Auctions with Untrustworthy Bidders

Sviatoslav Braynov
Department of Computer Science
and Engineering
State University of New York at Buffalo
Buffalo, NY 14260
sbraynov@cse.buffalo.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
sandholm@cs.cmu.edu

Abstract

The paper analyzes auctions which are not completely enforceable. In such auctions, the winning bidders may fail to carry out their obligations, and parties involved cannot rely on external enforcement or control mechanisms for backing up a transaction. We propose two mechanisms that make bidders directly or indirectly reveal their trustworthiness. The first mechanism is based on discriminating bidding rules that separate trustworthy from untrustworthy bidders. That is, the auctioneer offers two types of auction rules which are designed so that all trustworthy bidders choose one bidding rule, while untrustworthy bidders choose another. This allows the auctioneer to discover trustworthy bidders and to transact with them. The second mechanism is a generalization of the Vickrey auction to the case of untrustworthy bidders. We prove that, if the winner is considered to have the trustworthiness of the second-highest bidder, truthfully declaring one's trustworthiness becomes a dominant strategy. We expect the proposed mechanisms to reduce the cost of trust management and to help agent designers avoid many market failures caused by lack of trust.

1 Introduction

Auctions have been extensively used in e-commerce as a means for price determination for multilateral trading without market intermediaries. They are particulary useful in markets with incomplete and asymmetric information, where the bidders' private information is the main factor determining strategic behavior.

As mechanisms for distributed optimization, auctions can offer several computational challenges. Determining the winners in combinatorial auctions, for example, is a complex optimization problem that has been recently studied [20, 17, 16, 8]. Several bidding languages have been proposed in an effort to reduce the communication overhead [9, 18]. Another important thread tries to identify auction protocols limiting the preferences that are to be revealed by bidders [1, 24, 6].

Most of the literature on auction theory has focused on enforceable auctions. It is usually assumed that auction results are binding for the auctioneer and bidders. That is, each party behaves as expected, and carries out their obligations. Many on-line auctions do not usually meet this assumption. For example, a winning bidder may fail to make a payment or deliver a product. An economic agent can fail to perform his task or to meet his commitment due to lack of incentives, lack of ability, or circumstances beyond his control. Internet users still fear the possibility of fraud, identity change, misuse of private information, etc. Complaints about online auction transactions have skyrocketed, accounting for 87 percent of the Internet fraud reports made to the Internet Fraud Watch in the first six months of 2002, compared to 70 percent in 2001 (www.fraud.org).

This 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. That is, there is no outside enforcement mechanism or a trusted third party that could compel the auction winners to perform their duties. An auction winner is free not to deliver, or not to provide the promised quality at the promised time, and so forth.

An important characteristic of these settings is the risk of losses due to failure, fraud, or inability of other parties to fulfill their contractual obligations. Another important characteristic is the presence of asymmetric information. That is, untrustworthy agents may not truthfully communicate private information concerning their contractual abilities or intentions. Disclosing such information could hurt future business, and agents usually prefer to overstate their trustworthiness in order to enjoy more benefits of future cooperation.

In the paper, we analyze a reverse multidimensional auction in which a trustworthy buyer faces many sellers with varying degree of trustworthiness. The buyer does not know the bidders' trustworthiness and has to move first after the auction has been closed. That is, the buyer has to make the payment without having any guaranties of delivery. Such a setting raises several important questions. The first one is how to evaluate bids. The profitability of each bid depends on the bidder's trustworthiness (including quality of product, timeliness of delivery, etc.), which is privately known to the bidder. That is, the auctioneer does not have enough information to evaluate bid profitability, unless bidders truthfully declare their actual trustworthiness. Another question is what kind of incentive-compatible schemes are possible and what is their economic efficiency.

Many applications of mechanism design [15, 13, 14] consider schemes that provide sufficient incentives to parties to 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 bellow, we show that the standard Vickrey auction fails to provide bidders with sufficient incentives to truthfully declare their trustworthiness.

In the paper, we study two mechanisms that make agents truthfully reveal 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 to 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 to the case of untrustworthy bidders. In the auction, the highest bidder wins and the terms of trade 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 frequently and have been extensively studied [5, 3]. For example, many defense procurement auctions involve multidimensional bids on promised technical characteristics, delivery date, estimated project costs, etc.

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 varying degree of trustworthiness. A discriminating auction based on several bidding sched-

ules is described in Section 4. In the auction, agents reveal their trustworthiness by choosing different schedules. 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 concept of trust has been a subject of continuous interest in different research areas, including multi-agent systems [12, 4, 2], game theory and economics [7, 23] and risk-analysis [22]. The notion of trust is also closely related to the design and implementation of multi-stage safe exchanges [19, 21].

Trust has different connotations and has been used in different meanings in different contexts by different authors. Many authors [10] consider 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 assume 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 is an event Γ that the trustor cannot control and that depends on the trustee. The trustee may have partial or full control over Γ.

- The trustor voluntarily decides to put himself in a position dependent on Γ in the sense that the trustor will benefit if Γ occurs, otherwise he will lose.

In other words, the trustor depends on the trustee for some event Γ which is controlled by the trustee. We assume that trustworthiness could be measured by the probability of Γ. For example, the trustee could be an untrustworthy seller and $\Gamma = \{$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 $\Gamma = \{$The trustee delivers accurate and truthful information$\}$. Another interpretation is $\Gamma = \{$The quality of the merchandize meets the buyer's expectation$\}$.

In general, two types of trustworthiness can be identified: perceived and actual. Perceived trustworthiness is defined as the trustor's subjective belief in Γ which could be different from the objective trustworthiness, that is, the objective probability of Γ. For example, an agent might believe that a seller will deliver promised merchandize with probability $\hat{\theta}$, while the actual probability of delivery is θ.

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

$$U(\hat{\theta}, \Gamma(p_1, ..., p_n)) \tag{1}$$

where U is the trustor's utility, $p_1, ..., p_n$ are parameters describing the event Γ, and $\hat{\theta}$ is the degree of perceived trustworthiness, i.e., the degree in which Γ is expected to happen.

The event Γ is favorable to the trustor:

$$\frac{\partial U(\hat{\theta}, \Gamma(p_1, ..., p_n))}{\partial \hat{\theta}} \geq 0$$

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

$$U(1, \Gamma(p_1, ..., p_n)) > 0$$
$$U(0, \Gamma(p_1, ..., p_n)) < 0$$

If we assume that utility is a continuous function of trustworthiness, then there is a threshold level $\hat{\alpha}_0 \in [0, 1]$ that separates trustworthiness from untrustworthiness:

$$U(\hat{\theta}, \Gamma(p_1, ..., p_n)) \geq 0 \qquad \text{for all} \qquad \hat{\theta} \geq \hat{\theta}_0$$

That is, he trustor is always better off if the other agent's trustworthiness exceeds the threshold $\hat{\theta}_0$ which depends on the event Γ and its parameters $p_1, ..., p_n$. This defines a natural *participation constraint*: the trustor will place trust on the trustee (or will voluntarily agree to depend on the trustee) if the the trustee's perceived trustworthiness exceeds $\hat{\theta}_0$. The 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 (the level of acceptable trustworthiness), which depends on the interaction context (through parameters $p_1, ..., p_n$) and on the trustor (through the trustor's utility function U). In other words, the threshold $\hat{\theta}_0$ 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 Γ. The model is general enough to 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 $\hat{\theta}$ (or θ) as a measure of trustworthiness we do not mean that trust always depends on a single factor. The event Γ may have a complex structure represented by parameters $p_1, ..., p_n$. In another work of ours [11] we experimentally validated a multidimensional model of trust in on-line exchanges. We 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 section describes a reverse auction with untrustworthy bidders. A buyer solicits bids from sellers with two different levels of trustworthiness α and β, $\alpha < \beta$; $\alpha, \beta \in [0, 1]$. Both α and β are 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 exogenously or endogenously determined. For example, the seller may depend on a shipping intermediary for a fast delivery. In another interpretation, α or β could be the probability that the contract between the buyer and the seller can be enforced by a third party. Many Internet transactions, for example, are difficult or costly for legal enforcement.

Each bidder knows only his own type (α or β), and the set of possible types and the joint probability distribution over types are common knowledge among the buyer and the sellers. Throughout the paper, we refer to bidders of type α and β as untrustworthy and trustworthy bidders, respectively, assuming that the variation in trustworthiness is significant enough to make a difference.

The buyer is 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. Since by moving first the buyer explicitly discloses his type, the assumption of complete buyer's trustworthiness does not limit the generality of the model.

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, $(p, q) \in \mathbb{R}_+^2$:

$$U(p, q, \theta) = V(q, \theta) - p \qquad (2)$$

where θ is the bidder's trustworthiness, $\theta \in \{\alpha, \beta\}$, and $V(q, \theta)$ is the buyer valuation function, $V_q > 0, V_{qq} < 0$, and $V_q(0, p) = 0$ to ensure an interior solution. Subscripts denote partial derivatives. That is, $V_q(q, \theta) = \frac{\partial V(q, \theta)}{\partial q}, V_{qq}(q, \theta) = \frac{\partial^2 V(q, \theta)}{\partial q^2}$, and so forth.

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

$$W(p, q, \theta) = p - C(q, \theta) \qquad (3)$$

where W and C are the bidder's utility and cost functions, respectively. We assume $C_{qq} > 0, C_\theta > 0$, and $C_{q\theta} > 0$. Thus, both the total and the marginal cost increase with θ. To understand the intuition behind these assumptions, it is convenient to view one's trustworthiness as a measure of quality or probability of delivery. Production costs usually increase with quality, all other things being equal. In addition, trustworthy agents may incur added costs for establishing and keeping a good reputation.

If the auctioneer uses a scoring function equal to his utility, defined by Equation 2, and asks bidders to reveal their types, then untrustworthy bidders ($\theta = \alpha$) may have an incentive to report a higher type ($\theta = \beta$). The problem is that the scoring function (and the auctioneer's utility) increases in θ. For example, in a standard Vickrey auction, the winner has to match the price and the quality of the second-score bidder. This, however, does not prevent an untrustworthy bidder from reporting higher trustworthiness. Reporting a higher type 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.*[1]

The problem with untrustworthy bidders is that the buyer's utility depends on the trustworthiness of the winner, which is only privately known. In this case, the informed winner's trading decision depends on privately held information in a manner that negatively affects the uninformed auctioneer. By declaring higher trustworthiness, a bidder can manipulate the way bids are evaluated. Without knowing bidders' types, the auctioneer cannot precisely evaluate the utility of a bid, and therefore determine the auction winner. 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 reputation database or history of previous interactions.

4 A separating auction

In this section we study the problem of how to perform bid evaluation and winner determination based solely on information contained in bids.

A natural way to approach the problem with untrustworthy bidders is to assume that the auctioneer adopts a play-safe strategy and decides to insure himself against the worst case possible. That is, the auctioneer evaluates bids on the assumption that all bidders are untrustworthy (type α).

---

[1] Due to space limitations the proofs are omitted in this version of the paper

Definition 1 *In a **distrust-based auction**, every bidder submits a bid on price and quality. 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 strictly positive probability that an untrustworthy bidder wins in a distrust-based Dutch, English, first-score, and a second-score sealed bid auctions. If the difference in trustworthiness, $\beta - \alpha$, between the two agent types is sufficiently large, then only untrustworthy bidders win.*

Proposition 2 can be explained using the cost differences between agent types. If the difference in trustworthiness is sufficiently large, trustworthy bidders incur 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, $E(\theta)$. Unfortunately, a similar proposition holds here as well.

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

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

There is a strictly positive probability that an untrustworthy bidder wins in Dutch, English, first-score, and second-score sealed bid auctions. If the difference in trustworthiness, $\beta - \alpha$, between the two agent types is sufficiently large, then only untrustworthy bidders win.

Propositions 2-3 show that, in some cases, trustworthy agents will be driven out of the market, thereby causing a 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*

$$\phi_k(t_1, ..., t_n) \text{ for } k = 1, ..., m$$

where $\{\phi_k\}_{k=1}^m$ 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:

$$q = q_0, \text{ and } p \in [p_{min}, p_{max}] \tag{4}$$

One possible interpretation is that the maximal price is the auctioneer's reservation level, and the minimal price is the

reservation level for a bidder of a certain type. In our setting, constraints (4) reduce a two-dimensional auction on price and quantity to a unidimensional auction on price.

One important characteristic of constrained auctions is that the bidders' expected utility can be ex-ante limited by the auction rules. For example, constraints (4) impose an upper bound, $p_{max} - C(q_0, \theta)$, and a lower bound $p_{min} - C(q_0, \theta)$ for type θ 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 intuition behind this assumption is that every bidding strategy for a 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 $[a_{min}^1, a_{max}^1]$ and an auction A_2 with a utility range $[a_{min}^2, a_{max}^2]$, where the only difference between A_1 and A_2 is*

$$
\begin{aligned}
a_{min}^1 &> a_{min}^2 \\
a_{max}^1 &> a_{max}^2 \\
a_{max}^1 - a_{min}^1 &= a_{max}^2 - a_{min}^2
\end{aligned}
$$

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

In other words, in both auctions, a bidder has the same strategy set, faces the same opponents and the same rules, with the only difference being the range of strategy payoffs. The utility range of auction A_1 could be viewed as a result of affine transformation on 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.

In the next section, we will drop Assumption 1 and propose a generalized Vickrey auction in which truth-telling is a dominant strategy.

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, the auctioneer offers two sets of bid constraints. A bidder chooses a set of constraints and strictly follows the set throughout the auction. All other auction rules remain the* *same for all bidders. A bidder is not allowed 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 choose 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 inidimensional (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, $\beta - \alpha$, 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 may 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:

$$q^{untr} = arg \max_q \left(V(q, \alpha) - C(q, \alpha) \right)$$

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

$$q^{trust} = arg \max_q \left(V(q, \beta) - C(q, \beta) \right)$$

or to

$$q^{untr} = arg \max_q \left(V(q, \alpha) - C(q, \alpha) \right)$$

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, $V(q, \theta)$, and the cost function, $C(q, \theta)$. 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 the Vickrey auction

In this section we describe a generalization of the Vickrey auction to 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 $(p, q, \hat{\theta})$. The auction uses a constrained-bidding schedule where each bidder is required to submit the maximal price for each combination of quantity and price:

$$p = C(q, \hat{\theta}) \tag{5}$$

We assume that the cost function is known to the auctioneer who can then check Condition 5 for each bid and verify its validity. The score function is equal to the auctioneer's utility, assuming that every bidder truthfully declares

his type, i.e., $\hat{\theta} = \theta$. 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 score. For example, if the highest rejected score is 20, the winner must choose quantity and price that generate a score of 20. 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. More formally:

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

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

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

$$p - V(q, \hat{\theta}^s) = p^s - V(q^s, \hat{\theta}^s) \tag{6}$$
$$V_q(q, \hat{\theta}^s) - C_q(q, \hat{\theta}) = 0 \tag{7}$$

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

Condition 6 guarantees that the score of winner, $p - V(q, \hat{\theta}^s)$, matches the second-highest score, $p^s - V(q^s, \hat{\theta}^s)$. 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, $\hat{\theta}^s$. 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 trustworthiness $\hat{\theta}$. The Equations 6 and 7 are defined so that the winner maximizes his utility only if he truthfully declares his trustworthiness $\hat{\theta}$.

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 Conclusions

In the 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' 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 and simplify many complex and costly infrastructures for risk assessment and fraud protection like reputation databases, recommender 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] S. Bikhchandani and J. Ostroy. The package assignment model. UCLA, 2001. Working Paper Series, mimeo.

[2] S. Brainov. An incentive compatible trading mechanism for trust revelation. In *Proceedings of the IJCAI'01 Workshop on Economic Agents, Models and Mechanisms*, pages 62–70, 2001.

[3] F. Branco. The design of multidimensional auctions. *Rand Journal of Economics*, 28(1):63–81, 1997.

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

[5] Y. Che. Design competition trough multidimensional auctions. *Rand Journal of Economics*, 24(4):668–680, 1993.

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

[7] P. Dasgupta. Trust as a commodity. In D. Gambeta, editor, *Making and Breaking Cooperative Relations*. Basil Blackwell, 1990.

[8] 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.

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

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

[11] D. Kim, Y. Song, S. Braynov, and H. Rao. A B-to-C trust model for on-line exchange. In *Americas Conference on Information Systems*, pages 784–787, 2001.

[12] S. Marsh. *Formalizing Trust As a Computational Concept*. PhD thesis, University of Stirling, U.K., 1994.

[13] E. Maskin and J. Riley. Optimal auctions with risk-averse buyers. *Econometrica*, 52:1473–1518, 1984.

[14] R. McAfee and J. McMillan. Competition for agency contracts. *RAND Journal of Economics*, 18(2):296–307, 1987.

[15] R. Myerson. Optimal coordination mechanisms in generalized principle-agent problems. *Journal of mathematical economics*, 10:67–81, 1982.

[16] N. Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of ACM Conference on Electronic Commerce*, pages 1–12, 2000.

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

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

[19] T. Sandholm and V. Ferrandon. Safe exchange planner. In *Proceedings of the Fourth International Conference on Multiagent Systems*, pages 255–262, Boston, 2000.

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

[21] T. Sandholm and X. Wang. (Im)possibility of safe exchange mechanism design. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 338–344, 2002.

[22] P. Slovic. Risk perception and trust. In V. Molak, editor, *Trust: Fundamentals of Risk Analysis and Risk Management*. Lewis Publishers, 1997.

[23] C. Snijders. *Trust and Commitments*. ICS, 1996.

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