December 2000

# Economic Mechanism Design for Securing Online Auctions

Wenli Wang
*Emory University*

Zoltan Hidvegi
*IBM Corporation*

Andrew Whinston
*University of Texas*

# ECONOMIC MECHANISM DESIGN FOR SECURING ONLINE AUCTIONS

**Wenli Wang**
Goizueta Business School
Emory University
U.S.A.

**Zoltán Hidvégi**
Enterprise Server Group
IBM Corporation
U.S.A.

**Andrew B. Whinston**
Department of Management Science and Information Systems
University of Texas, Austin
U.S.A.

**Abstract**

*Enhancing e-commerce security through computing technology alone is not sufficient. E-commerce designers should apply economic mechanisms to design proper digital processes that accommodate new perspectives raised in e-commerce. For instance, traditional auction mechanisms, such as the Generalized Vickrey Auction, are vulnerable to false-name bidding, an online fraud exploiting the lack of authentication over the Internet. We develop a Sealed-bid Multi-round Auction Protocol (S-MAP), which sells multi-unit identical goods. S-MAP is not only robust against false-name bidding but also simple and efficient.*

## 1. RESEARCH QUESTIONS

The increased popularity of Internet online auctions has become one of the most noticeable phenomena in the emerging digital economy. Although auctions have long been applied in a remarkable range of situations in pre-Internet commerce, their scope is far less than that of online auctions. Auctions can be applied whenever there is a need to establish individual prices for each item sold. For Internet transactions, such a need prevails. The wide accessibility of the Internet introduces an unpredictable scope of demand and supply, which favors dynamic price discovery. The inexpensive computing power allows for low-cost auctions, which makes such mechanisms possible and popular.

Unfortunately, the increased popularity of online auctions has been accompanied by inevitable growing pains. According to the statistics from the Internet Fraud Watch, online auction fraud has become the number one Internet fraud over the last two years. In 1999, online auction fraud accounted for 87% of the reported incidents, increasing from 68% in 1998. In fact, online auctions seem to attract fraud. Fraud related to online auctions soared between 1998 and 1999 by 76% while fraud related to other types of online transactions declined by 44%.

Online auction fraud may take many forms. Undelivered products, inaccurate description of auctioned items, and fraudulent comments undermine the trust between auction sellers and buyers and affect the credibility of online auction intermediaries. One thing to notice is that these forms of fraud are not specific to auctions. They exist for fixed-price online sales as well. Our research does not consider these general "lack of trust" issues but focuses on false-name bidding, a form of auction fraud that is rooted in the uniqueness of online auctions, that is, variable pricing executed in a distributed computing environment.

Variable pricing aims to discover the efficient and stable price for a product when private valuation of the product is uncertain. Such flexibility and dynamic ability are a strength of an auction but also a weakness because they leave room for illegitimate price manipulation. On the other hand, the distributed computing environment of the Internet provides new ways of illegitimate price manipulation, often exploited by dishonest sellers and bidders. A popular exploitation is false-name bidding, which takes advantage of the lack of authentication over the Internet by bidding under false identities. A bidder can create several identities and "collude" among herself in order to maximize her surplus.

Take the Generalized Vickrey Auction (GVA) as an example. GVA is a traditional incentive compatible combinatorial auction mechanism, which selects an allocation of a set of goods to a set of agents. Its objective function is to maximize social surplus, which is calculated as the sum of the valuations for the set of goods obtained by bidders. For each bidder we calculate how much she contributes to the social surplus by calculating the difference between the optimal social surplus and the surplus that can be achieved in her absence. Each bidder gets credit for this contribution to the social surplus, and only pays her bid minus this credit. Of course, the auctioneer only knows the bids of the participants and not their true willingness to pay, so to calculate the social surplus and payments, he assumes that bidders bid their true valuations. This is a correct assumption, since in such an auction it is a dominant strategy for each bidder to bid their true valuations.

In a reserve-price GVA, the social surplus of an allocation is calculated as the sum of bids corresponding to the allocation minus the reservation price for the set of goods allocated to bidders. The payment is calculated as before; each bidder's bid is discounted by her contribution to the modified social surplus.

Despite the incentive compatibility and other strength of GVA, GVA is vulnerable to online auction fraud. GVA assumes that each agent tries to maximize her own valuation without colluding with other agents. However, the assumption no longer holds for online GVA auctions, where agents can easily collude among themselves by creating fictitious identities. Therefore, GVA is no longer incentive compatible because an agent can reduce her payment by submitting bids under her false names; that is, even though agent $x$'s valuation is not included in the calculation of $x$'s payment, $x$ can still manipulate her payment through her valuations under false identities.

> ***Example One.*** *On a reserve-price GVA auction with reserve price $r$ < \$10, there are two tickets for sale and two bidders, Alice and Bob. Alice would like to buy two tickets for \$30 or one ticket for \$15. Bob would like to buy two tickets for \$20 but he does not want a single ticket.*

**Table 1. A Simple GVA Auction**

| Bidder | Qty. | Bid |
|---|---|---|
| Alice | 2 | \$30 |
| Bob | 2 | \$20 |
| Alice | 1 | \$15 |
| Allocation | Alice | 2 |
| | Bob | 0 |
| Payment | Alice | \$20 |
| | Bob | \$0 |
| Utility | Alice | \$10 |
| | Bob | \$0 |
| | Seller | \$20-2r |
| Social Surplus | | \$30-2r |

The optimal allocation of example one gives both tickets to Alice, which achieves a social surplus of \$30-2$r$. If Alice had not bid, the optimal allocation would have given two tickets to Bob for \$20-2$r$ social surplus. Therefore, Alice's bid improves the social surplus by \$10. Alice's payment is \$20, which is her bid minus \$10.

> ***Example Two.*** *On a reserve-price GVA auction, where r < \$10, there are two tickets for sale and three bidders, Alice, Fanny, and Bob. Both Alice and Fanny would buy one ticket for \$15 each. Bob would like to buy twotickets for \$20 but he does not want a single ticket.*

**Table 2. A GVA Auction with False-name Bidding**

| Bidder | Qty. | Bid |
|---|---|---|
| Alice | 2 | – |
| Fanny | 2 | – |
| Bob | 2 | \$20 |
| Alice | 1 | \$15 |
| Fanny | 1 | \$15 |
| Bob | 1 | – |
| Allocation | Alice | 1 |
| | Fanny | 1 |
| | Bob | 0 |
| Payment | Alice | max{\$5, r} |
| | Fanny | max{\$5, r} |
| | Bob | 0 |
| Utility | Alice | min{\$10, \$15-r} |
| | Fanny | min{\$10, \$15-r} |
| | Bob | 0 |
| | Seller | max{\$10-2r, 0} |
| Social Surplus | | \$30-2r |

Here, the optimal allocation gives one ticket to Alice and one ticket to Fanny for \$30-2$r$ social surplus. If Alice had not bid, there would have been two possible allocations: one allocation would have given one ticket to Fanny and the social surplus, in this case, would have been \$15-$r$; the other allocation would have given two tickets to Bob and achieved \$20-2$r$ social surplus. If \$20-2$r$ ≥ \$15-$r$, which happens if $r$ ≤ \$5, the optimal allocation would have given two tickets to Bob. In this case, Alice's bid improves the social surplus by \$10 and her payment is \$5. If r > \$5 and Alice had not bid, the optimal allocation would have given one ticket to Fanny; therefore, Alice's bid improves the social surplus by \$30-2$r$ – (\$15-$r$) = \$15-$r$, thus her payment is \$15 – (\$15-$r$) = $r$. In the latter case, Fanny also obtains one ticket and pays the same as Alice and the payment is the larger of \$5 and the reserve price $r$.

Now let us assume that Fanny is really a false identity created by Alice. In this case, Alice really obtains two tickets for the larger of \$10 and 2$r$. Alice's and Bob's valuations are the same as in example one, the allocation is also identical, but Alice's payment is less than \$20, for $r$ < \$10. This means that Alice can reduce her payment by creating a false identity and splitting her bid between her identities.

These illegal practices of false-name bidding are hard to detect because of the extreme difficulty to authenticate an agent in a distributed Internet environment. Most online auction houses do rely on some kind of computing technologies for authentication, such as IP address monitoring or false-name registration detection through statistical analysis. However, dishonest agents can still possibly circumvent these technologies.

In general, e-commerce developers and computer scientists have been trying to develop and refine technologies (i.e., firewalls, cryptographic protocols) to enhance e-commerce security and prevent online fraud. Wang et al. (2000) have suggested the use of formal verification methods, such as model checking, to help guarantee the correctness of e-commerce systems at the code level. Unfortunately, hackers discover new computing ways to attack or circumvent e-commerce systems every day and correctness through model checking is only as complete as users' specifications. It seems that technology alone cannot guarantee the secure operations of e-businesses. Instead of building a defense against certain fraudulent activities, we should accept their existence and create economic mechanisms that can eliminate their negative effects.

## 2. RESEARCH METHODOLOGY AND OBJECTIVES

Applying mechanism design is a proactive way to handle online fraud. Following the revelation principle, mechanisms can be designed to discourage agents from using anonymity maliciously. In some sense, mechanism design can be regarded as a way of providing incentives for agents to voluntarily reveal their true identities. This is more important than the typical concept of authentication: verify one is who she claims to be. Even though current authentication methods may apply the most rigorous cryptographic protocols to verify agents' identities, it might strongly authenticate false identities and cannot prevent these false identities from damaging the desired system outcome.

We expect that economic theories would apply to the e-commerce world even better than to physical commerce. Economic theories work best when the participating agents are rational. In e-commerce, agents are often automated and operate on well-defined rational rules instead of acting on emotions as humans sometimes do. Therefore, the actual outcome of an e-commerce mechanism can be much closer to the theory than in traditional commerce.

Despite the fit between mechanism design and e-commerce, challenges still exist. Mechanisms that function well in traditional commerce do need to be updated to accommodate new perspectives raised in e-commerce. As an example, traditional mechanism designs assume a fixed set of known agents (e.g., sellers and buyers) who normally look out for their own interests with precise objective functions. However, in e-commerce, such a set has to be enlarged to include some unknown or unrevealed agents such as the following:

(1) Additional internal agents created by the participating agents for strategic purpose, such as false-name bidders created by a participating agent.

(2) External malicious agents (e.g., hackers) whose interests are to disturb the mechanism, cause problems for other agents, and manipulate the process outcome.

The need for update is inevitable because environmental change requires the modification of trading rules (Wilson 1987). Traditional mechanisms, if applied for online processes, have to be rethought with respect to distributed computing issues.

Another inevitable fact is the trade-off between optimization results and added constraints. Traditional mechanisms discussed in economics are often first-best choices. They are based on mathematical optimization, abstract from execution environment and technology. Consideration of special environments and special technologies will lead to sacrifices in objectives, such as the falls in social surplus. This is because of the added constraints, which always worsen the optimization results. Hence, our goal is to search for second-best mechanisms that are still nearly optimal in the e-commerce setting. To achieve resilience to online fraud, trade-offs have to be made to the best traditional mechanisms.

Even though false-name bidding in traditional physical auctions does exist, its scope is very limited because it is impossible for a bidder to bid under two different identities in one auction room and it is difficult to hire someone and make him/her act as expected all the way through the auction. Therefore, little literature on auction theory has examined false-name bidding and suggested mechanisms against this type of fraud. With the development of the Internet economy, few economists have recognized the issues related to false-names. Friedman and Resnick (forthcoming) argued that there is an inherent social cost to free name changes and suggested a dues-paying strategy and a cryptographic mechanism to limit name changes and encourage pseudonym commitments. Contrary to the approach of Friedman and Resnick, we are interested in how to prevent agents from taking advantage of creating false identities.

Yokoo et al. (2000) were the first to take into account false-name bidding in auctions and introduced a combinatorial auction protocol, called the Leveled Division Set (LDS) protocol, that is robust against this type of fraud. LDS uses a set of predefined division sets of the auctioned goods to restrict the possible allocations. The division sets are arranged into levels, where each level is a "refinement" of the previous levels. The lowest level that allows the allocation of some goods is chosen to calculate the outcome of the auction. This protocol guarantees that bidders cannot benefit from submitting several bids under false identities, but to achieve that, it drastically reduces the social surplus achieved by the auction.

The efficiency of LDS depends heavily on the choice of division sets and the reservation prices. Unfortunately, Yokoo et. al. did not give any advice on how to choose the best leveled division sets and reservation prices that either minimize computation complexity or maximize social surplus or both.

## 3. CURRENT STATUS

We designed a Sealed-bid Multi-round Auction Protocol (S-MAP), which sells multiple-unit identical goods and is robust against false-name bidding. S-MAP is defined below:

> The auctioneer has N identical goods for sale with the reserve price r for each item. Let $2^n$ be the highest power of 2 that is less than or equal to N: $2^n \le N < 2^{n+1}$. The auction is conducted in at most n + 1 rounds, i = 0,...,n. At round i, each bidder can only submit a single bid for a bundle of $2^{n-i}$ items of goods and label the bid either "conditional" or "unconditional." A conditional bid is considered only if the current round is the last round. All the bids at the last round n are considered unconditional. Let $N_i$ be the number of goods for sale at the beginning of round i. For i = 0,...,n, let $I_i = [N_i/2^{n-i}]$, where [x] denotes the largest integer not greater than x. By definition, $N_0 = N$ and $I_0 = 1$.
>
> Round i is conducted as follows:
>
> 1. If $I_i = 0$, no bundles of goods are for sale and no bids are accepted, and the auction proceeds to round i + 1.
> 2. If there are $I_i$ or less unconditional bidders who bid at least the reservation price, i.e., $2^{n-i}r$, each of them obtains $2^{n-i}$ goods at the reservation price. Conditional bidders do not get anything and the auction proceeds to round i + 1. If i = n or if there is no goods left for sale, the auction terminates.
> 3. If there are more than $I_i$ unconditional bidders who bid at least the reservation price, the bidders, including the conditional bidders, are ordered in a decreasing bid order. The $I_i$ highest bidders get $2^{n-i}$ goods each. All of them pay the $I_{i+1}$ th highest bid, and the auction terminates.

The key point is that S-MAP prefers to sell larger bundles which guarantees that bidders can never gain by splitting their bids. Our protocol also has a low computational complexity, O(*b* log *N*), where *b* is the number of bids and *N* is the number of goods for sale. We have already analyzed the strengths and weaknesses of S-MAP through mathematical analysis and examples. In future research, we will implement S-MAP in an online ticket sales prototype. We will also extend our S-MAP design concept into a robust market mechanism.

## References


Friedman, E. J. and Resnick, P. "Social Costs of Cheap Pseudonyms: Fostering Cooperation on the Internet," *Journal of Economics and Management Strategy*, forthcoming.

Wang, W., Hidvégi, Z., Bailey, A. D., and Whinston, A. B. "E-Process Control and Assurance Using Model Checking," *IEEE Computer*, October 2000.

Wilson, R. "Game-Theoretic Analyses of Trading Processes," Chapter 2 in *Advances in Economic Theory: Fifth World Congress*, Cambridge, UK: Cambridge University Press, 1987, pp. 30-70.

Yokoo, M., Sakurai, Y. and Matsubara, S. "Robust Combinatorial Auction Protocol Against False-name Bids," in *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, 2000.