An Online Auction Mechanism with Tradeoffs Between Bid Privacy, Cognitive Cost and Number of Rounds

****Draft, February 21, 2003 ****

Helger Lipmaa

Helsinki University of Technology
Laboratory for Theoretical Computer Science
Department of Computer Science and Engineering
P.O.Box 5400, FI-02015 Espoo, Finland
helger@tcs.hut.fi

Abstract. We propose a new, cryptographically protected, multi-round auction mechanism that is specifically tailored for online auctions. Our auction mechanism is designed to provide (in this order) security, cognitive convenience and round-effectiveness. One can vary internal parameters of the mechanism to trade off bid privacy and cognitive costs, or cognitive costs and the number of rounds.
Keywords: auctions, cognitive costs, cryptography, mechanism design, privacy

1 Introduction

Mechanism design has much in common with the construction of cryptographic protocols, both in goals and in methodology. As an important example, significant amount of research has been concentrated on nonmanipulable mechanism design. Here, a mechanism is called *nonmanipulable*, if rational agents will maximize their utility by revealing their true type (or preferences). However, to be applicable in practice, it is not sufficient for a mechanism to be nonmanipulable, it must also satisfy other properties like simplicity and privacy. In this paper, we will show how one can weave cryptography into mechanism design to achieve many desirable properties.

More concretely, we will concentrate on online auctions. Online auctions can be organized over the Internet or a local wireless network. The bidders can use software agents that do the computationally intensive parts of the bidding, while the human beings stay in full control over the prices. The most fundamental new feature of *online* auctions is that the software agents have, as compared to the human beings, the necessary computing power and "willingness" to participate in more resource-consuming auction types. This increases the flexibility of mechanism design, making it possible for the sellers (or auctioneers) to choose auction mechanisms that are technically impossible to implement in conventional auctions. In particular, it becomes possible to use public-key cryptography [DH76] to ensure both the auction correctness and bid privacy.

At the expense of mitigated computational costs, the importance of other properties of auction mechanisms will grow in online auctions. Importantly, *cognitive costs* of computing one's valuation and one's strategy will dominate over the computational costs. Therefore, to further simplify participation in online auctions, one must alleviate

the cognitive costs by devising an auction mechanism that neither requires the bidders to do an elaborated precomputation to calculate their precise valuation (or their best bid, given information about the other bidders), nor extensive online calculations to react properly to the bidding strategies of other participants.

Another important concern in auctions is security, including both the auction correctness and bid privacy. Auction fraud was the most common complaint to Internet Fraud Complaint Centre (IFCC) in 2001, according to the annual report by the IFCC [CoI02]. The number of frauds can be hopefully decreased by devising an auction mechanism with better security properties. For example, an online auction mechanism should be secure against a malicious seller (or auctioneer) and different possible attacks (shills, collusive bids, jump bidding). Additionally, only a minimal amount of information should be leaked to the auctioneer or other bidders.

All the mentioned goals are not achievable at the same time. As we will see in Section 2, one must trade off both between cognitive costs and resource-effectiveness, and between cognitive costs and privacy. In particular, to have small cognitive costs, one should have both a large number of rounds but also some (otherwise unnecessary) privacy leakage.

We believe that a good auction mechanism should emphasize privacy and correctness over the cognitive costs. The main (although somewhat informal) reason for our belief is that it is easier to define what is the privacy (and what is a privacy leak) than to model the cognitive costs, the latter being largely a psychological notion. For example, if instead of a single bid, information about two competing bids will be leaked, then this is certainly a privacy leak. But does it help to alleviate the cognitive costs? Can the bidders use this additional information to regulate their information about their own values? Probably yes, but how much exactly do they gain? If one cannot guarantee that deliberate loss of privacy will decrease the cognitive costs, it is better not to lose any. (Cognitive cost *is* modelled in some publications [Par99,LS01], but there the authors are more concerned with the agents doing the computations, not the human beings.)

Therefore, we argue that when constructing an online auction mechanism, one should first make sure that the auction is Pareto-efficient, correct and (almost-ideally) privacy-preserving. The next goal is to mitigate the cognitive costs as much as possible, without hurting the correctness and the privacy-preserving properties. For example, to minimize the cognitive costs, a mechanism should be nonmanipulable. At last, one should make sure that the mechanism is sufficiently effective—that is, that it does not have more rounds than an English auction, or require superpolynomial-time computations. We will base our mechanism on those guidelines, but we will also introduce parameters that make it possible to have a conscious trade off between the privacy and the cognitive costs, and between the cognitive costs and the number of rounds.

An example mechanism that is tailored for agent-mediated online auctions is *proxy bidding*. Proxy bidding decreases the cognitive cost (as compared to Vickrey auctions) and the bid leakage (as compared to English auctions). We will discuss other desired and existing properties of (online) auctions in Section 2. There, we will point out why proxy bidding is less than ideal.

As a motivating example, note that if in a proxy bidding *all* users use the software agents, then the proxy bidding will become equivalent to a multi-round auction, where

every round is a first-price sealed-bid auction. The highest bidder of a round is declared to be the winner, unless some other bidder wants to continue in the next round. Even such a multi-round mechanism does not completely solve the problem of revealed statistics that is characteristic to proxy bidding, even when every round is cryptographically secured. (For example, this mechanism will reveal information about persons who participate in every round, and therefore leaks partial information about their valuations.) Moreover, being a first-price mechanism, this mechanism is manipulable. And finally, if many bounded-rational people participate, the auction can start from small prices, and progress very slowly. So, while such a multi-round mechanism together with an adequate cryptographic protection increases privacy and efficiency, compared to pure English auctions, it is still not an ideal mechanism.

Our new mechanism handles these problems as follows. First, every round of our mechanism is a second-price mechanism (i.e., a Vickrey auction). This makes the mechanism incentive-compatible. Second, during every round only the currently second highest bid is revealed. The revealing helps to alleviate cognitive costs (compared to a Vickrey auction), and the hiding of other bids protects privacy (compared to an English auction or proxy bidding). Third, our auction mechanism is parameterized by the cognitive error coefficient $0 \leq \varepsilon < 1$, that forces the bidders to precompute their values at least to some extent. Additionally, our mechanism is cryptographically protected, and includes some sensible ending conditions that provide protection against shills and collusive bids. Some protection is also provided against the jump bids.

Our mechanism has the same privacy properties as the cryptographically secured Vickrey mechanism (indeed, the choice $\varepsilon = 0$ results in a Vickrey auction), while the cognitive costs are comparable to the ones in English auctions. See Section 3 for a closer description of our mechanism that is followed by a precise analysis.

As far as we know, this is the first auction mechanism that has been designed from the scratch to provide correctness, bid privacy and cognitive costs at the same time. We think that our work has relevance to classical auction theory, and helps to converge the different research lines of game-theoretic, cognitive and cryptographic properties of auctions.

Road-map. Section 2 gives a short overview of the known auction mechanisms. Section 3 describes our new auction mechanism, followed with discussions and analysis. Section 4 explains the difference with related work. We finish the paper with a section on further work and acknowledgments.

Notation. Let $\mathcal{V} = \{v_1, \ldots, v_V\}$ be the set of possible valuations (bids). Let $\mathcal{B} = \{1, \ldots, B\}$ be the set of all possible bidders. Throughout this paper, let $(X_1, \ldots, X_B)$ be the vector of bids in non-decreasing order, and let Y_i be the bidder whose bid was X_i. In a multi-round auction, let $(X_1^r, \ldots, X_B^r)$ be the list of bids made in the $r \geq 1$-st round, in non-increasing order, and let Y_i^r be the bidder who made the bid X_i^r. We assume that $X_i^0 = 0$ and that (Y_i^0) is an arbitrary permutation of all bidders.

2 State of the Art

An auction mechanism is basically a set of rules (or an algorithm) for an auction that has a motivational ingredient. In particular, nobody should have a negative payoff when following the auction mechanism. We refer to [Kri02] for a good overview of different auction mechanisms. We call a participant (either a bidder or the seller), who dutifully follows the auction mechanism and does not share her private information with other parties, *honest*, as it is common in cryptographic literature.

An ideal auction mechanism should provide at least the next properties. First, *Pareto-efficiency*. The auctioneered item will be awarded to the bidder who values it the most, at least when she follows the auction rules. Second, *resource-effectiveness*. The auction takes a small number of rounds. The auction rules should be sufficiently simple so that the seller and the bidders can follow them in "reasonable" time. Third, *correctness*. A cheating seller will be caught. In particular, she will not be able to increase the final price or change the winner by will. Fourth, *privacy*. No information about the bids of honest bidders will be revealed, except the information that follows from the name of the winner and the contract price. Fifth, *minimal cognitive cost*. The cognitive cost of computing the valuation should be minimal. Other properties are security against shills, collusive bids, jump bidding, etc [Kri02].

Some of the mentioned properties are somewhat contradictory. For example, sealed-bid auctions are resource-effective but they do involve a significant cognitive cost. Namely, the first-price auctions force the bidders to use complicated strategies to compute their bid as a function on the possible valuations of other bidders. Even the Vickrey auctions, that are known to be nonmanipulable (also called *incentive-compatible* [Vic61]), are such only under the independent private-values assumption: that is, only in the case when the private values of clients are independent of each other. Moreover, Vickrey auctions cease to be Pareto-efficient if the bidders are *bounded-rational*, that is, if they do not know their private values.

In general, one cannot assume that the bidders know their valuations, since the auctioned items do not have a well-known utility. The cognitive cost of strategy planning is especially important in online auctions [UPF98,PUF98]. Since other participating costs decrease considerably due to use of software agents, *cognitive costs* of computing one's valuation (and one's strategy) start to dominate. Therefore, it becomes desirable to decrease cognitive costs by devising an auction mechanism that neither requires the bidders to do an elaborated homework to compute their precise valuation (or their best bid, given information about the other bidders), nor requires them to do extensive calculations online to react properly to the bidding strategies of other bidders. Such a mechanism should still have other properties like Pareto-efficiency and security (e.g., the bid privacy). Additionally, it should have an "acceptable" computational complexity.

A well-known fact is that the more information about competitors' valuations is revealed to every bidder, the higher will be the contract price [MW82], and possibly the more Pareto-efficient will be the auction. When the bidders are bounded-rational, one-round mechanisms create therefore the least revenues. The largest revenues are in English auctions that are on the other hand, ineffective in the number of rounds.

An interesting trade-off was proposed in [PWZ00], that constructed a two-round sealed-bid auction mechanism (that we call an PWZ mechanism) with the same seller

revenues as the English auctions. Both rounds are a Vickrey auction. The two highest bidders of the first round continue in the second round. The distribution of first round losers' bids will be revealed to the two winners before the second round. The PWZ mechanism is resource-effective. It is also slightly better than the Vickrey mechanism in cognitive cost. However, if the bidders are bounded-rational, then it is not Pareto-efficient. This is since the $(B-2)$ players were never given the second chance, and since the remaining 2 players only got one more chance.

Given this discussion, it may seem that the "ideal" online auction mechanism is indeed the English auction. However, the English mechanism is very ineffective in the number of rounds, since it can last up to V rounds. Of course, one can trade off the "precomputational" (or cognitive) efficiency against the number of rounds since a precise *a priori* approximation of one's valuation will decrease the number of possible human-interacted rounds. However, in English auctions one is usually not motivated jump-bind, since this may potentially increase the contract price significantly.

It is clearly desirable to improve upon the round-ineffectiveness of English auctions. Recently, another mechanism that does that in the agent-mediated case has received a lot of attention. In agent-mediated *proxy bidding*, bidders use a software agent with a fixed upper bound on the price. The agents participate in an English auction until this upper bound has been reached. Only after that the agents consult with their owner, who has to decide whether to continue to bid (by setting a new upper bound) or not (by passing). This can last many rounds, until the final price will not raise anymore. Proxy bidding has smaller cognitive costs than one-shot auctions, and on the other hand, has smaller participation costs than English auctions. Hence, proxy bidding offers a balance between the cognitive cost and the resource-effectiveness of the English and Vickrey auction mechanisms. It is noteworthy that proxy bidding is used very successfully in Internet auctions. As early as in 1999, Lucking-Reiley surveyed 142 auction sites and found that 65 of them use a form of proxy bidding [Luc00, Section VIII.A].

However, even the proxy bidding is not ideal. First, it is a form of first-bid auctions. Therefore, it is manipulable: since the winner pays as much as he bid, the players are motivated to bid less. As already mentioned in the introduction, another big concern is privacy.

Bid Privacy and Correctness in Auctions. While Vickrey auctions are attractive from the theoretical viewpoint, they are rarely used in practice. We already explained the first major reason, the involved high cognitive costs. The second major reason is the possibility of having a cheating seller, who could either (a) change the outcome of auctions (invalidate the *correctness* property), or (b) reveal bidders' private information (invalidate the *privacy* property). As argued in [RTK90,RH95], in the first case, a honest bid taker will not choose a Vickrey auction, while in the second case, a cheating bid taker eventually destroys the trust on which the use of Vickrey auctions depends. Therefore, Vickrey auctions seem to become more widely applicable when secured cryptographically, so that the seller is forced to follow the auction mechanism and no extra information is revealed to him. These observations have motivated a huge body of research on cryptographic Vickrey auction schemes, starting with [NS93].

Clearly, protecting privacy is important also in other auction mechanisms. However, the PWZ mechanism, proxy bidding and English auctions are (designed to be) bad from the privacy viewpoint, since they reveal the bid statistics to alleviate the cognitive cost. If one protects such mechanisms cryptographically so that they will completely protect losers' privacy, these mechanisms will lose their main raison d'tre. Recall that revealing losers' bids made it possible to mitigate cognitive costs! Therefore, one has a unsurprising trade-off between privacy and cognitive costs, or security and convenience.

Cryptographic auction schemes. *Cryptographic auction schemes* are cryptographic algorithms to support concrete mechanisms, that, when correctly followed by a honest party, ensure that certain well-defined privacy/correctness properties will be held w.r.t. her. In particular, a good auction scheme must ensure that neither a cheating auctioneer nor cheating bidders can make the auction non-Pareto-efficient.

Andrew Yao [Yao82] was the first to consider cryptographic (English) auctions. The first cryptographically secure Vickrey auction scheme that provides losers' privacy was proposed in [NS93]. A large number of cryptographic Vickrey auction schemes have been proposed since that. (See [NPS99,LAN02] for some examples and overview of related literature.) Such schemes would satisfy all desired properties that were described in the beginning of this chapter, except that they do not minimize the cognitive cost. In particular, the best cryptographic auction schemes guarantee the auction correctness, and privacy, to the extend required by the auction mechanism.

In the following we will shortly describe a simplified version of the LAN auction scheme by Lipmaa, Asokan and Niemi. The full version of this scheme [LAN02] incorporates, in particular, protection against the replay attacks. The LAN scheme has B bidders, a seller S and an auction authority A. Anybody who wishes to sell something can act as S (this means in particular that no trust can be put on S) while the authority is an established business party with a reputation history. In this scheme, a bid b is encoded as B^b, B being the (maximum allowed) number of bidders. The ith bidder encrypts the encoding B^{b_i} of his bid b_i with A's public key by using a suitable homomorphic encryption scheme, and sends the result to S. S multiplies all the received encrypted bids, and sends the resulting encryption of $B^{\sum_i b_i}$ to A. After decrypting this result, A finds out the bid statistics (that is, how many bidders bid b for any possible bid b) but is not able to connect any bidders with their bids. Then, A sends the second highest bid to S. Every action in this scheme is accompanied with an *efficient* (statistical) zero-knowledge correctness proof. By using recently proposed cryptographic range proofs [Bou00,Lip01], both the bidder-seller and the seller-authority communication complexity of the LAN scheme are of order $\Theta(V \cdot \log_2 B)$ bits, where V is the maximum possible number of different bids.

Summary of auction mechanisms There are many well-known auction mechanisms, like the English, Dutch, first-price sealed-bid and Vickrey [Vic61] auctions. (A description of these mechanisms can be found in [Kri02]) Different auction mechanisms satisfy different desiderata that are summarized in Table 1. We do not know any mechanism-scheme combinations that satisfy all the previously described auction desiderata. Note that not all four desiderata, as described in the beginning of the current section, are

Mechanism	Pareto-e.	Round-effect.	Correctness	Priv.	Cogn. c.
English	+			-	+
Dutch			+	+	
First-price		+	(+)	(+)	
Vickrey [Vic61]	+	+			
Proxy bidding	+	(+)		-	+
PWZ [PWZ00]	+	+		-	(+)
Secure Vickrey	+	+	+	+	
Secure proxy bidding	+	(+)	+	-	+
Our mechanism	+	(+)	+	+	+

Table 1. Comparison of different existing auction mechanisms and our mechanism in the mentioned five categories: A "+" means that the mechanism performs well in this category, "(+)" means that the mechanism enjoys slightly better properties than the unmarked mechanisms, and "-" means that this property is undesirable by the design.

equally important in all situations. Traditionally, one has mainly been stressing the first two properties. In this paper, we will concentrate on online auctions, where, as we will see, the last three properties will gain in importance.

3 New Mechanism

3.1 High-Level Description

Next, we will describe our cryptographically secured multi-round sealed-bid auction mechanism. Discussions and explanations will follow. Some of the following notation was defined at the end of introduction. Let $C_K(x)$ denote a commitment of x by using a suitable commitment scheme.

Setup. Our mechanism is parameterized by a public value $\varepsilon < 1$, selected by the seller S and announced to everybody before the auctions. There are B bidders $1, \ldots, B$, one seller S and the auction authority A. The participants obtain a commitment key, an encryption key and a signature key of other relevant parties, depending with whom they will start to communicate. Otherwise, auctions are set up as usual.

Auction round $r \geq 1$**.** At the beginning of the r-th round, bidders receive a signal e_i^r about their true private value. This signal depends on their initial estimation e_i^1 and on the public information, obtained during the previous rounds. Bidders enter $b_i^r = e_i^r$ into their mobile device. After that, the devices participate in a cryptographically secured sealed-bid auction protocol between bidders, the seller and the authority. Every bidder i submits an encrypted bid, and argues in zero-knowledge that

$$\frac{1}{1-\varepsilon} b_i^1 \geq b_i^r \geq \max(b_i^{r-1}, X_2^{r-1} - 1) \ . \tag{1}$$

At the end of rth round, the authority outputs a signed tuple $(X_2^r; C_K(X_2^r))$. The authority accompanies this with a zero-knowledge argument that the values X_2^r and $C_K(X_1^r)$ are correctly computed. All this is published in an authenticated manner.

End criteria. The auction lasts $R \geq 2$ rounds and stops iff $X_2^R = X_2^{R-1}$ or $X_1^R = X_1^{R-1}$. The contract price will be X_2^R, unless $X_1^R = X_1^{R-1}$ and $R > 1$; in the latter case, the winning price will be X_2^{R-1}. Then Y_1^R is established by using another (interactive) cryptographic protocol. If there is a tie-break, one of the winners is selected by using the equal probability rule.

Cryptographic implementation. Next, we outline some cryptographic implementation details. Every round of our mechanism is a cryptographically secured Vickrey auction with some added bells and whistles. We base our example implementation on the LAN scheme [LAN02], although we stress that this is just an example cryptographic implementation. The main benefits of the LAN scheme, as compared to the competitors, are: (a) It does not rely on threshold trust between > 2 machines, possibly operated by the (occasional and thus untrusted) auctioneer himself, but rather on the assumption that the auctioneer and a trusted auction authority do not collaborate (see [NPS99,LAN02] for explanations why this model makes more sense than the threshold-trust-based auctioneering model); and (b) It is severely more efficient than other existing cryptographic auction schemes without the threshold trust.

As a consequence of using the LAN scheme, some of the information is moved around in an encrypted and some of it is moved around in a committed form; one needs to prove occasionally that the encrypted value is equal to the committed value. We will assume it implicitly in what follows. We will, as well, assume that the seller creates the commitment key K. Currently, the Damgård-Jurik homomorphic encryption scheme [DJ01] and the Damgård-Fujisaki integer commitment scheme [DF02] seem to be the best candidates for the cryptographic primitives E (encryption) and C (commitment). Since the Damgård-Fujisaki commitment scheme [DF02] is statistically hiding and computationally binding, the corresponding zero-knowledge arguments will be statistically hiding and computationally convincing. This suits the auction scenario perfectly well, since one might want to have bid privacy for a long time, while the binding (and convincing) property have a more online character.

In every round the bidders send their bids, in an encrypted (with A-s public key) form to the seller S, by using an authenticated channel. This is accompanied by a non-interactive statistical zero-knowledge (NISZK) argument that the bid was correctly formed [LAN02], and that Equation (1) holds. The latter can be done efficiently by using an efficient range argument [Bou00,Lip01]. Both the bids and the NISZK arguments are stored on a cryptographic bulletin-board [Rei94,Rei95].

Next, the seller forwards the encrypted bids to the authority, who decrypts the bids, finds out the two highest bids (X_1^r, X_2^r) and sends X_2^r and $C_K(X_1^r)$ (his commitment on X_1^r) back to the seller over an authenticated channel. X_1^r will not be revealed to the seller. This is accompanied with an NISZK argument that X_2^r was the second highest bid and that the committed bid was the highest bid (this corresponds precisely to a protocol from [LAN02]), and an NISZK range argument for either $X_1^r = X_1^{r-1}$ or $X_1^r > X_1^{r-1}$. After verifying the NISZK arguments, the seller posts $(X_2^r, C_K(X_1^r))$ together with the NISZK arguments and her own and authority's signatures on the bulletin-board. The bidders verify the signatures and the NISZK arguments.

The bulletin-board contents (that is, the tuple $(C_K(b_1^r), \ldots, C_K(b_B^r), X_2^r, C_K(X_1^r))$ together with the signatures and NISZK arguments) is stored.

Alternative cryptographic implementations. Alternatively, one can implement this mechanism by using Yao's model of general two-party computation [Yao82]. This would involve the design a specific circuit that is suitable for this mechanism. Such an approach was successfully used by Naor, Pinkas and Sumner [NPS99] for Vickrey auctions. The LAN auction scheme is more efficient (especially when the number of bidders is large), while the Naor-Pinkas-Sumner scheme will not reveal any unwarranted information to A. (Note that one can use any of the available cryptographic auction schemes that rely on threshold trust.)

3.2 Discussion

The meaning of ε. We call the bidders who are able to ε-approximate their true valuation *ε-rational*. Intuitively, one may assume that it is in common knowledge that non-ε-rational rational bidders will not participate. A value of ε, relevant in practical auctions, can be $0.1 \ldots 0.6$. Setting $\varepsilon \leftarrow 0$ would result in Vickrey auctions. A smaller ε will raise the time-efficiency of auctions and (as we will see) make the auctions less subject to jump bidding, while a greater ε has the potential to attract more bounded-rational bidders. If the seller wants to have a greater participation at the expense of risking to have longer auctions and jump bidding, she might set $\varepsilon \leftarrow 1 - 10^{-6}$.

Equilibria. Setting $b_i^r > e_i^r$ can occasionally result in negative payoffs. Thus, if the bidders are conservative then $e_i^r \leq v_i$. Therefore, choosing a value $b_i^r < e_i^r$ will not increase the payoffs. Hence the strategy of choosing $b_i^r = e_i^r$ is not dominated by any other strategy and therefore results in a non-dominated equilibrium.

Cognitive cost. Out mechanism becomes Pareto-efficient as soon as all bidders are able to calculate their valuations with an arbitrary large but *a priori* known accuracy, given that the bidders are rational enough to avoid some "weird" strategies. More precisely:

Lemma 1. *Our auction mechanism is Pareto-efficient if (a) The bidders are able to distinguish between the cases $v_i > e_i^r$ and $v_i = e_i^r$, (b) The ith bidder never bids more than v_i; (c) The bidders do not set $e_i^r \leq X_2^{r-1}$ if $v_i > X_2^{r-1}$; and (d) The highest bidder of a round does not overbid himself in the next round.*

Note that one can trade off the cognitive cost versus the privacy by publishing the tuple $(X_2^r, \ldots, X_m^r)$, $m > 2$, instead of X_2^r. This is an important property of our mechanism that makes it possible to have an almost continuous tradeoff between the cognitive costs and the privacy.

Computational efficiency. The two inequalities in Equation (1) are introduced, in particular, to increase the computational efficiency. The leftmost inequality enforces bidders to do at least some homework to estimate their valuation with precision ε. This can decrease the number of rounds. The rightmost inequality enforces the sequence (b_i^r) to be nondecreasing in r, and hence also helps to decrease the number of rounds. Bidding

$b_i^r = X_2^{r-1} - 1$ intuitively equals to passing: by doing so, one is guaranteed not to win at round r, since $X_1^r \geq X_2^r > X_2^{r-1} - 1$. Our chosen solution is superior to the one where the bidders can pass, since in this case some of the private information of bidders will become public. (Additionally, it would make it possible the bidders to collude by signaling each other.)

One can additionally decrease the number of expected rounds by requiring that if b_i^r increases, then $b_i^r > (1+\delta)b_i^{r-1}$ for some public constant δ that may depend on the currently second highest bid X_2^{r-1}. This solution is common in English auctions, and can also be employed in conjunction with our mechanism to achieve additional effectiveness. However, it also has the potential to decrease the revenues of the seller by a factor of $(1+\delta)$.

Expected revenue. Intuitively, in our mechanism, it is possible that the revenue X_2^R is smaller than the second highest valuation V_2, since $\Pr[X_2^R = V_2 - i] \sim 2^{-i-1}$. (As shown in [MW82], truthful revealing of information can never decrease the revenues of the seller, under the standard game-theoretic assumptions like omnipotency and rationality of the bidders.)

3.3 Security Analysis

By using a secure cryptographic implementation, the auction will be correct and privacy-preserving. Additionally, it will have some mechanism-centric properties that are not shared (say) by cryptographically secured English auctions.

We say that a bidder is *antisocial* if, may be knowing that he cannot win, he bids more than his value solely to increase the contract price of other players. That is, an antisocial bidder acts not to maximize his utility, but to minimize the utility of other players. A *shill* is an antisocial bidder that is manipulated by the seller to possibly drive up the price.

Lemma 2. *The proposed mechanism is secure against shills and antisocial bidders, as soon as all signatures and zero-knowledge arguments are verified and the highest bidder of round $r-1$ does not increase his bid in round r.*

Proof. In the round r, knowing the value X_2^{r-1}, a shill j will make some bid b_j^r. If $b_j^r \leq X_2^{r-1}$ then the second highest bid will not increase. Assume $b_j^r > X_2^{r-1}$. If $b_j^r > X_1^{r-1}$ then the shill j has to pay for the price himself. Bidding $X_2^{r-1} < b_j^r \leq X_1^{r-1}$ will raise the second highest bid but not the highest bid. Therefore, by the auction rules, the contract price will be X_2^{r-1}. In particular, if $b_j = X_1^{r-1}$ then the shill will neither increase the second highest bid nor become one of the tie-breakers. □

Security against premature finishing. A possible alternative to requiring everybody to decrease their bids over time is to instead have the same scheme without this requirement, but with declaring the winner of the previous round as the winner of the auction whenever $X_2^R < X_2^{R-1}$. However, then the highest bidder Y_1^{R-1} could in some cases prematurely finish the auctions (and thus decrease the revenues of the seller) by bidding X_2^{R-1} in round R. Given that only $Y_1^R = Y_2^{R-1}$ will bid $\geq X_2^{R-1}$ at round R, X_2^R will

be equal to X_2^{R-1}. If Y_2^{R-1} bid less than X_1^{R-1} in round R, Y_1^{R-1} will obtain the item for X_2^{R-1}, which might be less than the valuation of Y_2^{R-1}. Our mechanism does not have this problem.

Security against collusive bids. The proposed auction mechanism is secure against collusive bids by the same reasons why it is secure against shills' bids: the collusive bidders must bid more than the current highest bid to get their signal trough. However, this also means that they might have to pay for the item. This is at least the case when the previous round highest bidder had approximated her value sufficiently precisely.

Security against jump bidding. English auctions are subject to jump bidding, where one bidder bids very high in the beginning of the auction just to scare other bidders away. Our auction mechanism does not feature complete security against the jump bidding. Indeed, it might be the case that one participant (the one who values the item the most) approximates his bid relatively well during the first round, while other bidders will not do a worthy homework. In such a case, when $X_1^1 > v_2$, this one participant might obtain the item with as low price as $(1-\varepsilon)V_2$, where V_i is the actual valuation of the V_ith highest valuator.

More precisely, if the highest valuator i has the initial bid $b_i^1 \geq (1-\varepsilon)b_j^1$, for all $j \neq i$, he can get the item $(1-\varepsilon)$ times cheaper than in the case when somebody else would also be doing the homework. A similar phenomenon happens in the Vickrey auctions, except that there $\varepsilon = 0$. (English auctions are even worse in this sense than Vickrey auctions.) The larger is ε, the less can be gained by jump bidding. Thus we improve upon both Vickrey and English auctions by taking a moderately large ε. A cautious seller might have ε to be relatively high if she is afraid of jump bidding in the case when the richest client is also the most diligent. (Alternatively, she can just increase the initial price.) On the other hand, if rich but oblivious customers can be expected, a smaller ε will be more beneficial to the seller.

4 Comparison with Related Work and Conclusions

The first paper that emphasized the cognitive costs in online auctions is by Parker, Ungar and Foster [PUF98]. Their paper analyzed the existing mechanisms from this aspect and concluded that the English auctions are the best in the context of bounded rationality. A large body of research has been following, see [Par99,LS01] for some examples and further references. However, most of the papers in this line of research do not actually propose new mechanisms, but instead propose criteria on how to choose between the already existing and well-known mechanisms.

Moreover, the mentioned papers are more concerned about fully autonomous agents, and they assume that the agents can somehow quantify their computational costs of regulating their beliefs. This is often not the case.

A completely different line of research has been focusing on the correctness and privacy properties of the online auctions. Many different authors have been proposing completely different cryptographic schemes that guarantee correctness and privacy of many different auction mechanisms under different assumptions, including and excluding threshold trust. Again, the focus has been on the existing mechanisms.

Our approach was different. We first asked what is relevant in online auctions. Our conclusion was that correctness and privacy are more important than cognitive cost (since the latter cannot be precisely modeled), while the latter is more important than the computational effectiveness (e.g., the number of rounds). We proposed a new mechanism that has all the mentioned properties, but puts emphasis on the security over the cognitive convenience, and on the cognitive convenience over the computational convenience.

Moreover, our mechanism makes it possible to trade off cognitive costs versus computational costs (by changing the parameter ε), and cognitive costs versus privacy (by increasing the amount of published data $(X_2^R, \ldots, X_m^R)$).

Our mechanism can be used together with any reasonable cryptographic auction scheme. We described an implementation based on [LAN02], since we agree with its authors that avoiding threshold is more important than its bid statistics leakage to an established authority. Moreover, the scheme of [LAN02] is very efficient and easy to understand. However we stress that many other concrete cryptographic schemes can be used.

It has been long argued that security issues [RTK90,RH95] and huge cognitive cost [Par99] are two main reasons why incentive-compatible auction mechanisms like the Vickrey auction are not widely used in practice. Our scheme mitigates both concerns and is still nonmanipulable.

Further Work and Acknowledgments

We hope that this paper will stimulate more work in the direction of designing new auction mechanisms, suited for online auctions. We also expect to see some convergence between the until-now separate lines of research on the game-theoretic, cognitive and cryptographic properties of auctions. The trade-off between privacy, cognitive costs and computational efficiency is especially important in combinatorial auctions, and has received some attention in this context [HKMT02].

We would like to thank N. Asokan and Valtteri Niemi for fruitful discussions while writing the first version of this paper in 2001, under partial support of Nokia Research.

References

[Bou00] Fabrice Boudot. Efficient Proofs that a Committed Number Lies in an Interval. In Bart Preneel, editor, *Advances in Cryptology — EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 431–444, Bruges, Belgium, May 14–18 2000. Springer-Verlag. ISBN 3-540-67517-5.

[CoI02] National White Collar Crime Center and Federal Bureau of Investigation. Ifcc 2001 internet fraud report. Available at `http://www1.ifccfbi.gov/strategy/IFCC_2001_AnnualReport.pdf`, as of November 2002, 2002.

[DF02] Ivan Damgård and Eiichiro Fujisaki. An Integer Commitment Scheme Based on Groups with Hidden Order. In Yuliang Zheng, editor, *Advances on Cryptology — ASIACRYPT 2002*, volume 2501 of *Lecture Notes in Computer Science*, pages 125–142, Queenstown, New Zealand, December 1–5 2002. Springer-Verlag.

[DH76] Whitfield Diffie and Martin E. Hellman. New Directions in Cryptography. *IEEE Transactions Information Theory*, IT-22:644–654, November 1976.

[DJ01] Ivan Damgård and Mads Jurik. A Generalisation, a Simplification and Some Applications of Paillier's Probabilistic Public-Key System. In Kwangjo Kim, editor, *Public Key Cryptography '2001*, volume 1992 of *Lecture Notes in Computer Science*, pages 119–136, Cheju Island, Korea, 13–15 February 2001. Springer-Verlag.

[HKMT02] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling Equilibrium in Combinatorial Auctions. Working paper, `http://iew3.technion.ac.il/~moshet/rndm11.ps`, May 2002.

[Kri02] Vijay Krishna. *Auction Theory*. Academic Press, 2002.

[LAN02] Helger Lipmaa, N. Asokan, and Valtteri Niemi. Secure Vickrey Auctions without Threshold Trust. In Matt Blaze, editor, *Financial Cryptography — Sixth International Conference*, volume ? of *Lecture Notes in Computer Science*, pages ?–?, Southhampton Beach, Bermuda, March 11–14 2002. Springer-Verlag. To appear.

[Lip01] Helger Lipmaa. Statistical Zero-Knowledge Proofs from Diophantine Equations. Cryptology ePrint Archive, Report 2001/086, November 20 2001. http://eprint.iacr.org/.

[LS01] Kate Larson and Tuomas Sandholm. Costly Valuation Computation in Auctions. In Johan van Benthem, editor, *Eighth Conference of Theoretical Aspects of Knowledge and Rationality (TARK VIII)*, Certosa di Pontignano, University of Siena, Italy, July 8–10 2001. Morgan Kaufmann.

[Luc00] David Lucking-Reiley. Auctions on the Internet: What's Being Auctioned, and How? *Journal of Industrial Economics*, 48(3):227–252, September 2000.

[MW82] Paul R. Milgrom and Robert J. Weber. A Theory of Auctions and Competitive Bidding. *Econometrica*, 50(5):1089–1122, September 1982.

[NPS99] Moni Naor, Benny Pinkas, and Reuben Sumner. Privacy Preserving Auctions and Mechanism Design. In *The 1st ACM Conference on Electronic Commerce*, Denver, Colorado, November 1999.

[NS93] Hannu Nurmi and Arto Salomaa. Cryptographic Protocols for Vickrey Auctions. *Group Decision and Negotiation*, 2:363–373, 1993.

[Par99] David C. Parkes. Optimal Auction Design for Agents with Hard Valuation Problems. In Alexandros Moukas, Carles Sierra, and Fredrik Ygge, editors, *Agent Mediated Electronic Commerce II, Towards Next-Generation Agent-Based Electronic Commerce Systems, IJCAI 1999 Workshop*, volume 1788 of *Lecture Notes in Computer Science*, pages 206–219. Springer-Verlag, 1999.

[PUF98] David C. Parkes, Lyle H. Ungar, and Dean P. Foster. Accounting for Cognitive Costs in On-line Auction Design. In Pablo Noriega and Carles Sierra, editors, *Agent Mediated Electronic Commerce, First International Workshop on Agent Mediated Electronic Trading, AMET-98*, number 1571 in Lecture Notes in Computer Science, pages 25–40, Minneapolis, MN, USA, May10 1998. Springer-Verlag. Selected papers.

[PWZ00] Motty Perry, Elmar Wolfstetter, and Shmuel Zamir. A Sealed-Bid Auction that Matches the English Auction. *Games and Economic Behaviour*, 33(2):265–273, November 2000.

[Rei94] Michael K. Reiter. Secure Agreement Protocols: Reliable and Atomic Group Multicast in Rampart. In *2nd ACM Conference on Computer and Communications Security*, pages 68–80, Fairfax, Virginia, USA, 2–4 November 1994. ACM Press.

[Rei95] Michael K. Reiter. The Rampart Toolkit for Building High-integrity Services. In Kenneth P. Birman, Friedemann Mattern, and André Schiper, editors, *Theory and Practice in Distributed Systems*, volume 938 of *Lecture Notes in Computer Science*, pages 99–110, Dagstuhl Castle, Germany, 5–9 September 1995. Springer-Verlag, 1995. ISBN 3-540-60042-6.

[RH95] Michael H. Rothkopf and Ronald M. Harstad. Two Models of Bid-Taker Cheating in Vickrey Auctions. *Journal of Business*, 68(2):257–267, April 1995.

[RTK90] Michael H. Rothkopf, Thomas J. Teisberg, and Edward P. Kahn. Why are Vickrey Auctions Rare? *The Journal of Political Economy*, 98(1):94–109, February 1990.

[UPF98] Lyle H. Ungar, David C. Parkes, and Dean P. Foster. Cost and Trust Issues in On-Line Auctions. In *Agents'98 Workshop on Agent Mediated Electronic Trading (AMET'98)*, Minneapolis/St.Paul, MN, 1998.

[Vic61] William Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1):8–37, March 1961.

[Yao82] Andrew Chi-Chih Yao. Protocols for Secure Computations (Extended Abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 160–164, Chicago, Illinois, USA, 3–5 November 1982. IEEE Computer Society Press.