill)	m_{TLB}(~shill)	m_{TLB}(U)	m_{AS}(shill)	m_{AS}(~shill)	m_{AS}(U)	m_{CBA}(shill)	m_{CBA}(~shill)	m_{CBA}(U)
e*e**	0	0.5	0.5	0	0.9236	0.0764	0	0.95	0.05
O***i	0	0.9	0.1	0	0.7917	0.2083	0.3458	0	0.6542
s*h**	0.0003	0.6	0.4	0	0.9236	0.0764	0	0.95	0.05
f*a**	0	0.592	0.408	0	0.9236	0.0764	0	0.95	0.05
s*l**	0.0184	0	0.9816	0.7917	0	0.2083	0.8693	0	0.1308
6*o**	0.0542	0	0.9458	0	0.8708	0.1292	0	0.95	0.05
P*P**	0.1183	0	0.8817	0	0.9236	0.0764	0	0.95	0.05
P*k**	0.1403	0	0.8597	0	0.8972	0.1028	0	0.95	0.05
a*l**	0.1399	0	0.8601	0	0.9236	0.0764	0	0.95	0.05
i*e**	0.1973	0	0.8027	0	0.9236	0.0764	0	0.95	0.05
n*0**	0.3481	0	0.6519	0	0.8708	0.1292	0	0.95	0.05
v*i**	0.5451	0	0.4549	0	0.9236	0.0764	0	0.95	0.05

TABLE VIII.
THE BASIC MASS ASSIGNMENTS FOR BID-LEVEL EVIDENCE WPB, BIA, AND AF

Bidder	m_{WPB}(shill)	m_{WPB}(~shill)	m_{WPB}(U)	m_{BIA}(shill)	m_{BIA}(~shill)	m_{BIA}(U)	m_{AF}(shill)	m_{AF}(~shill)	m_{AF}(U)
e*e**	0	0.45	0.55	0	0.8	0.2	0	0.58884	0.4112
O***i	0	0.0139	0.9861	0	0.0333	0.9667	0.6863	0	0.3137
s*h**	0	0	1	0	0.8	0.2	0.7	0	0.3
f*a**	0	0	1	0	0.8	0.2	0.7	0	0.3
s*l**	0	0.0482	0.9518	0.795	0	0.205	0.51467	0	0.4853
6*o**	0.9	0	0.1	0	0.0323	0.9677	0.6794	0	0.3206
P*P**	0	0.0829	0.9171	0	0.0667	0.9333	0.6451	0	0.3549
P*k**	0	0	1	0	0.0091	0.9909	0.7	0	0.3
a*l**	0	0	1	0	0.066	0.934	0.56272	0	0.4373
i*e**	0	0	1	0	0.04	0.96	0.65195	0	0.348
n*0**	0	0	1	0.7949	0	0.2051	0.6451	0	0.3549
v*i**	0	0	1	0	0.0101	0.9899	0.7	0	0.3

Besides evidence from bid level, we also use evidence from the auction level, such as starting price (SP) and number of bids (NB), to facilitate the certification of auction bidders. According to Eq. (27.1), Eq. (27.2), Eq, (28.1), and Eq. (28.2), the masses for SP and NB are assigned as shown in Table IX.

TABLE IX.
BASIC MASS ASSIGNMENTS FOR AUCTION-LEVEL EVIDENCE

m_{NB}(shill)	m_{NB}(~shill)	m_{NB}(U)	m_{SP}(shill)	m_{SP}(~shill)	m_{SP}(U)
0.65360	0	0.3464	0.799	0	0.201

The shill certification results are shown in Table X. Recall that functions *bel* and *pl* define belief and plausibility, respectively, as presented in Section B. Each bidder is recognized with one of the three certifications: *Shill*, *Suspect*, and *Trusted Bidder*. The certification levels ensure that each bidder is certified and fraudulent bidders are identified. In this example, we set R= {0.95, 0.5}. To reduce the number of false positives generated from our proposed approach, the shill threshold φ should be sufficiently high. These certification results are assigned in accordance with the rule shown in Figure 5.

C. Certification Result Analysis and Discussion

We now analyze the auction data and the certification results by considering three levels of certification. In the Xbox auction, there are totally 12 bidders and 61 bids in the auction. At the end of the certification process, 8 bidders are certified as *Trusted Bidder*, 3 bidders are certified as *Suspect*, and 1 bidder is certified as *Shill*. We first examine the auction-level evidence. The auction attracted 61 bids. This number is much higher than the average number of bids, which is 7.67, in the same type of auctions. Note that this is not due to the lower price of this auction than those of the concurrent auctions, because the final price of the auction is $167.5, which is higher than the average final price, $138.94, of auctions selling the same item. Furthermore, the average starting price of auctions selling the same product is $40.64, but the starting price of this auction is merely $0.01. While lower starting price may attract bidders to the auction, there is also a higher probability that the seller planed to employ shills to set up a hidden reserve price in order to sell the item at a satisfactory price. The auction-level analysis supports that the auction under investigation may involve shills.

TABLE X.
SHILL CERTIFICATION RESULTS

Bidder	$bel_{(shill)}$	$pl_{(shill)}$	$bel_{(\sim shill)}$	$pl_{(\sim shill)}$	Results
e***e(642)	0.00115	0.00124	0.99876	0.99885	Trusted Bidder
O***i(2)	0.57803	0.58641	0.41359	0.42197	Suspect
s***h(0)	0.01398	0.01428	0.98572	0.98602	Trusted Bidder
f***a(0)	0.01440	0.01471	0.98529	0.98560	Trusted Bidder
s*l(27)**	**0.99981**	**0.99999**	**0.00001**	**0.00019**	**Shill**
6***o(3)	0.74710	0.74868	0.25132	0.25290	Suspect
P***P(8)	0.12798	0.13083	0.86917	0.87202	Trusted Bidder
P***k(0)	0.21782	0.22180	0.77820	0.78218	Trusted Bidder
a***l(20)	0.11713	0.12028	0.87972	0.88287	Trusted Bidder
i***e(7)	0.15599	0.15909	0.84091	0.84401	Trusted Bidder
n***0(8)	0.66078	0.66298	0.33702	0.33922	Suspect
v***i(0)	0.28270	0.28542	0.71458	0.71730	Trusted Bidder

We now investigate various bidders with different certifications to see if our manual investigation is consistent with the certification result.

Shill: The system only certifies one bidder, *s***l* as a shill in this auction. Given Eq. (33) – (37), the degrees of belief from available evidence are combined to obtain the joint belief of shill. The belief of shill for *s***l* is 0.9998, which is greater than φ (0.95), thus the system updates the certification of *s***l* to *Shill*. There are reasons to consider this ultimate result as reasonable because this bidder's behavior is very suspicious. First, *s***l* has a very obvious bidding pattern as shown in Table XI. In most of the Xbox 360 auctions that were hosted by the same seller, *s***l* joined the auction in the middle, placing a proxy bid at $75, and then when outbid, this bidder increases the bid to $100. Most of the time, this bidder stops bidding at $125. From this bidding pattern, it looks like *s***l* is driving up the price and outbidding the potential buyers until the price is relatively high (e.g., $125) and the risk of winning is high. Second, *s***l* has placed bids in 30 out of 36 auctions that were hosted by the same seller. The high number indicates *s***l* and the seller have a strong business relationship. However, the winning ratio for *s***l* is low. Even though he has won 5 auctions, the winning price is relative low. These wins are most likely accidental wins because most of the winning prices are $100 or $125, which looks to be the habitual bid values *of s***l*. So, bidder *s***l* was forced to win the auction when nobody placed a higher bid. Third, in almost all of the auctions, *s***l* has placed bids with increments 40 times the minimum increment. Note that typical normal bidders tend to bid cautiously, with the minimum increment. Fourth, *s***l* placed as high as 70 abnormal concurrent bids in one month. Bidder *s***l* might know that the price in the specific seller's auction was higher than that of a concurrent auction and that this auction would end later than the other one, yet *s***l* still placed bids on the seller's auction.

TABLE XI.
A SUSPICIOUS BIDDING HISTORY FOR S***L

Winner	Winning bid	Num. of Bidders	Num. of bids	Bids of s***l(27)
a*a(4)**	168.49	7	20	80-115
g*o(20)**	202.5	8	18	NA
t*e(2)**	149.5	13	31	NA
s*l(27)**	100	6	14	100
z*u(1)**	147.5	7	15	100-110-115-120-125-140-145
3*6(163)**	171	15	27	70-100-110
d*d(195)**	168.5	8	16	100-120-130
0*1(0)**	152.5	12	32	130-140
e*e(642)**	167.5	12	61	100-116-128
o*i(2)**	129.5	9	23	88-127
6*o(3)**	137.5	6	27	100-120-127-135
o*d(9)**	137.5	15	24	100-125
b*d(27)**	152.5	11	22	100-130-140

All in all, the evidence is consistent with the certification results. Interestingly, several days later after we collected the data, we found that *s***l* became labeled as "No Longer A Registered User (NLARU)" in eBay.

According to eBay's explanation, NLARU means the bidder's account is suspended by eBay due to violations of eBay's policy, such as shill bidding, selling counterfeit item, keyword spamming, transaction outside eBay. Although we would not be able to know the actual reason why the account was suspended, as a bidder, the most possible reason for *s***l* to be labelled NLARU is due to shill bidding. This once again helps confirm the shill certification result.

Suspect: For suspects, the evidence is not sufficient enough to support the bidders as shills but it is still more sufficient than evidence that supports *~shill*. In the case study, the system gives the certification of *Suspect* to three bidders. They are bidders *o***i*, *6***o*, and *n***0*. We first justify bidder *o***i*. The statistics shows that *o***i* placed bids in 6 out of 36 auctions hosted by the same seller on different dates. He placed many bids in Xbox auctions but won only once at a very low price ($125.5). This behavior matches one of the most significant shilling behavior characteristics: bids frequently, but seldom wins. Besides, *o***i* placed 11 abnormal concurrent bids in the seller's auction. This evidence further enforces the belief that *o***i* is suspicious. However, *o***i* placed his last bid at the final stage of the auction, and also placed bids in many other sellers' auctions. There is strong evidence that supports both sides. Therefore, this bidder is certified as a shill suspect. For the other two bidders, we can observe that the wining ratio (Wins Per Bid) of *6***o* for the specific seller is lower than his average winning ratio, and the increment of *n***0* is exceptionally high. These behaviors make the shill evidence more sufficient than the honest evidence. Therefore, bidder 6***o and bidder *n***0* are certified as suspect at this moment. The certification result is again consistent with our manual investigation.

Trusted Bidder: We show what kind of bidders is considered honest. Without much doubt, *e***e* is not a shill because *e***e* only participated once in the seller's auctions, and won at the end. This win is not accidental since *e***e* placed all of his bids in the final stage of the auction. Now we consider bidders *s***h* and *f***a*. Even though they contribute all of their bids to the seller and they did not win any auction, both of them only participated in only one auction, and they placed bids in the final stage of the auction. Such bidding behavior indicates that they are at least not afraid of winning. Therefore, they are most likely typical cautious new bidders. The reason why *p***p* is not a shill or shill suspect is that *p***p* placed bids in five different sellers' auctions and *p***p* finally won one Xbox game system auction that was hosted by other seller. For the other two bidders, *p***k* and *a***l*, it is easily to see that they placed bids on more than one Xbox system auctions hosted by different sellers. Although they did not win any of the auctions, they did not place any abnormal concurrent bid and they increased their bids cautiously. To sum up, our calculation results are consistent with our manual investigation for identifying trusted bidders.

To study the importance of the auction-level evidence, we performed the shill certification again, but using only bid-level evidence. The shill certification results for the same auction data collected from eBay are listed in Table XII. Based on the experimental results, we found that with auction-level evidence, the values for belief-of-shill can be amplified, so some bidders (O***i(2), 6***O(3), and n***0(8)) considered as normal bidders (without auction-level evidence) are now considered as Suspects. This would be valuable to identify any suspicious bidders who can then be subjected to further investigation. Note that a detailed discussion on the impacts of adopting different levels of evidence is beyond the scope of this dissertation, but it is envisioned as useful for future work.

TABLE XII.

SHILL CERTIFICATION RESULTS WITHOUT AUCTION-LEVEL EVIDENCE

Bidder	bel(shill)	pl(shill)	bel(~shill)	pl(~shill)	Result
e***e(641)	0	8.63722E-05	0.9999	1	Trusted Bidder
O***i(2)	0.0714	0.0899	0.9102	0.9286	Trusted Bidder
s***h(0)	0.0007	0.0010	0.9990	0.9993	Trusted Bidder
f***a(0)	0.0007	0.0010	0.9990	0.9993	Trusted Bidder
s*l(27)**	**0.9972**	**0.9999**	**0.0001**	**0.0028**	**Shill**
6***O(3)	0.1666	0.1718	0.8282	0.8334	Trusted Bidder
P***P(8)	0.0071	0.0104	0.9896	0.9929	Trusted Bidder
P***k(0)	0.0144	0.0195	0.9805	0.9856	Trusted Bidder
a***l(20)	0.0059	0.0094	0.9906	0.9941	Trusted Bidder
i***e(7)	0.0094	0.0130	0.9869	0.9906	Trusted Bidder
n***0(8)	0.1147	0.1205	0.8795	0.8852	Trusted Bidder
v***i(0)	0.0234	0.0271	0.9729	0.9766	Trusted Bidder

The Dempster-Shafer theory based shill certification approach reported in this case study is also employed to identify shills in Section C of Chapter VI, in which we empirically study if the difference between actual and expected auction price is related to shill bidding. In the experiments, shills are found in 49 auctions. We have manually confirmed the existence of shills in the 49 identified auctions. The manual investigation was done in a similar way as shill analysis presented earlier in this section.

To manually investigate shilling behaviors, we focused on looking for three different groups of evidence. The first group of evidence is the shill related evidence reported in Section C of Chapter IV. It is important to verify the existence of this group of evidence because we need to make sure that the D-S theory based shill certification approach works as desired. In addition, this group of evidence is very indicative of shill biddings. After investigation, we noticed that all of the identified shills had some abnormal behaviors as the evidence used in the shill certification process, such as abnormal concurrent bidding activity, low wins-per-bid ratio, large bid increment, etc.

In the second group of evidence, we tried to discover independent suspicious bidding patterns and collaborations evidence that are proposed by other researchers [FXV12] but were not used in our approach, namely Elapsed Time before First Bid (ETFB) and Number of Bids in different stage (NBDS).

ETFB is the time that elapses from the start of an auction to a bidder's first bid. A large value of ETFB indicates that the bidder started participating late in the auction, whereas a small value of ETFB indicates that the bidder participated very early in the auction. Although it is possible for a normal bidder to place bids early in an auction, placing bids extremely close to the start of the auction implies the bidder's possible prior knowledge about the auction. Thus, a bidder with a very small value of ETFB is suspicious.

NBDS refers to the number of bids placed by a bidder in a particular auction stage. Following the definition in [XBS09], the early stage refers to the first quarter of the auction duration; the middle stage refers to [0.25, 0.9] of the auction duration; and the final stage refers to the last 10% of the auction duration. Note that we adopted the number of bids (NB) of the entire auction as an auction-level shill indicator in the shill certification approach. However, NBDS is quite different from NB according to the definitions. The NB in different stage of an auction by a bidder indicates different possibilities of shill bidding. A large value of NB at the early stage typically indicates a suspicious bidder's desire to raise the auction price quickly or set up a hidden reserve price. A large value of NB at the middle stage is not considered as a strong indicator of shill bidding because a normal bidder may legally participate in a bid war in the middle stage. A large value of NB at the final stage is against a shill's intent to avoid winning the auction, thus does not indicate shill bidding.

When investigating the 49 auctions involving shills, we observed small value of ETFB and large number of bids in early stage of auctions frequently.

The third group of evidence we targeted is also independent from the evidence we used in the D-S theory based shill certification approach. Although the "NLARU" label is not proof for shilling behavior, it indicates likely shill bidding as we analyzed earlier. Therefore, in the manual investigation process, the NLARU label is used as an additional piece of supportive evidence for shill bidding. Among the 49 suspicious auctions, four bidders were labeled as NLARU. Except "NLARU", we also looked at a seller's feedback rating (SFR) as discussed in [FXV12]. A bidder who placed a bid in an auction hosted by a seller with zero feedback score is suspicious since bidders are most likely to bid in an auction if the seller has significant positive feedback. Supplementing to the two pieces of auction-level evidence that is used in the shill certification approach, SFR can be considered as one more piece of independent auction-level evidence for shill bidding.

Overall, we have investigated the identified 49 auctions. In each auction, we found several pieces of supportive evidence on either auction level or bid level, or both, which indicates some belief that the auction is suspicious in terms of shill biddings.

VI. PRICE COMPARISON: A RELIABLE APPROACH TO IDENTIFYING SHILL BIDDING IN ONLINE AUCTIONS?

This chapter discusses an approach to identifying shill bidding by comparing actual auction price to expected auction prices, and is based on work in [DSX10, DSX11].

A. Background

In order to alleviate the effects of shill bidding and provide bidders convenient assessment rules, in this chapter we study the relationship between shill bidding and final selling price in online auctions and also examine rules that may help detect shill bidding. Note that we use the terms *actual auction price* and *final auction price* interchangeably, to mean the final bid in first price auctions. The *expected auction price* is the predicted auction price, which can be learned by the observation of historical auction data from the same category of items. We explore answers to the following research questions:

- If there is a significant price difference between actual auction price and expected auction price, is this difference related to shill bidding? Does this difference provide a clue for shill bidding?

- What if an auction's final auction price is equal to or lower than the expected auction price? Does this imply unlikely shill bidding?

We formally examine the association of final auction price and shill bidding using chi-square statistic, and carefully inspect their relationship by adopting a logistic regression model. Our empirical study closely follows the guidelines for empirical research in software engineering defined in [KPP02]. We report data sampling criteria, hypotheses formation, parameter significance test and goodness fit of models.

B. Hypotheses Formulation

In this section, we first present different opinions of how shill bidding is related to final prices in online auctions. Based on these opinions, hypotheses concerning how shill bidding is decided by final auction prices are proposed.

1. Conflicting opinions

Generally, two conflicting opinions can be observed. On one hand, there is the possibility of a so-called "Shiller's curse," which presumes that when there are shill bids in an auction, the auction's final price will be lower than it otherwise would be – as if the shills were cursed. On the other hand, there is a view that shill bids might be interpreted as signals that the item is highly valued, thus inflating the final auction price.

Shiller's curse. "Shiller's curse" is a term that was proposed by Wang et al. [WHW02]. The idea is that when buyers realize or suspect the existence of shill bidding, they may shield their bid and quit the auction or wait for the seller to sell the item for a lower price in the next round of the auction. In accordance to this logic, we may expect that under such a situation the final auction price would be suppressed. This is because bidders become conservative when considering shill bidders' existence and may even quit the bid war, so the auction price is not bid up. This scenario is practical since most online auction websites allow sellers to relist their items several times until the items are sold. Relisting occurs for many reasons, for example no bidders having placed bids in the auction, the final price of the auction not reaching the reserve price, or a seller engaged in shilling and accidentally "won" the auction. Kosmopoulou and De Silva [KDS07] examined the effect of shill bidding in online auctions on the seller's payoff and on the price. Through a series of experiments, they found that bidders tended to overpay, but prices decreased, as they anticipated the behavior of sellers and adjusted their bidding strategies.

Signaling Effect. There are also researchers who believe that shill bidding will inflate the final price of an auction. Wang et al. [WHW01] showed that private-value English auctions with shill bidding could result in a higher expected seller profit than other types of auctions, violating the classical revenue equivalence theory. Kauffman and Wood [KW05] examined the effects of shill bidding on final bid price in rare coin auctions and showed that some bidders might view shill bids as signals that an item is worth more, thus they might pay more than other bidders who cannot see such signals. Roth and Ockenfels [RO02] noted that shill bids could be considered as valuation indicators for an item. The earlier a shilling signal appears for an item, the higher the final price will be. Thus, bidding high early in the auction could possibly invoke a bidding war.

2. Hypotheses

We can see that there are contrasting views about how shill bidding affects final price in an auction. Before we investigate whether the difference between actual auction price and expected auction price can imply shill bidding, we want to first test if shill bidding has effect on the final auction price. In case the existence of shill bidding in an auction has no effect on the final auction price, the final auction price should be close to the expected price. Under this circumstance, there is no incentive to infer shill bidding by comparing final actual price to expected price in an auction. However, if the existence of shill bidding results in an increase of the final auction price, the final auction price should be significantly higher than expected price; on the other hand, if the existence of shill bidding results in a decrease of the final auction price, the final auction price should be significantly lower than the expected price. Hence, whether the difference between actual and expected auction price can imply shill bidding depends on if shill bidding has an effect on the final auction price.

From the theoretical point of view, it is possible that rational and sophisticated bidders can properly evaluate the end price of an auction in order to adjust their bids accordingly, so shill bidding does not affect the final auction price. For example, when the bidders know there are shill bids in an auction, they may shave their bids and keep bidding until they have to bid more than their own valuation of the item. As such, shill bid cannot allure competing bids, and the final bid for an item will be the same as it would have been without the shill bid. This leads to our first hypothesis.

□ ***H_1*:** *The difference between actual and expected auction price in online auctions is not related to shill bidding.*

Since a shill bidder's primary goal is to drive up the final price, it seems reasonable that the final auction price should be significantly higher than it would have been if no shill bids were placed in the auction. In addition, Roth and Ockenfels [RO02] claimed that bids might act as signals to other bidders about the quality of the item, as well as the quality of service and trustworthiness of the seller. As such, shill bidding may lead to a high final price of the auctioned item. Kauffman and Wood [KW05] also found that shill bidding acted as a signal for other bidders to place higher bids and thus increased the auction's winning bid. In other words, if shill bidding occurs in an auction, the auction will very likely end with a higher final auction price. Based on these views, we have our second hypothesis:

□ ***H_2* :** *A lower-than-expected or as-expected final auction price indicates unlikely shill bidding.*

Inspired by the logic of Shiller's Curse, sophisticated bidders will reduce their bids to a greater extent if they observe a suspicious bidder, so the auction's final price will be decreased. This logic implies that shill bidding is not likely to inflate the auction price. Therefore, we have the third hypothesis:

□ ***H_3* :** *A higher-than-expected final auction price indicates unlikely shill bidding.*

Auction prices can reflect the current market of the auctioned item. If an auction's final price falls in the normal price range – for example within +/-20% of the average price – the final auction price conforms to the market discipline. In contrast, if the market for a type of item is depressed, the final auction prices for these items are not expected to be high. Under this circumstance, if the final prices of a particular seller's auctions are significantly and consistently higher than those of the same-item auctions, the seller is suspected of employing shill bidding or other types of fraudulent bidding activities. Therefore, the final price of an auction can be considered as an indicator of the trustworthiness of the auction. An auction that does not involve shill bidding is likely to have a lower-than-expected, or as-expected, final auction price. This leads to the fourth hypothesis.

□ ***H_4* :** *A higher-than-expected final auction price indicates possible shill bidding.*

Because there is no related research supporting that a lower-than-expected final auction price indicates shill bidding, and this condition is counter-intuitive, we leave out this hypothesis in this dissertation. Overall, we have four competing hypotheses, each based upon previous research results, that predict if shill bidding has influence on final auction price and different possibilities of shill bidding when there is a significant difference between final auction price and expected auction price.

C. Experiments

We now turn our attention to the development of experiments to investigate the relationship between shilling behavior and final selling price in online auctions.

1. Experiment design

First, we built and trained a neural network to predict auction prices. When the neural network based price predictor achieved good performance, we employed the price predictor to predict the final prices of new auctions that were not used in the training and testing phases. Since the price predictor can achieve a relatively high accuracy, we consider the predicted prices as "expected" prices. We randomly selected 192 auctions that are not used for training and testing the price predictor. A predicted price is computed for each of the auctions and compared to the actual price. The relationship between actual price and expected price is recorded. A shill score is also computed for each of the 192 auctions. The shill score of an auction is defined as the highest belief of a bidder being a shill among those who participate in the auction. The shill score is a number between 0 and 1 (inclusive) that indicates the likelihood of an auction involving shills. In brief, the predicted price is compared with the actual final auction price, and the relationship is recorded together with shill information. We performed chi-square test and tests based on logistic regression on the data to test hypotheses formulated in Section B.2. In the following sections, we describe the methods for obtaining expected auction prices and deciding shills.

2. Data collection and sample choice

A major challenge of this research is the lack of annotated auction datasets. For this reason, we designed a software agent to collect auction data from eBay.com, which offers a broad range of auctioned items and provides detailed information of each auction as well as some limited information on sellers and bidders. The data we gathered is under the category of Nintendo WII game console systems. There are two reasons why we choose this category. First, Nintedo WII game console systems were popular among shoppers when the data was collected, and thus auctions of WII game systems attracted ample bidders and bids for this study. Second, the popularity and the price range made WII game systems good targets for shilling. In other words, if an item is less popular or the price range of an item is low, such as several dollars, the item is less likely to be the target of shill bidding. Although the broad categorization of the data is WII game console system, the items bundled with the game systems vary from auction to auction. For example, the items for sale in one of the auctions include a WII game console and a new WII FIT, which is an accessory of WII game system; while in another auction the auctioned item is a bundle of Nintendo WII System Console, Steering Wheel and 13 Games. Therefore, the category of WII console system is still a broad category that contains many types of items. This partially explains why the prices of this category cover a wide range.

We gathered a dataset of 1792 distinct auctions. The data was separated into two disjoint groups randomly. One group containing 1600 shill-free auctions was used to train and test the price prediction model that is to be introduced in Section C.3. Another group containing 192 auctions is used to fit the logistic regression model, to be introduced in Section D.2.

For each auction, we collected data that was provided by the seller, including information about the seller, details of the item (name, specifications, description, photos, etc.) and attributes about the auction (length, starting bid, reserve price, shipping charges, etc.). The data is processed to extract attributes and create new features that are then used to predict the final prices. The data features, classified in four different groups, are listed in Table XIII.

TABLE XIII.
DATA FEATURES IN 4 DIFFERENT GROUPS

Group Name	Item	Seller	Bid Details	Category Specific
Features	Condition (new, used, refurbished, unspecified)	Reputation percentage (%)	Number of bids	Bundle (Yes/No)
			Number of bidders	
	Number of pictures	Has a web-store (Yes/No)	Starting bid ($)	Game (number)
	Shipping cost ($)	Reputation score	Month	Wheel (Yes/No)
			Day	
			Starting time	
	Description	Top-Rated (Yes/No)	End Time	Fit (Yes/No)
			Duration	

3. Price prediction

We built a price predictor based on a neural network, specifically the Large Scale Memory Storage and Retrieval (LAMSTAR) [GK98]. The LAMSTAR, which combines Self Organization Map (SOM) and statistical decision tools, has been successfully applied to diagnosis, prediction and detection applications [G07]. The trained network for predicting auction final prices is shown in Figure 8.

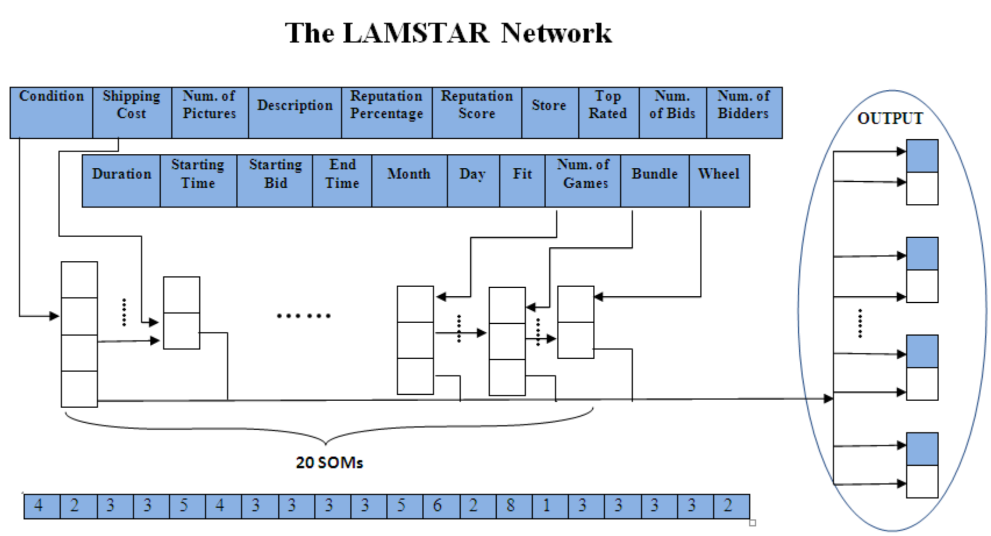

Figure 8. The nerual network for predicting auction final prices

The grids on the top of the figure are the subwords, representing the features in Table XIII. The features are preprocessed and then provided as inputs to the neural network. For every subword there is an associated SOM module (in the middle of the figure) that is used to store and retrieve information in the training process. For each subword, a winning neuron in the associated SOM module is determined based on the similarity between the input and a storage-weight vector (stored information).

In the middle of the figure, the arrows between SOM modules encode the correlations between them. The link-weights between different SOM modules and the link-weights from the SOM modules to the output decision layer are continuously trained during normal operational runs. The output decision layer is depicted as the grids on the right side of the figure (pointed to by arrows and also circled by an oval shape). The link-weights are adjusted on a reward/punishment principle. Specifically, for the weights of links to the output layer, if the output of the particular output neuron is desired, the link weight for that neuron is rewarded by a small non-zero increment, while if the output is not desired, the link weight is punished by a small non-zero decrement. The link-weights between SOMs are trained in a similar way.

The network is designed with multiple output layers, and each layer consists of two neurons, so each layer represents a binary classifier: whether the final auction price is within certain $*X* range or not. The value of *X* in this study is set as 50. In other words, the price predictor is trained to predict if an auction's final price is in a $50 range or not, such as ($185, $235], rather than a specific numeric value. Since the minimum price in the collected data set is $135, and the maximum price is around $410, there are 6 output layers in the neural network.

Because the actual final price is a specific number, while the expected price is defined using a range, to compare these prices, the actual price is compared to the average price of the expected range. For example, if a predicted price falls in the range ($185, $235], $210 is used as the comparator. In this study, we define "higher" as at least $25.01 more and "lower" as at least $25.01 less.

A total of 1600 auctions were randomly selected from auctions that were tested to be shill free, based on the shill analysis to be introduced in Section C.4. After training the neural network on 1000 auctions and testing it on another 600 auctions, the neural network achieved a precision of 95.5%. This means that, given an auction, the price predictor can determine with a small chance of error if the final auction price will fall in a price span of $50. In this experimental study, we only focus on predicting the final price of one specific category of items. We leave the work of predicting the final price of general items as future work.

4. Shill analysis

Because this work aims to present an empirical study of how the difference between final auction price and expected auction price can provide clues for shill bidding, shill analysis results are needed to feed the logistic regression model so as to examine the hypotheses H_1, H_2, H_3, and H_4. Chapter IV introduced a shill certification method based on the mathematical theory of evidence, Dempster-Shafer (D-S) theory. Six bid-level properties and two auction-level properties are quantified to compute the *belief of shill* for every bidder in an auction. The bid-level properties include the time of a bidder's last bid in an auction, the bidder's concurrent bidding activities, the bidder's reputation score, the bidder's average bid increment, the bidder's winning ratio, and a bidder's affinity for the seller. The auction-level properties include the number of bids and the starting price of the auction.

In the experiments, we compute the shill score for every bidder in an auction. The shill score for an auction is defined as the highest shill score among the bidders. The auction is suspected to involve shill bidding when the shill score for the auction is higher than 0.9; while the auction is considered to be free of shills when the shills score for the auction is less than 0.5. For more details on this approach, refer to Chapter IV.

D. Hypothesis Tests and Results

In this section, we present the model development procedures and results of the hypotheses tests. As stated earlier, chi-square test of independence and logistic regression analysis are carried out on the collected auction data that directly relate to the research questions.

1. Testing of hypothesis H_1

To test that whether shill bidding is associated with the difference between final auction prices and expected auction prices, the collected auction data are analyzed using chi-square statistics. The chi-square ($\chi 2$) test is a non-parametric statistical method that can be used to test the significance of association between two variables. The $\chi 2$ value is computed from a contingency table, which records the counts or frequencies of different combinations of attribute values. As a rule of thumb, statistics references Mason et al. [MLM98] recommend the use of chi-square test only if (1) all cells in the contingency table have expected values greater than 1, and (2) at least 80% of the cells in the contingency table have expected values greater than 5. The $\chi 2$ statistics is approximately distributed with a chi-square distribution, which is often tabulated in statistics texts and available in statistical software packages.

Given an $r \times c$ contingency table, the $\chi 2$ is computed as follows:

$$\chi^2 = \sum \frac{(E_{ij} - O_{ij})^2}{E_{ij}} \qquad (38)$$

$$E_{ij} = \frac{R_i * C_j}{N} \qquad (39)$$

where $\sum$ is taken over all cells of the table, E_{ij} is the expected frequency for the cell in the i^{th} row and the j^{th} column, O_{ij} is the observed frequency in the cell, R_i is the total number of subjects in the i^{th} row, C_j is the total number of subjects in the j^{th} column, and N is the total number of subjects in the whole table. The degree of freedom is equal to $(r\text{-}1)(c\text{-}1)$ where r is the number of rows, and c is the number of columns.

The numbers of auctions that are detected to involve shill bidding and not involve shill bidding for different relationships between the actual auction price and the expected auction price are provided in Table XIV. Recall that auctions' prices are predicted in ranges. When comparing expected price with actual price, the expected price takes the value of midpoint of the expected range. A total of 192 WII online auctions are sampled for this study.

TABLE XIV.
THE DISTRIBUTION OF AUCTION DATA

	Shill	**Non-shill**
Actual Price -Expected Price > 25	44	22
| Actual Price - Expected Price | ≤ 25	1	61
Actual Price - Expected Price < -25	4	60

Since the Hypothesis H_1 is non-directional, the first row and third row of Table XIV are merged in order to apply a non-directional chi-square test. Table XV shows a contingency table for testing Hypothesis H_1. The number in parenthesis is the expected number for that cell, calculated using Eq. (39).

TABLE XV.
THE CONTINGENCY TABLE FOR SHILL BIDDING AND AUCTION PRICES

	Shill	**Non-shill**	**∑row**
| Actual Price - Expected Price | > 25	48 (33.18)	82 (96.82)	130
| Actual Price - Expected Price | ≤ 25	1(15.82)	61(46.18)	62
∑column	49	143	192

Applying Eq. (38) to the statistical data in Table XV, we get $\chi 2 = 27.54$ and degree of freedom = 1. The result is significant at 99% confidence level. In other words, the result indicates that there is a strong relationship between final auction prices and shill bidding. Therefore, Hypothesis H_1 is rejected.

2. The logistic regression model

To better study the relationship between final auction prices and shill bidding, an empirical model based on logistic regression is developed. The reason why we choose logistic regression (aka logit model) is that the outcome variable in logistic regression is binary or dichotomous, such as shill or non-shill. In addition, logistic regression allows a discrete outcome to be predicted from values of one or more variables that may be continuous, discrete and binary, or a mixture of any of these [HL00].

The general form of the logit model is shown as in Eqs. (40) and (41), in which *f(z)* represents the expected value of the outcome variable given the independent variables, x_i .

$$f(z) = \frac{e^z}{1+e^z} \qquad (40)$$

$$z = \alpha + \sum \beta_i x_i \qquad (41)$$

where α is the model intercept, and the β's are the regression parameters for the independent variables.

The logit transformation is defined in terms of *f(z)* as in Eq. (42).

$$g(z) = \ln[\frac{f(z)}{1-f(z)}] = \alpha + \sum \beta_i x_i \qquad (42)$$

The transformation is important because *g(z)* has many of the desirable properties of a linear regression model.

The function *g(z)* is linear in its parameters that may be continuous, and may range from -∞ to +∞, depending on the range of x_i. The conditional distribution of the outcome variable follows a binomial distribution with the probability given by *f(z)*. For more details, refer to [HL00].

3. The conceptual model

We will build a model to predict whether an auction involves shill bidding using its actual and expected auction price. Our goal is to show that online auction users can easily determine if an auction is trustworthy (i.e., no presence of shill bidding) with the help of an auction price predictor. Recall that the price prediction model described in Section C.3 is trained by four groups of features such as item, seller, bid and category specific features, and each feature-group includes several detailed features such as the seller's reputation, date and time when an auction begins, starting bids, etc. Since the price prediction model has already factored in many crucial features except shill bidding, we focus on studying the relationship of shill bidding and online auction price, specifically, how shill bidding might be detected by the difference between actual auction price and expected auction price. This explains why only auction prices are considered in our model.

First, we define two indicator variables as in Eqs. (43) and (44).

$$x_1 = \begin{cases} 1, \text{ if actual price - expected price} > 25 \\ 0, \text{ otherwise} \end{cases} \tag{43}$$

$$x_2 = \begin{cases} 1, \text{ if actual price - expected price} < \text{-25} \\ 0, \text{ otherwise} \end{cases} \tag{44}$$

The reason why x_1 *and* x_2 are treated as categorical variables is that the relationship of higher or lower is more straightforward and meaningful for bidders to identify shilling behaviors than numbers. For example, it may be too subtle for bidders to understand the difference between \$30 higher and \$55 higher. However, when bidders are provided a range of the lowest price and the highest price, they can easily know if the actual price is out of the reasonable price range. For the data collected from auction listings of the WII game systems, we consider a \$25 difference as significant because each predicted price range spans \$50, and an actual price is in the expected range only if it is \$25 (or less) higher or lower than the expected price.

The logit model [HL00, GKR94] is defined as the following:

$$prob[ShillBidding = 1] = f(x_i) = \frac{e^{\alpha + \beta_1 x_1 + \beta_2 x_2}}{1 + e^{\alpha + \beta_1 x_1 + \beta_2 x_2}} \tag{45}$$

Eq. (45) implies the relationship between the probability of shill bidding and the difference of actual price and expected price as shown in Eq. (46).

$$prob[ShillBidding = 1] = \begin{cases} \frac{e^{\alpha}}{1+e^{\alpha}}, \text{if } |\text{actual price - expected price}| \leq 25 \\ \frac{e^{\alpha+\beta_1}}{1+e^{\alpha+\beta_1}}, \text{if actual price - expected price} > 25 \\ \frac{e^{\alpha+\beta_2}}{1+e^{\alpha+\beta_2}}, \text{if actual price - expected price} < -25 \end{cases} \quad (46)$$

4. Results

We used the statistical package *R* (http://www.r-project.org/) to perform logistic regression. The model parameters were estimated using the *maximum-likelihood* method. Table XVI shows the parameter estimation results.

TABLE XVI.

LOGIT MODEL PARAMETER ESTIMATION RESULTS

Model Variables	Estimated Coefficient	Std. Error	Z Value	Pr(>|Z|)
β_1	4.804	1.041	4.614	3.95e-06***
β_2	1.403	1.132	1.239	0.215
α (Intercept)	-4.111	1.008	-4.079	4.53e-05***
Note: ***=significant at p<0.0001; Log Likelihood = -62.09163 (df=3);				

As we can see from Table XVI, β_1 is statistically significant, however, β_2 is not. This means that higher-than-expected auction price has a significantly positive relationship to shill bidding while the lower-than-expected auction price does not indicate likely shill bidding. Therefore, we accept H_4, but we cannot accept H_3, which states a contradictory hypothesis to H_4. In addition, we found no evidence to reject H_2 (The conclusions remain valid even with Bonferroni correction [HT87] for multiple corrections). Table XVII shows the predicted probabilities of shill bidding and non-shill bidding for different relationships of actual auction price and expected auction price.

TABLE XVII.

PREDICTED PROBABILITY OF SHILL BIDDING

x_1	x_2	Relationship	Prob. of Shill Bidding	Prob. of non-shill bidding
0	1	**Actual Price - Expected Price< -25**	0.0625	0.9375
0	0	**| Actual Price - Expected Price | ≤ 25**	0.0161	0.9839
1	0	**Actual Price -Expected Price > 25**	0.6667	0.3333

The probability of shill bidding is 66.67% when the actual price of an auction is significantly higher than the expected auction price. This figure is quite suggestive. Honest bidders could possibly reduce their risks of being victimized to a great extend if they quit auctions with actual price significantly higher than the expected price. . Meanwhile, when the auction price is less-than or equal-to the expected auction price, the likelihood of shill bidding is much lower but the likelihood of non-shill bidding is quite high. Thus, the lower-than-expected or equal-to-expected auction prices are also valuable signals, indicating an auction is likely to be trusted.

5. Model fit

Now we consider the effectiveness of the model to predict the outcome. This is referred to as goodness of fit. There is no generally accepted goodness of fit measure, such as R-square for linear regression, for binary outcome models such as logit [KK98]. However, chi-square and deviance are often adopted by researchers to access the goodness of fit for logistic regression. The test statistic examines whether the model with exploratory variables fits significantly better than the null model. For our model ($\chi 2=93.92$, $p<0.0001$), the chi-square of 93.92 with two degrees of freedom and an associated p-value of less than 0.0001 demonstrate that the model as a whole fits significantly better than a model without the exploratory variable. Thus, we have convincing evidence that x1 is a significant variable in predicting shill bidding. Note this chi-square test is different from the chi-square test of independence presented in Section D.1. For details of the chi-square and deviance test, refer to [HL00].

E. Implications and Threats to Validity

1. Implications

Our findings have several practical implications. Shill bidding can be detrimental to online auctions in the long run. Since honest bidders tend to bid at a price much lower than their evaluation when they factor in shill bidding, sellers will not be able to sell items at the proper market prices and have to reduce the prices. As a result, auction houses will earn less commission fees. Therefore, auction houses need to actively and aggressively detect online auction fraudsters. Our research results offer a time-efficient way to access the likelihood of shill bidding in online auctions. Since auction houses have access to related historical auction data, it is possible for auction houses to utilize the data to obtain an accurate predicted auction price for each auction. By comparing the final auction price to the predicted auction price, an auction can be classified into one of the two groups: auctions that have higher-than-expected prices and auctions that have lower-than or equal-to expected prices. An auction in the first group might be then subjected to some appropriate action such as suspension of the auction and further investigation if needed.

For bidders in online auction markets, this research examines how a final auction price might reflect shill bidding and provides simple rules to discover signals of shill bidding. A better understanding of the relationship between final auction prices and shill bidding can help honest bidders protect themselves from being cheated and thus reduce the chance of monetary loss incurred by fraudulent bidding behavior. Perhaps auction houses or third parties should provide price estimation or prediction as a service to assist honest bidders in diagnosing auctions for shill bidding.

When a bidder finds that the final price of an auction is higher than the expected price, bidders can decide whether they want to continue participating in the auction.

The difference between the final auction price and the expected auction price is in fact reliable in identifying auctions involving shilling behavior because there is a very slight chance for a shill to counteract this indicator. If shills want to counteract this indicator, they can either devise some ways to impair the accuracy of price prediction or lower the final auction prices intentionally. However, neither way is practical to be put into action. To impair the accuracy of the price prediction, shills need to pollute the training data but they can neither know which auctions will be included in the training data nor constantly place shilling bids without being caught. Reducing the final auction price is against the purpose of shill bidding so shill bidders probably will not take actions to achieve this. Therefore, it appears that the difference between final auction price and expected auction price is a reliable indicator of shill bidding.

2. Threats to validity

During the experiments and results interpretation, we held certain assumptions that could lead to validity threats. In this section we list the major issues that could pose some threats to the validity of our arguments and findings. The main threats to validity in this study relate to auction price prediction, shill detection, and data collection, each detailed below.

Expected Price The statistic analysis process was carried out with the assumption that the expected price range of an auction is known precisely when an auction ends. As introduced in Section C of Chapter VI, the expected auction price is the same as the auction price predicted by a trained neural network. Hence, the accuracy of the auction price predictor may be a potential threat to validity of the statistic analysis results. In other word, if an auction's price cannot be predicted in high accuracy, the hypotheses test results in this dissertation may be invalid. To limit this threat, the LARMSTAR neural network based price predictor is carefully trained and tested using real auction data. It can predict the final auction price in a range with high precision. However, to generalize this approach to a broader category, more research efforts are still needed to refine the accuracy of auction prices prediction.

Detection of Shill The data used for logistic regression analysis depends critically on the accuracy of identifying shill bidders in an auction. An ineffective and low-precision shill detection technique could affect the hypotheses test results and in turn affect the effectiveness of the rules derived from the hypotheses tests. Nonetheless, the Dempste-Shafer theory of evidence based shill detection approach [DSX09, DSX10b] has been compared with manual investigation results and demonstrated to be effective. Thus, this threat should be limited.

Data Collection Including only auction data from WII gaming systems that were listed in eBay platform limits the possibility to generalize the application of this approach to other categorization of auction items. This also introduces the risk of missing information that could be helpful in improving auction prices prediction. However, as most auctions have similar listing structures, and the gaming system auctions features plenty of accessories and console combinations, the lessons and experiences that are learned from gaming system auctions can be applied to other categories of items. We will leave generalization of this approach to other category of items as future work.

VII. CONCLUSION AND FUTURE WORK

A. Conclusions

The research reported in this dissertation involved three major topics: 1) We provided a state of art review of the online auction fraud, especially shill bidding; 2) We discussed reasoning under uncertainty for shill detection in online auctions using Dempster-Shafer theory; 3) We explored price comparison as a possible approach to identifying shill bidding in online auctions.

With the prevalence of Internet auctions, auction fraud has become one of the major concerns in electronic commerce. Shill bidding happens during the auction process and is often covered up. Hence it often makes victims suffer without notice. Shill bidding produces undesirable effects not only on the auction participants but also on the auction mechanism itself as a resource allocation market. In the worst case scenario, shill bidding could lead to auction market failure. Because most of the existing online auction systems suffer from user distrust, trustworthy systems that could provide reliable services are highly desired. Current work on Internet auction fraud prevention and detection has taken a simplistic approach, which is not rigorous or complete enough to solve the problem. To prevent in-auction fraud, robust auction rules need to be proposed by economists. On the computer technology side, there is a need of airtight transaction process design to foil the efforts of fraudsters. In this dissertation, we summarized the indicators of shill bidding, and pointed out that because no single indicator will be accurate or strong enough to assure the presence of in-auction fraud, a combinatorial way using multiple indicators would be more effective and precise.

For the second topic, based on the conceptual framework of Dempster-Shafer theory, a unique practical shill detection approach has been proposed. This method in essence takes evidence from different levels, i.e., auction-level and bid-level, into consideration. The knowledge from auction properties and bidding behaviors are represented and quantified. Using Dempster's rule of combination, we combined evidence that enforces each other and resolved the conflicts between different pieces of evidence. The case study shows that our proposed approach is accurate and practical for real world deployment.

For the third topic of this dissertation, we studied whether auction users can infer shill-bidding behavior from the difference between actual auction price and expected auction price. By employing chi-square test of independence and a logistic regression model, we examined the contrary predictions made by extensions to existing auction theory in an attempt to explain how bidders can infer shill bidding from the difference of final selling price. The results show that the relationship of final auction price and expected auction price could be considered as a reliable indicator of shill bidding. We also found that a lower-than or equal-to expected final auction price is quite persuasive in concluding unlikely shill bidding and a higher-than-expected final auction price suggests possible shill bidding. We believe that the rules derived from the hypotheses tests in this dissertation are helpful and applicable for both auction houses and auction users. Honest bidders can protect themselves from being cheated and reduce the risk of monetary loss in online auctions. In addition, auction houses can adopt the rules to complementing existing shill detection techniques and enhancing the confidence of auction users towards auction houses.

B. Future Work

The volume of transactions in online auctioning business is compounding each year, and unfortunately so is shill bidding. Now is the time to act to reduce and stop in-auction fraud. We hope that the ideas from many researchers summarized in the review can help auction policy makers and information technology designers develop future trustworthy environments for online auctions.

Future research can be continued to complement the work completed in this dissertation. For example, the work on reasoning under uncertainty can be extended by the design of a shill detection agent with self-learning capability so that the parameters used in our approach can be optimized automatically. One can also consider user profiles to assist the process of shill detection. We believe that the Dempster-Shafer theory, as a theoretically generalized Bayesian inference method, can provide a practical approach and enhance system performance for shill detection in online auctions.

To extend the research on price comparison, one can further study auction data for different categories of items, refine our price prediction techniques and further validate the signals that can help identify shill bidding activities. It would also be useful to study how the proposed method can be implemented in actual auction environments so that auction houses can provide trustworthy services to their customers.

CITED LITERATURE

[A02] M.R. Albert., "E-Buyer Beware: Why Online Auction Fraud Should Be Regulated," *American Business Law Journal*, vol. 9, no. 4, 2002, pp. 575-643.

[AAP03] Y. W. Ahn, H. J. Ahn, and S. J. Park, "Knowledge and Case-Based Reasoning for Customization of Software Processes - A Hybrid Approach," *International Journal of Software Engineering and Knowledge Engineering,* vol.13, no. 3, 2003, pp. 293-312.

[AJM97] A. J. Menezes, *Handbook of Applied Cryptography*, PC Van Oorschot, SA Vanstone, 1997.

[B92] L.S. Bagwell, "Dutch Auction Repurchases: An Analysis of Shareholder Heterogeneity," *Journal of Finance*, vol. 47, no. 1, 1992, pp. 71-105.

[BH03] P. Bajari, and A. Hortaçsu, "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions," *Rand Journal of Economics,* vol. 2, 2003, pp.329-355.

[BJZ05] B.K. Bhargava, M. Jenamani, and Y.H. Zhong, "Counteracting Shill Bidding in Online English Auction," *International Journal of Cooperative Information Systems*, vol. 14, no. 2-3, 2005, pp. 245-263.

[BP02] S. Ba and P. Pavlou, "Evidence of the Effect of Trust Building Technology in Electronic Markets: Price Premiums and Buyer Behavior," *MIS Quarterly*, vol. 26, no. 3, 2002, pp.243-268.

[CA04] Consumer Affairs, "Shill Bidding Exposed in Online Auctions," November 2004, Retrieved on July 5, 2009 from http://www.consumeraffairs.com/news04/shills.html.

[CD00] C. Dellarocas, "Immunizing Online Reputation Reporting Systems against Unfair Ratings and Discriminatory Behavior," *Proc. of the ACM Conference on Electronic Commerce*, Minneapolis 2000, pp.150-157.

[CD03] C. Dellarocas, "The Digitization of Word-of-Mouth: Promise and Challenges of Online Feedback Mechanisms," *Management Science*, vol. 49, no. 10, 2003, pp. 1407–1424.

[CFL10] A. Clark, C. Fox and S. Lappin, *The Handbook of Computational Linguistics and Natural Language Processing*, Wiley-Blackwell, 2010.

[CPF06] D.H. Chau, S. Pandit, and C. Faloutsos, "Detecting Fraudulent Personalities in Networks of Online Auctioneers," *Principles and Practice of Knowledge Discovery in Database*, 2006, pp. 103-114.

[CSK05] C. Chua, D.W. Straub, H.M. Khoo, S. Kadiyala, and D. Kuechler, "The Evolution of e-Commerce Research: A Stakeholder Perspective," *Journal of Electronic Commerce Research*, vol. 6, no. 4, 2005, pp. 262-280.

[CW96] E.M. Clarke and J.M. Wing, "Formal Methods: State of the Art and Future Directions," *ACM Computing Surveys (CSUR)*, vol.28, no. 4, 1996, pp. 626-643.

[CW04] C. Chua and J. Wareham, "Fighting Internet Auction Fraud: An Assessment and Proposal," *IEEE Computer*, vol. 37, no. 10, 2004, pp. 31-37.

[CX08] K. Chui and R. Xwick, "Auction on the Internet-A Preliminary Study," Manuscript available at http://repository.ust.hk/dspace/handle/1783.1/1035, July 7, 2008

[CV05] T. M. Chen, and V. Venkataramanan, "Dempster-Shafer Theory for Intrusion Detection in Ad Hoc Networks" *Proc. of the IEEE Internet Computing*, 2005, pp. 35-41.

[D68] A. P. Dempster, "A Generalization of Bayesian Inference," *Journal of the Royal Statistical Society*, Series B vol. 30, 1968, pp. 205-247.

[DSX09b] F. Dong, S. M. Shatz, and H. Xu, "Combating Online In-Auction Fraud: Clues, Techniques and Challenges," *Computer Science Review (CSR),* vol. 3, no. 4, Nov. 2009, pp. 245-258.

[DSX09] F. Dong, S. M. Shatz, and H. Xu, "Inference of Online Auction Shills Using Dempster-Shafer Theory," *Proc. of the 6th International Conference on Information Technology: New Generations (ITNG 2009)*, April 2009, pp. 908-914.

[DSX10] F. Dong, S. M. Shatz, and H. Xu, "An Empirical Evaluation on the Relationship Between Final Auction Price and Shilling Activity in Online Auctions." *Proceedings of 22nd Intl. Conf. on Software Engineering & Knowledge Engineering*, San Francisco Bay, 2010, pp. 286-291.

[DSX10b] F. Dong, S. M. Shatz, and H. Xu, "Reasoning Under Uncertainty for Shill Detection in Online Auctions Using Dempster-Shafer Theory," *International Journal of Software Engineering and Knowledge Engineering (IJSEKE)*, vol.20, no. 7, 2010, pp. 943-973.

[DSX11] F. Dong, S. M. Shatz, H. Xu, and D. Majumdar "An Empirical Study of Online Auctions with Unexpected Prices," To appear in *Electronic Commerce Research and Applications (ECRA)*.

[eba08] eBay, "Retracting a Bid at eBay," eBay Help, eBay Inc., 2008, Retrieved on July 7, 2008 from http://pages.ebay.com/help/buy/bid-retract.html

[ebay08] eBay, "Shill Bidding Examples," eBay Help, eBay Inc., 2008, Retrieved on July 7, 2008 from http://pages.ebay.com/help/policies/seller-Shill-bidding.html

[ebay09] eBay, "eBay Services: Selling and Buying Tool: Power Seller Welcome," Retrieved on July 5, 2009 from http://pages.ebay.com/services/buyandsell/welcome.html

[F07] B.J. Feder, "Jeweler to Pay $400,000 in Online Auction Fraud Settlement," *The New York Times*, June 9, 2007, Retrieved on July 7, 2008 from http://www.nytimes.com/2007/06/09/business/09auction.html

[FT90]D. Fudenberg and J. Tirole, *Game Theory*, MA: MIT Press, Cambridge, 1990.

[FTC03] Federal Trade Commission, "Internet Auction Fraud Targeted by Law Enforcer: Thousands of Consumers Fleeced by Auction Scams," Office of Public Affairs, Federal Trade Commission, April 2003, Retrieved on July 7, 2008 from http://www.ftc.gov/opa/2003/04/bidderbeware.shtm

[FXV12] B. J. Ford, H. Xu, and I. Valova, "A Real-Time Self-Adaptive Classifier for Identifying Suspicious Bidders in Online Auctions," To appear in *The Computer Journal (COMPJ)*, 2012.

[G04] S. Grazioli, "Where Did They Go Wrong? An Analysis of the Failure of Knowledgeable Internet Consumers to Detect Deception Over the Internet," *Group Decision and Negotiation*, vol. 13, no. 2, 2004, pp. 149-172.

[G07] D Graupe, *Principles of Artificial Neural Networks* (2nd Edition) World Scientific Publishing Co. Pte. Ltd., Singapore, 2007.

[GA08] Consumer Affairs, "Shill Bidding Exposed in Online Auctions," Retrieved on July 7, 2008 from http://www.consumeraffairs.com/news04/shills.html

[GK98] D Graupe, and H Kordylewski, "A Large-Memory Storage and Retrieval Neural Network for Adaptive Retrieval and Diagnosis," *International Journal of Software Engineering and Knowledge Engineering*, vol. 6, no. 1, March 1998, pp. 115-138.

[GKR94] J. Geweke, M. Keane, and D. Runckle, "Alternative Computational Approaches to Inference in the Multinomial Probit Model," *Review of Economics and Statistics*, vol. 76, no. 4, Nov.1994, pp. 609-632.

[GM87] D.A. Graham and R.C. Mashall, "Collusive Bidder Behavior at Single-Object Second-Price and English Auctions," *Journal of Political Economy*, vol. 95, no. 6, 1987, pp. 1237-1239.

[GS04] R. Ghani and H. Simmons, "Predicting the End-price of Online Auctions," *Proceedings of* I*nternational Workshop on Data Mining and Adaptive Modelling Methods for Economics and Management held in conjunction with the 15th European Conference on Machine Learning* (ECML/PKDDD, 2004, Pisa, Italy

[GS06] D.G. Gregg and J.E. Scott, "The Role of Reputation Systems in Reducing On-Line Auction Fraud," *International Journal of Electronic Commerce*, vol. 10, no. 3, 2006, pp. 95-120.

[GS08] D.G. Gregg and J.E. Scott, "Typology of Complaints about eBay Sellers," *Communications of the ACM* (CACM), vol. 51, no. 4, 2008, pp. 69-74.

[GT08] B. Gavish and C.L. Tucci, "Reducing Internet Auction Fraud," *Communications of the ACM* (CACM), vol. 51, no. 5, 2008, pp. 89-97.

[HK01] J. Han and M. Kamber, *Data Mining: Concepts and Techniques*, Morgan Kaufmann Publishers, 2001.

[HPW07] D. V.Heijst, R. Potharst, and M. V. Wezel, "A Support System for Predicting eBay End Prices," *Decision Support Systems*, vol. 44, no. 4, 2007, pp. 970-982.

[HT87] Y Hochberg, and A. Tamhane, Multiple *Comparison Procedures*. New York: Wiley, 1987.

[HL00] D. W. Hosmer, and S. Lemeshow, *Applied Logistic Regression*, 2nd edition, Wiley-Interscience Publication, 2000.

[JZB07] M. Jenamani, Y. Zhong, and B.K. Bhargava, "Cheating in Online Auction - Towards Explaining the Popularity of English Auction," *Electronic Commerce Research and Applications*, vol. 6, no. 1, 2007, pp.53-62.

[K04] P. Klemper, *Auctions: Theory and Practice*, Princeton University Press, 2004.

[KDS07] G. Kosmopoulou and D. G. De Silva The Effect of Shill Bidding upon Prices: Experimental Evidence. *International Journal of Industrial Organization,* 25, 2007, 291-313.

[KK98] M. Kamstra and P. Kennedy, "Combining qualitative forecasts using logit," *International Journal of Forecasting*, Elsevier, vol. 14, no.1, 1998, pp.83-93.

[KPP02] B. A Kitchenham, S. L. Pfleeger, L. M. Pickard, , P. W Jones, D. C. Hoaglin, K El Emam, and J Rosenberg, "Preliminary Guidelines for Empirical Research in Software Engineering," *IEEE Trans. on Software Engineering*, vol. 28, no. 8, 2002, pp.721-734.

[KW03] R. J. Kauffman and C. A. Wood, "Running up the Bid: Detecting, Predicting, and Preventing Reserve Price Shilling in Online Auctions," *Proc. of the 5th International Conference on Electronic Commerce*, 2003.

[KW05] R. J. Kauffman and C. A. Wood, "The Effects of Shilling on Final Bid Prices in Online Auctions," *Electronic Commerce Research and Applications*, vol. 4, no. 1, 2005, pp. 21-34.

[L99] D. Lucking-Reiley, "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet." AMERICAN ECONOMIC REVIEW, vol. 89, no. 5, 1999, pp. 1063-1080.

[LAH08] D. Lim, P. Anthony and C. Ho, "Agent for Predicting Online Auction Closing Price in a Simulated Auction Environment," *Proc. of Pacific Rim International Conference on Artificial Intelligence - PRICAI*, 2008, pp. 223-234.

[MLG12] A. Mukherjee, B. Liu and N. Glance, "Spotting Fake Reviewer Groups in Consumer Reviews," *Proc. of the 21st international conference on World Wide Web*, April 16–20, 2012, Lyon, France.

[M07] R. McMillan, "Auction Fraud Tops Net Crime," *PC World*, May 2007, Retrieved on July 7, 2008 from http://www.pcworld.com/article/129940/auction_fraud_tops_net_crime.html

[MLM98] Mason, R. D., Lind, D. A. and Marchal, W. G. *STATISTICS: An Introduction*, 5th edition. Duxbury Press, 1998.

[MM87] R.P. McAfee and J. McMillan, "Auctions and Bidding," *Journal of Economic Literature*, vol. XXV, 1987, pp. 699-738.

[MM05] F.M. Menezes and P.K. Monteiro, *An Introduction to Auction Theory*, Oxford University Press, 2005, pp. 10-11.

[MW82] P. Milgrom and R. Weber, "A Theory of Auctions and Competitive Bidding," *Econometrica*, vol. 50, no. 5, 1982, pp. 1089-1122.

[NB08] A. Nikitkov and D. Bay, "Online Auction Fraud: Ethical Perspective," *Journal of Business Ethics*, vol. 79, no. 3, 2008, pp. 235-244.

[NCL08] National Consumer League, "Online Auctions Survey-Summary of Findings," Retrieved on July 7, 2008 from http://www.nclnet.org/shoppingonline/auctionsurvey.htm

[PCW07] S. Pandit, D.H. Chau, S. Wang, and C. Faloutsos, "Netprobe: a Fast and Scalable System for Fraud Detection in Online Auction Networks," *Proc. of World Wide Web*, 2007, 201-210.

[PS03] R. Porter and Y. Shoham, "On Cheating in Sealed-Bid Auctions," *Proc. of the ACM Conference on Electronic Commerce*, 2003, pp. 76–84.

[PXG07] Patel, R., Xu, H. and Goel, A. Real-Time Trust Management in Agent Based Online Auction Systems. *Proceedings of the 19th International Conference on Software Engineering and Knowledge Engineering*, July 9-11, 2007, Boston, USA, 244-250.

[RCG05] S. Rubin, M. Christodorescu, V.J. Ganapathy, T. Griffin, L. Kruger, H. Wang, and N. Kidd, "An Auctioning Reputation System based on Anomaly Detection," *Proc. of the 12th ACM Conference on Computer and Communications Security*, 2005, pp. 270–279.

[RKZ00] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman, "Reputation Systems," *Communications of the ACM*, vol. 43, no. 12, 2000, pp. 45-48.

[RO02] Roth, A. E. and Ockenfels, A. Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet. *American Economic Review* 92, 4, September 2002, 1093-1103.

[RTK90] M.H. Rothkopf, T.J. Teisberg, and E.P. Kahn, "Why are Vickrey Auctions Rare?" *Journal of Political Economy*, vol. 98, no. 1, 1990, pp. 94-109.

[RZ02] P. Resnick and R. Zeckhauser, "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System," *The Economics of the Internet and E-Commerce*, vol. 11 of Advances in Applied Microeconomics, 2002.

[S76] G. Shafer, *A Mathematical Theory of Evidence*, Princeton University Press, 1976.

[SJS03] H. S. Shah, N. R. Joshi, A. Sureka, and P. R. Wurman, "Mining eBay: Bidding Strategies and Shill Detection," *Proc. of the 4th International Workshop on Mining Web Data for Discovering Usage Patterns and Profiles*, July 2003.

[TR07] J. Trevathan and W. Read, "Detecting Collusive Shill Bidding," *Proc. of International Conference on Information Technology: New Generations*, 2007, pp.799-808.

[TR09] J. Trevathan and W. Read, "Detecting Shill Bidding in Online English Auctions," *Handbook of Research on Social and Organizational Liabilities in Information Security*, M. Gupta and R. Sharman (eds.), Information Science Reference, 2009, pp. 446-470.

[V61] W. Vickrey, "Counter speculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, vol. 16, 1961, pp. 8-37.

[W01] M.J. Woolridge, *Introduction to Multiagent Systems*, John Wiley & Sons, Inc. NY, 2001.

[WHW01] W.L. Wang, Z. Hidvègi, and A.B. Whinston, *Shill Bidding in English Auctions*, Technical report, Emory University, 2001, Available at http://oz.stern.nyu.edu/seminar/fa01/1108.pdf

[WHW02] Wang, W. L., Hidvègi, Z. and Whinston, A.B. Shill Bidding in Multi-Round Online Auctions. *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*, January 2002.

[WWW98] P.R. Wurman, W.E. Walsh, and M.P. Wellman, "Flexible Double Auctions for Electronic Commerce: Theory and Implementation," *Decision Support Systems*, vol. 24, no. 1, 1998, pp. 17-27.

[XSB08] H. Xu, S. M. Shatz, and C. K. Bates, "A Framework for Agent-Based Trust Management in Online Auctions," *Proc. of the 5th International Conference on Information Technology: New Generations (ITNG 2008)*, April 2008, pp. 149-155.

[XBS09] H. Xu, C. K. Bates, and S. M. Shatz, "Real-Time Model Checking for Shill Detection in Live Online Auctions," *Proc. of the International Conference on Software Engineering Research and Practice (SERP'09)*, July 13-16, 2009, Las Vegas, Nevada, USA, pp. 134-140.

[XC07] H. Xu and Y.-T. Cheng, "Model Checking Bidding Behaviors in Internet Concurrent Auctions," *International Journal of Computer Systems Science & Engineering*, vol. 22, no. 4, 2007, pp. 179-191.

[YS00] X. Yi and C.K. Siew, "Secure Agent-Mediated Online Auction Framework," *International Conference on Consumer Electronics*, 2000, pp. 114-115.

[Z07] R. Zeithammer, "Strategic Bid-Shading and Sequential Auctioning with Learning from Past Prices," *Management Science*, vol. 53, no. 9, 2007, pp. 1510 - 1519.

[ZZF08] B. Zhang, Y. Zhou, and C. Faloutsos, "Toward a Comprehensive Model in Internet Auction Fraud Detection," *Proc. of the 47th Hawaii International Conference on System Sciences*, 2008, 79.

VITA

NAME: Fei Dong

EDUCATION: B.S., Software Engineering

Harbin Institute of Technology, China, 2006

PUBLICATIONS: **Fei Dong**, Sol M. Shatz, Haiping Xu, and Dibyen Majumdar "Price Comparison: A Reliable approach to identifying shill bidding in online auctions?," To appear in *Electronic Commerce Research and Applications (ECRA*); available online 19 December 2011.

Fei Dong, Sol M. Shatz, and Haiping Xu, "An Empirical Evaluation on the Relationship Between Final Auction Price and Shilling Activity in Online Auctions," *Proceedings of 22nd Intl. Conf. on Software Engineering & Knowledge Engineering*, San Francisco Bay, 2010, 286-291.

Fei Dong, Sol M. Shatz, and Haiping Xu, "Reasoning Under Uncertainty for Shill Detection in Online Auctions Using Dempster-Shafer Theory," *International Journal of Software Engineering and Knowledge Engineering (IJSEKE)*, Vol. 20, Issue. 7, 2010, pp.943-973.

Fei Dong, Sol M. Shatz, and Haiping Xu, "Combating Online In-auction Fraud: Clues, Techniques, and Challenges," *Computer Science Review (CSR),* Vol. 3, Issue 4, Nov. 2009, pp. 245-258.

Fei Dong, Sol M. Shatz, and Haiping Xu, "Inference of Shill Bidding in Online Auction Systems with Dempster-Shafer Theory," *Proc. of the 6th International Conference on Information Technology: New Generations* (ITNG 2009*)*, Las Vegas ,USA, pp. 908-914.

Hsinyi Jiang, Tien N. Nguyen, Carl K. Chang, and **Fei Dong**: "Traceability Link Evolution Management with Incremental Latent Semantic Indexing," *Proc. of the 31st Annual International Computer Software and Applications Conference (COMPSAC 2007)*, Beijing, China, pp. 309-316.