Sorting out Trustworthy and Untrustworthy Bidders

Waqas M. Khan and Muhammad Anwar Usman
SZABIST
Karachi, Pakistan

Abstract: *This paper focuses on those auctions which are not enforceable at all. In such auctions, there is always a risk that the winning bidder might fail to complete the commitment and parties involved can not depend on external methods of control mechanism for supporting the transaction. We propose two different methods in which the bidders directly or indirectly reveal their trustworthiness. The first method is based on discerning bidding rules that separate trustworthy and untrustworthy bidders. In this method, the auctioneer offers two different sets of bidding rules, which are designed in such a way that all trustworthy bidders choose one while all untrustworthy bidders choose another. This provides the auctioneer a set of trustworthy bidders, so that he can transact with them. The second method is the generalization of Vickery auction in the case of untrustworthy bidders. We try to prove that if the winning bidder is measured to have the trustworthiness of second highest bidder, declaring the trustworthiness truthfully becomes one's dominant strategy. We expect that the proposed methods can be used to reduce the market cost of trust management and to help the designers of agents to avoid most of the market failures caused by lack of trust.*

Keywords: *Trust in e-commerce, trustworthy bidders, generalized Vickrey Auction*

1. INTRODUCTION

For multilateral trading, without market intermediaries, auctions are the only way that has been extensively used for price determination, particularly in those areas where bidders' private information is the major factor determining strategic behavior.

As mechanisms for distributed optimization, auctions can offer several computational challenges. For example, determining the winners in combinatorial auctions is a complex optimization problem that has been recently studied [1], [2], [3], [4]. Quite a few bidding languages have been created in an effort to minimize the communication operating cost [5], [6]. Another important thread tries to identify auction protocols restraining the preferences that are to be revealed by bidders [7], [8], [9].

Most of the text on auction theory are paying attention on enforceable auctions. It is usually perceived that auction results are obligatory for the auctioneer and bidders. That is, each party behaves as anticipated, and carries out their obligations. Most of the online auctions do not usually meet this supposition. For example, the winning bidder may fail to deliver a product or make a payment. An economic agent, due to lack of incentives or lack of the paper analyzes auctions, which are not completely enforceable. In such auctions, economic agents may fail to carry out their obligations and parties involved cannot rely on external enforcement or control mechanisms for backing up a transaction.

An important feature of these settings is the risk of losses due to fraud, failure, or inability of other parties to fulfill their contractual obligations. Another important feature is the presence of asymmetric information, i.e. untrustworthy agents may not truthfully communicate their private information concerning their contractual intentions or abilities. Revealing such information could hurt future business, and agents usually prefer to over emphasize their trustworthiness in order to enjoy more reimbursement of future co-operation.

In this paper, we examine a reverse multidimensional auction in which a trustworthy buyer faces many sellers with anecdotal degree of trustworthiness. The buyer does not identify the bidders' trustworthiness and has to move first after the auction has been closed, i.e. the buyer has to make the payment without having any guaranties of delivery.

Many applications of mechanism design [10], [11], [12] consider schemes that provide sufficient incentives to parties to truthfully reveal privately known information. The problem in our case is that the auctioneer faces uncertain profits and has to move first without being able to condition his payment on contractual performance. If an auctioneer asks bidders to declare their trustworthiness, they could lie and declare high trustworthiness in order to win the auction. In Section 3, we show that the standard Vickrey auction is unsuccessful to provide bidders with sufficient incentives to truthfully declare their trustworthiness.

In this paper, we study two mechanisms that make agents truthfully declare their trustworthiness. The first mechanism is based on constrained bidding, in which the auctioneer offers different bidding rules for different types of bidders. The rules are designed in a way that it can separate trustworthy from untrustworthy bidders. That is, all trustworthy bidders choose one rule, while untrustworthy bidders choose another. This eliminates information asymmetry, and allows the auctioneer to evaluate bids using the actual bidders' trustworthiness.

The second mechanism is a generalization of the Vickrey auction, in order to cater the case of untrustworthy bidders. In the auction, the highest bidder wins and the rules and

terms of transaction are chosen as if the winner had the trustworthiness of the second-highest bidder.
The auction analyzed in the paper is three-dimensional, where sellers bid on price and quantity, besides reporting their trustworthiness. Multidimensional procurement auctions arise commonly and have been widely studied [13], [14]. The paper is organized as follows: Section 2 provides a brief formalization of trust in the context of e-commerce; Section 3 defines the problem setting in which an auctioneer faces many bidders with anecdotal degree of trustworthiness. A discriminating auction based on several bidding schedules is described in Section 4; and Section 5 presents a generalization of the Vickrey auction to the case of untrustworthy bidders. Finally, the paper concludes by summarizing the results and providing directions for future research.

2. A FORMAL FRAMEWORK OF TRUST

The idea of building trust in different types of applications dealing with transactions is a subject of continuous interest in different areas including game theory and economics [15], [16], multi-agent systems [17], [18], [19] and risk-analysis [20]. The notion of trust is also closely related to the design and implementation of multi-stage safe exchanges [21], [22].
Trust has different connotations and has been used in different meanings in different contexts by different authors. One group of authors [23] considers trust as a belief or cognitive stance that could eventually be quantified by a subjective probability. We give a brief conceptualization of trust that will help avoid confusion and will facilitate further exposition. We suppose that trust is a bilateral relation that involves an entity manifesting trust called the 'trustor'; and an entity being trusted called the 'trustee'.

Further, we assume that there exist events G that cannot be controlled by the trustor and that depends on the trustee. The trustee may have partial or full control over G. The trustor voluntarily puts himself in a position dependent on G, in the sense that the trustor will benefit if G occurs, otherwise he will lose. In other words, the trustor depends on the trustee for some event G that is controlled by the trustee. We assume that trustworthiness could be presented as a measure of the probability of the occurrence of G. For example, the trustee could be an untrustworthy seller and G = {the seller delivers promised merchandize after it has been paid for}. In another example, the trustor could depend on the trustee for some information and G = {the trustee delivers accurate and truthful information}. Another interpretation is G = {the quality of the merchandize meets the buyer's expectation}.

In general, trustworthiness can be identified in two broad types: perceived and actual. Perceived trustworthiness is defined as the trustor's personal belief in occurrence of G, which could be different from the actual trustworthiness; i.e. the real probability of G. For example, an agent might believe that a seller will deliver promised merchandize with probability $\hat{\mu}$, while the actual probability of delivery is μ.

Formally, the trustor's utility function can be denoted by:

$$U(\hat{\mu}, G(p_1, \ldots\ldots, p_n)) \qquad (1)$$

Where, U is the trustor's utility, $p_1 \ldots pn$ are parameters describing the event G, and $\hat{\mu}$ is the degree of perceived trustworthiness, i.e. the degree in which G is expected to happen.
The trustor will be benefited from the event G, if:

$$\frac{\partial\, U(\hat{\mu}, G(p_1, \ldots\ldots, p_n))}{\partial\, \hat{\mu}} \geq 0$$

That is, the trustor benefits from higher trustworthiness. The case of complete trustworthiness is represented by $\hat{\mu} = 1$ and the trustee is completely untrustworthy when $\hat{\mu} = 0$:

$$U(1, G(p_1, \ldots\ldots, p_n)) > 0$$
$$U(0, G(p_1, \ldots\ldots, p_n)) < 0$$

If we assume that utility function is a continuous function of trustworthiness, then there exists a threshold level $\hat{\mu}_0 \in [0, 1]$ that can be used to separate trustworthiness from untrustworthiness:

$$U(\hat{\mu}, G(p_1, \ldots\ldots, p_n)) \geq 0 \qquad \text{for all} \quad \hat{\mu} > \hat{\mu}_0$$

That is, the trustor will always be in a benefited state if the other agent's trustworthiness exceeds the threshold level of trust $\hat{\mu}_0$ which depends on the event G and its parameters $p_1, \ldots\ldots, p_n$. This can be defined a natural participation constraint: the trustor will place trust on the trustee if the trustee's perceived trustworthiness exceeds $\hat{\mu}_o$. The natural participation constraint corresponds to the intuition that an agent will only engage in an interaction if the trustworthiness of the other party exceeds some threshold level, which depends on the interaction context (through parameters $p_1, \ldots\ldots, p_n$) and on the trustor (through the trustor's utility function U). In other words, the threshold $\hat{\mu}_0$ is both objectively and subjectively determined.

Such a formalization of trust is domain independent and captures a wide range of applications, where the trustor believes that the trustee will behave in some expected way specified by the event G. This model is so general enough that it can capture, not only, auctions but also business contracts, negotiations etc. Depending on the context, the perceived (or actual) trustworthiness can be given different interpretations. For instance, it could be the probability of delivery, the probability of high product quality, probability that an agent will follow contract terms etc.

By choosing probability $\hat{\mu}$ (or μ) as a measure of trustworthiness, we do not mean that trust always depends on a single factor. The event G may have a complex structure represented by parameters $p_1, \ldots\ldots, p_n$. Other work [24] the authors experimentally validated is a multidimensional model of trust in on-line exchanges. They showed that the following six factors affect trust: information content, product, transaction, technology, institutions, and consumer-behavior. We assume that all these factors could be combined so as to produce a single measure of an agent's trustworthiness. In other words, we can think of α as a measure of the combined effect of different constituents and determinants of trust.

3. PROBLEM SETTING

This part of paper describes a reverse auction with untrustworthy bidders. A buyer gets bids from sellers with two different levels of trustworthiness α and β; $\alpha < \beta$; α, β ε [0, 1]. Both α and β are assumed to be normalized measures of a bidder's commitment to back up his bids. For the ease of interpretation, α and β could be thought of as probability of delivery, measure of quality, ability, etc. For example, in one interpretation, a less trustworthy bidder will deliver with probability α if he wins the auction, while a more trustworthy bidder will deliver with probability β. A second way to look at α and β is to see them as the sellers' ability or capacity to deliver which could be objectively or subjectively determined.
Each agent submits bid, knows only his own type (α or β), and the set of possible types; and the joint probability distribution over types are common knowledge between the buyer and the sellers. Throughout the paper, we assume that the variation in trustworthiness is significant enough to make a difference and refer to bidders of type α and β as untrustworthy and trustworthy bidders, respectively.

The buyer is supposed to be completely trustworthy and he makes the first move after the auction has been closed. That is, the buyer pays first without knowing the probability of delivery. By moving first the buyer openly discloses type. Each bid specifies an offer of promised quantity q and price p. The buyer and the sellers are risk-neutral, and the buyer derives utility from a bid, (p, q) ε R^2_+:

$$U(p, q, \mu) = V(q, \mu) - p \quad (2)$$

Where, μ is the bidder's trustworthiness; μ ε $\{\alpha, \beta\}$; and $V(q, \mu)$ is the buyer valuation function (denotes the amount of benefit buyer gets from the particular bid); $V_q > 0$; $V_{qq} < 0$; and $V_q(0, p) = 0$ to ensure an interior solution. Subscripts denote partial derivatives: that is, $Vq(q, \mu) = \partial v(q, \mu)/\partial q$, $V_q(q, \mu) = \partial^2 v(q, \mu)/\partial^2 q$ and so forth.

A bidder, upon winning, earns from a bid (p, q) the following profit:

$$W(p,q, \mu) = p - C(q, \mu) \quad (3)$$

Where, W and C are the bidder's utility and cost functions, respectively. We assume $C_{qq} > 0$; $C_\mu > 0$; and $C_{q\mu} > 0$. Thus, both the total and the marginal cost increase with μ. To completely understand the idea behind these assumptions, it is convenient to view one's trustworthiness as a measure of quality or probability of delivery. It is generally perceived that the production costs usually increase with quality, all other things being equal. In addition, trustworthy agents may include added costs for establishing and keeping a good reputation.

If the auctioneer uses a scoring function, which is equal to his utility, defined by Equation 2, and asks bidders to disclose their types, then untrustworthy bidders ($\mu = \alpha$) may have an encouragement to report a higher type ($\mu = \beta$). The main problem is that the scoring function (and the auctioneer's utility) increases in μ. For example, in a standard Vickrey auction, the winner has to match the price and the quantity of the second-score bidder. This, however, does not prevent an untrustworthy bidder from reporting higher trustworthiness. Reporting a higher type of trustworthiness increases the chance of winning the auction without affecting a bidder's utility.

Proposition 1: Truthfully declaring an agent's trustworthiness is not a dominant strategy in a standard Vickrey auction, where agents bid on price and quantity.

The problem with untrustworthy bidders is that the buyer's utility depends on the trustworthiness of the winner, which is only privately known. By declaring higher trustworthiness, an untrustworthy bidder can manipulate the way bids are evaluated. Without knowing the original bidders' types, the auctioneer cannot precisely evaluate the utility of a bid, and, therefore, might determine the auction winner incorrectly. Since the buyer is moving first, he cannot condition his payment on the seller compliance. We assume that the buyer does not have access to indirect indicators of a seller's trustworthiness such as history of previous interactions or reputation database.

4. A SEPARATING AUCTION

In this section, we study the problem of how to perform bid evaluation and winner determination based on information only contained in bids.
A much natural approach to solve the problem with untrustworthy bidders is to assume that the auctioneer adopts a play-safe strategy and decides to secure him against the worst case possible. That is, the auctioneer evaluates bids on the assumption that all bidders are untrustworthy (type α).

Definition 1: In a distrust-based auction, every bidder submits a bid on price and quantity. The auctioneer uses a scoring function S that treats each bidder as untrustworthy:

$$S(p, q) = V(q, \alpha) - p$$

Unfortunately, being overcautious does not help auctioneers to avoid untrustworthy bidders.
Proposition 2: There is a highly positive probability that an untrustworthy bidder wins in a distrust-based auction. If the difference in trustworthiness, β-α, between the two

agent types is sufficiently large, then only untrustworthy bidders win.

Proposition 2 can be explained by using the cost differences between agent types. If the difference in trustworthiness is sufficiently large, trustworthy bidders include sufficiently large costs compared to untrustworthy bidders; which prevents them from submitting competitive bids, and, therefore, from winning an auction.
Another way to solve the problem with untrustworthy bidders is to consider trustworthiness as a random variable and to evaluate bids using its expectation, E (μ). Unfortunately, a similar proposition holds here as well.

Proposition 3: Suppose that an auctioneer evaluates bids according to his expectation of agents' trustworthiness:

$$S(p, q) = V(q, E(\mu)) - p$$

There is a highly positive probability that an untrustworthy bidder wins in such auctions. If the difference in trustworthiness, $\beta-\alpha$, between the two agent types is sufficiently large, then only untrustworthy bidders win.
The above propositions 2 and 3 show that in some cases, trustworthy agents will be driven out of the market, thereby, causing market inefficiency. To fix the problem, we investigate constrained-bidding mechanisms.

Definition 2: In a constrained-bidding multidimensional auction, an eligible bid satisfies a set of constraints on bid parameters. That is, for every eligible bid b (t_1, ……, t_n), we have:

$$\Phi k(t_1, \ldots\ldots, t_n) \text{ for } k = 1, \ldots\ldots, m$$

Where, $\{\Phi k\}^m_{k=1}$ is a set of constraint predicates.
For example, the auction rules can fix the quantity to q_0 and define a minimal and a maximal price:

$$q = q_0; \text{ and } p \ \varepsilon \ [p_{min}, p_{max}] \quad (4)$$

One possible interpretation is that the minimal price is the reservation level for a bidder of a certain type, and the maximal price is the auctioneer's reservation level. In our setting, constraint 4 reduces a two-dimensional auction on price and quantity to a one-dimensional auction on price only.

One important characteristic of constrained auctions is that the bidders' expected utility could be limited by the auction rules in advance. For example, constraint 4 imposes an upper bound, $p_{max} - C(q_0, \mu)$, and a lower bound $p_{min} - C(q_0, \mu)$ for type μ bidders. By choosing a particular set of constraints, the auctioneer can affect the incentive structure of the auction, and therefore, can provide bidders with additional incentives. We will show that in our case, the bidders could be given incentives to reveal directly or indirectly, their type.
We assume that, if a seller faces a choice between two auctions, he will choose an auction, which gives him a better utility range, all other things being equal. For example, if a seller must choose between an auction A_1 with a utility range {2,10} and a auction A_2 with a utility range {0,8}, he would choose A_1; all other things being equal. The idea behind this assumption is that, every bidding strategy for an auction A_2 gives a better-expected utility when applied to auction A_1.

Assumption 1: Given an auction A_1 with a utility range [a^1_{min}, a^1_{max}] and an auction A2 with a utility range [a^2_{min}, a^2_{max}], where the only difference between A1 and A2 is:

$$A^1_{min} > a^2_{min}$$
$$A^1_{max} > a^2_{max}$$
$$A^1_{max} - a^1_{min} = a^2_{max} - a^2_{min}$$

Then a risk-neutral bidder prefers auction A_1 to auction A_2.

In other words, in both auctions, a bidder uses the same strategy set, faces the same opponents and the same rules, with the only difference being the range of strategy payoffs in terms of profit or satisfaction. The utility range of auction A_1 could be viewed as a result of higher transformation of the utility range of auction A_2. Therefore, the two auctions are strategically equivalent with the only difference being the scale of utility measurement. In other words, every bidding strategy has a higher expected utility in auction A_1 than in auction A_2.

Using bidders' preferences for auctions, the auctioneer can distinguish or screen various types of bidders by offering different bid constraints to different types of bidders.

Definition 3: In a separating constrained-bidding auction, two sets of bid constraints are offered by the auctioneer. A bidder chooses one set of constraints and strictly follows the set throughout the auction. All other auction rules remain the same for all bidders. A bidder does not have the ability to change his set of constraints during an auction.

In other words, there are two bidding rules, each bidder chooses and follows only one rule, and all bidders compete with one another. That is, each bidder competes with both the bidders following his rule and the bidders following the other rule. For example, in a separating constrained-bidding auction based on the first-score rule, the bidder with the highest score wins. In the beginning, the auctioneer offers two sets of bid constraints. A bidder either chooses a set of constraints and follows them, or leaves the auction.

According to the next proposition, sometimes it is possible to design two sets of bid constraints so that all trustworthy bidders prefer one set and all untrustworthy bidders prefer the other. Thus, by choosing a set of constraints, bidders disclose their type. This allows the auctioneer to evaluate the utility of each bid and to determine the winner. Since the auctioneer knows the bidders' types, he can associate every trustworthy bid with β and every untrustworthy bid with α.

Proposition 4: If $V_q(0, \alpha) > C_q(0, \beta)$, then there exists a constrained-bidding auction that separates trustworthy from untrustworthy bidders.

According to Proposition 4, auction rules can be designed to eliminate the strategic consequences arising from differences in bidders' types. In such cases, the auctioneer can offer two bidding rules and allow bidders to choose the more beneficial one. According to Proposition 4, the rules can be designed so that the first type bidders choose the first rule and the second type chooses the second one. It can be shown that after choosing a rule, both types of bidders face the same utility range and the same strategic choices.

One example of a separating auction is the following: Instead of bidding on price p and quantity q, bidders are given a choice from two auction rules. The first rule allows bids for fixed quantity q_0 and price $p \in [p_1, p_2]$. The second rule allows bids for quantity q_1 and $p \in [p_3, p_4]$. The quantities and price ranges can be chosen so that to make the auctioneer indifferent between the two rules. That is, both rules offer the same utility range for the auctioneer. While equally profitable to the auctioneer, the auction rules offer different utility to bidders. All trustworthy bidders are better off with the first rule, while all untrustworthy bidders prefer the second rule. By choosing a rule, each bidder unambiguously reveals his type. In this particular example, in order to separate trustworthy from untrustworthy bidders, the auctioneer splits a two-dimensional auction (on price and quantity) into two one-dimensional (on price only) auctions.

It should be pointed out that a separating auction does not prevent untrustworthy bidders from winning. What distinguishes a separating auction from distrust-based and expectation-based auctions is that the auctioneer can exactly evaluate bids and choose the most profitable bid. In addition, when the difference in trustworthiness, β-α, is sufficiently large, trustworthy agents are not driven out of the market, as is the case for the other auctions.

It should be pointed out that a separating auction might not maximize the social welfare. Obviously, some price has to be paid for the possibility to separate agent types. For example, in order to maximize his utility in a second-score auction, the auctioneer will choose bidding schedules with maximal utility range. That is, the auctioneer will choose quantity q^{untr} such that:

$$q^{untr} = \underset{q}{\arg\max}\,(V(q, \alpha) - C(q, \alpha))$$

If the auctioneer knew the type of each bidder, then he could fix the quantity to

$$q^{trust} = \underset{q}{\arg\max}\,(V(q, \beta) - C(q, \beta))$$

Or to

$$q^{untr} = \underset{q}{\arg\max}\,(V(q, \alpha) - C(q, \alpha))$$

depending on which agent type is more profitable for him. It is apparent, that in the case where trustworthy agents offer more utility to the auctioneer, the social welfare is not maximized. If, however, untrustworthy agents are more efficient, then a separating auction is socially optimal. Whether trustworthy agents are more efficient than untrustworthy ones depends on the value, $V(q, \mu)$, and the cost function $C(q, \mu)$. If the social cost of trustworthiness is less than its social value, then trustworthy agents will be more efficient, and vice versa.

5. A GENERALIZATION OF VICKREY AUCTION

In this section, we present a generalization of the Vickrey auction to cater the case of untrustworthy bidders. We drop Assumption 1 and the restriction of having only two types of bidders. The generalized auction is applicable to situations with a continuum of bidder types.

In the generalized auction, each bidder submits a bid on price, quantity, and a declaration of trustworthiness $(p, q, \hat{\mu})$. The auction uses a constrained-bidding schedule where each bidder is required to submit the maximal price for each combination of quantity and price:

$$p = C(q, \hat{\mu}) \qquad (5)$$

The score function is equal to the auctioneer's utility, assuming that every bidder truthfully declares his type, i.e., $\hat{\mu} = \mu$. The winner is the bidder with the highest score (ties are resolved randomly). The winning bidder matches the highest rejected score by choosing a price and a quantity, which generate the same score. That is, the exact price and quantity of the second highest bidder are not required, but only a price-quantity combination that generates the same utility for the auctioneer. The central point of the auction rules is that, in matching the second-highest score, the winner is assumed to have the same type as the highest-rejected bidder. In other words, the winner is allowed to choose a price and a quantity that generate the highest-rejected score using the declared trustworthiness of the highest-rejected bidder.

Definition 4: In the generalized Vickrey auction, each bidder submits a bid $b = (p, q, \hat{\mu})$. Bidding is constrained and eligible bids must satisfy Equation 5. The score is defined as:

$$S(p, q, \hat{\mu}) = V(q, \hat{\mu}) - p$$

The highest score wins. The price p and quantity q are chosen by the winner to satisfy:

$$p - V(q, \hat{\mu}^s) = p^s - V(q^s, \hat{\mu}^s) \qquad (6)$$

$$Vq(q, \hat{\mu}^s) - Cq(q, \hat{\mu}) = 0 \qquad (7)$$

Where, (p^s , q^s , $\hat{\mu}^s$) is the second-highest bid, and $\hat{\mu}$ is the winner's declared trustworthiness. As usual, subscripts denote partial derivatives.

Condition 6 guarantees that the score of winner, $p - V(q, \hat{\mu}^s)$, matches the second highest score, $p^s - V(q^s, \hat{\mu}^s)$. Note that Condition 6 requires the winner to match the second highest score under the assumption that he has the type of the second-highest bidder, $\hat{\mu}^s$. Equation 7 ensures that the marginal cost of the winner is equal to the marginal value, which the auctioneer could have received from the second highest bidder.

Proposition 5: In the generalized Vickrey auction, it is a dominant strategy for each bidder to truthfully report his trustworthiness.

The intuition behind the generalized Vickrey auction is as follows: Equations 6 and 7 define a system of simultaneous equations, which uniquely determine p and q (and hence the utility of the winner) for each declared level of trust-worthiness $\hat{\mu}$. The equations 6 and 7 are defined so that the winner maximizes his utility only if he truthfully declares his trustworthiness $\hat{\mu}$.

The generalized Vickrey auction provides a convenient solution to the problem of trust. During the auction, agents always report their true level of trustworthiness, even if they are untrustworthy. Honest reporting lets the auctioneer know the interaction risk and form realistic expectations about possible outcomes.

6. CONCLUSION

In this paper, we have analyzed a reverse multidimensional auction in which a trustworthy buyer faces sellers with different degrees of trustworthiness. We proposed two mechanisms that make bidders directly or indirectly reveal their trustworthiness. The first mechanism is based on discriminating bidding rules. We have proved that under certain conditions, it is possible to design bidding rules that separate trustworthy from untrustworthy bidders.

The second mechanism is a generalization of the Vickrey auction to the case of untrustworthy bidders. We proved that, if the winner is considered to have the trustworthiness of the second-highest bidder, truthfully declaring one's trustworthiness becomes a dominant strategy.

The mechanisms proposed in this paper provide several advantages. They do not require an estimation of other agents' trustworthiness. This could simplify individual decision-making and save deliberation costs. By eliminating the need to manipulate and speculate about other bidder's trustworthiness, the mechanisms could also simplify the architecture of economic software agents.

Another advantage of the mechanism is that it may reduce the cost of trust management, simplify many complex and costly infrastructures for risk assessment and fraud protection like reputation databases, recommended systems and trusted third parties. In risky environments, the mechanisms could enable mutually beneficial interactions, which are otherwise, costly to enforce or cannot be enforced.

REFERENCES

[1] T. Sandholm, S. Suri, A. Gilpin, and D. Levine. CABOB: A fast optimal algorithm for combinatorial auctions. In Pro-ceedings of the International Joint Conference on Artificial Intelligence, pages 1102-11082, Seattle, 2001.

[2] M. Rothkopf, A. Peckec, and R. Harstad. Computationally manageable combinatorial auctions. Management Science, 44 (8): 1131-1147, 1998.

[3] N. Nisan. Bidding and allocation in combinatorial auc-tions. In Proceedings of ACM Conference on ElectronicCommerce, pages 1-12, 2000.

[4] Y. Fujishima, K. Leyton-Brown, and Y. Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In Proceedings of the International Joint Conference on Artificial Intelligence, pages 548-553, 1999.

[5] H. Hoos and C. Boutilier. The package assignment model. Proceedings of the International Joint Conference on Artificial Intelligence, pages 1211-1217, Seattle, 2001.

[6] T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. Artificial Intelligence, 135:1-54, 2002.

[7] S. Bikhchandani and J. Ostroy. The package assignment model. UCLA, 2001. Working Paper Series, mimeo.

[8] P. Wurman and M. Wellman. AkBA: A progressive, anonymous-price combinatorial auction. In Proceedings of ACM Conference on Electronic Commerce, pages 21-29, 2000.

[9] W. Conen and Sandholm. Partial-revelation VCG mechanism for combinatorial auctions. In Proceedings of the National Conference on Artificial Intelligence (AAAI), 2002.

[10] R. Myerson. Revelation principle. Journal of mathematical economics, 10:67-81, 1982.

[11] Ye Diana Wang, Henry H. Emurian: An overview of online trust: Concepts, elements, and implications.

[12] Paul Klemperer : Auction Theory and its Applications to Economics Reading List. In The Current State of Economic Science, 2, 711-66. S. Dahiya (ed.)), 1999.

[13] Aditya V. Sunderam , David C. Parkes : Preference Elicitation in Proxied Multiattribute Auctions.

[14] Gregory E. Kersten, Sunil J. Noronha, Jeffrey Teich : Are All E-Commerce Negotiations Auctions? In *COOP'2000: Fourth International Conference on the Design of Cooperative Systems,*Sophia-Antipolis, France, May 23-26, 2000.

[15] P. Dasgupta. Trust as a commodity. In D. Gambeta, editor, Making and Breaking Cooperative Relations. Basil Black-well, 1990.

[16] C. Snijders. Trust and Commitments. ICS, 1996.

[17] Tiffany Y. Tang, Pinata Winoto and XiaoLin Niu: Investigating Trust between Users and Agents in A Multi Agent Portfolio Management System: a Preliminary Report.

[18] S. Braynov and T. Sandholm. Contracting with uncertain level of trust. Computational Intelligence, 14(4): 501-514,2002.

[19] Bernardo A. Huberman, Fang Wu and Li Zhang: Ensuring trust in one time exchanges: solving the QoS problem, In Springer 2006.

[20] Erhard John Dobesberger : Multivariate Techniques For Estimating The Risk Of Plant Pest Establishment In New Environments. In NAPPO PRA Symposium, Puerto Vallarta – Mexico, March 2002.

[21] T. Sandholm and V. Ferrandon. Safe exchange planner. In Proceedings of the Fourth International Conference on Multi-agent Systems, pages 255-262, Boston, 2000.

[22] T. Sandholm and X. Wang. Impossibility of safe exchange mechanism design. In Proceedings of the National Conference on Artificial Intelligence (AAAI), pages 338-344, 2002.

[23] A. Josang. Trust-based decision making for electronic trans-actions. In Proceedings of the Fourth Nordic Workshop on Secure Computer Systems (NORDSEC'99), pages 496-502,Sweden, 1999.

[24] Gadi Fibich , Arieh Gavious , Aner Sela : Asymptotic Analysis of Large Auctions with Risk-Averse Bidders. Econometrica, 52:1473-1518, 1984.