Interleaving Cryptography and Mechanism Design
The Case of Online Auctions

Edith Elkind[1] and Helger Lipmaa[2]

[1] Princeton University, Department of Computer Science,
35 Olden St, Princeton, NJ 08544, USA
{elkind}@cs.princeton.edu
[2] 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 for online auctions. This 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. We are aware of no previous work that interleaves cryptography explicitly with the mechanism design.
Keywords: auctions, cognitive costs, cryptography, mechanism design, privacy

1 Introduction

Traditionally, cryptography has been used to securify an existing auction mechanism—e.g., an English auction—by adding a layer of security and privacy on top of it. We show that introducing cryptography at the mechanism design level allows one to achieve many desirable properties. More precisely, we will concentrate on online auctions that can be organised over the Internet or a local wireless network. The bidders use software agents that do the computationally intensive parts of the bidding, while the human beings control the prices. Now, the software agents have, 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 (auctioneers) to choose between auction mechanisms that are infeasible to implement in conventional auctions. In particular, it becomes possible to use public-key cryptography [DH76] to ensure both security (correctness in the presence of malicious sellers) and bid privacy.

At the expense of mitigated computational costs, the importance of other mechanism properties will grow in online auctions. *Cognitive costs* of computing one's valuation will dominate over the computational costs. Therefore, to further simplify participation in online auctions, it is desirable to devise an auction mechanism that neither requires the bidders to do an elaborated precomputation to calculate their precise valuation, nor extensive online calculations to react properly to the bidding strategies of other participants.

Security is another important concern in auctions. Auction fraud was the most common complaint to Internet Fraud Complaint Centre (IFCC) during the last years [CoI03]. The number of frauds could be decreased by using an auction mechanism with better security properties. For example, an online auction mechanism should be secure against a malicious seller and various possible attacks (shills, collusive bids, jump bidding). Additionally, only a minimal amount of information should be leaked to the seller or to the other bidders. Unfortunately, not all goals are achievable at the same time. As we will see in Section 3, one must trade off cognitive costs and resource-effectiveness, as well as cognitive costs and privacy. In particular, to have small cognitive costs, one should allow a large number of rounds, but also introduce some (otherwise unnecessary) privacy leakage.

We argue that a good auction mechanism should emphasise privacy and security against the seller over cognitive costs. Hence, when constructing an online auction mechanism, one should first make sure that the auction satisfies the desired allocation criteria, is secure against sellers and (almost-ideally) privacy-preserving. The next goal is to mitigate the cognitive costs as much as possible, without hurting security against the seller and bid privacy. For example, to minimise the (online) cognitive costs, it is desirable to have a non-manipulable mechanism—otherwise, the strategies of participating bidders might become arbitrarily complex. On the other hand, also some information about other bidders' valuations must be leaked for this purpose. Finally, one should make sure that the mechanism is sufficiently effective—that is, that it does not have more (and desirably, has less) rounds with human interaction than say proxy bidding, another auction mechanism tailored for agent-mediated online auctions, or require super-polynomial-time computations.

We will propose a new auction mechanism that is based 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. We will discuss other desired and existing properties of (online) auctions in Section 3. There, we will point out why currently known mechanisms are less than ideal.

Briefly, every round of the new mechanism is a second-price auction (i.e., a Vickrey auction). This suffices to make the mechanism non-manipulable in the private value model, as well as in some interesting special cases of the common value model. Second, during every round only $m - 1$ bids are revealed, where m is a public auction parameter. The revelation helps 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, this auction mechanism is parameterised by the cognitive error coefficient $0 \leq \varepsilon < 1$, that forces the bidders to precompute their values at least to some extent and thus has the potential to reduce the number of rounds. Additionally, the described mechanism is cryptographically protected, and includes some sensible finishing conditions that provide protection against shills and collusive bids. Some protection is also provided against jump bids.

The proposed mechanism has the same privacy properties as the cryptographically secured Vickrey mechanism (indeed, the choice $m = 2$ and $\varepsilon = 0$ results in a Vickrey auction), while the cognitive costs are comparable to the ones in English auctions. See Section 4 for a fuller description of the new mechanism, followed by detailed analysis. Finally, the new mechanism seems to be the first one that has been designed from scratch to provide security against the seller and bid privacy, and to minimise cognitive costs at the same time.

This difference from the well-known methodology of adding a cryptographic protocol on top of an existing mechanism in that we are able to overcome some weaknesses of classical mechanisms. Therefore, our work has relevance to classical auction theory. We hope that it 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 and of mechanisms in general.

Road-map. Section 2 introduces some necessary cryptographic preliminaries. Section 3 gives a short overview of the different goals of auction mechanisms. Section 4 describes the new auction mechanism, followed by discussion and analysis. Section 5 explains the difference with related work.

2 Cryptographic Preliminaries

Public key cryptosystem is a triple $\Pi = (G_\Pi, E, D)$ of key generating, encryption and decryption algorithms. Commitment scheme $\Gamma = (G_\Gamma, C)$ is a tuple of key generating and commitment algorithms. We use standard notations like $E_K(m; r)$ and $C_K(m; r)$ to denote encryption/commitment of m by using a newly generated random value r. A public key cryptosystem Π (resp., a commitment scheme Γ) is *homomorphic* if $E_K(m_1; r_1)E_K(m_2; r_2) = E_K(m_1 + m_2; r_3)$ (resp., $C_K(m_1; r_1)C_K(m_2; r_2) = C_K(m_1 + m_2; r_3)$) for some r_3. For our purposes, we will use the homomorphic Damgård-Jurik cryptosystem [DJ01] that allows to flexibly encrypt large plaintexts. We will also use the homomorphic Damgård-Fujisaki (DF) statistically hiding and computationally binding integer commitment scheme [DF02] that allows to commit to arbitrary integers.

One can build efficient zero-knowledge arguments for a large class of languages by using an integer commitment scheme, as shown recently in [Lip03a]. In particular, there exist very efficient arguments for showing that (a) A committed number μ belongs to an arbitrary finite interval $[\ell, h]$. We call the corresponding argument a *range argument* and refer to [Lip03a] for precise proofs; and (b) A committed number has the form B^μ, where $\mu \in [\ell, h]$. We call the corresponding argument a *range argument in exponents* and refer to [LAN02,Lip03a] for a description. Due to the properties of the DF commitment scheme, the described zero-knowledge arguments will be statistically hiding and computationally convincing. This suits well the auction scenario, since one might want to have bid privacy for a long time, while the binding (and convincing) property is only needed for the duration of the auction.

We will also need to give range arguments (in exponents) for encrypted numbers. For this, we will assume that one accompanies all encryptions and operations on ciphertexts with similar commitments and operations on committed values. Now, when one needs to argue that the encrypted value satisfies some properties, one argues on the committed value instead, and then argues that the two values are equal. The latter argument is very standard.

3 Auction-Theoretic Goals for Mechanism Design

An auction mechanism is a protocol between the auction participants, with a motivational ingredient of monetary rewards for "proper" actions; in particular, it is required that nobody should have a negative payoff when following the auction mechanism. Some well-known mechanisms are English auctions (first-bid ascending auctions), Vickrey auctions (second-price sealed-bid auctions) and first-price sealed-bid auctions. We refer to [Kri02] for an overview of auction mechanisms. Auction theory usually assumes either the private value model (the bidders know their values or can compute them without using information about others' values) or the common value model (the valuation has a common component that is only partially known to bidders). 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*.

An ideal auction mechanism should aim for the following properties. First, with respect to allocation, usually the goal is either *Pareto-efficiency* or *revenue maximisation*. The former is equivalent to maximising the social welfare, i.e., awarding the item to the bidder who values it most, while the latter corresponds to maximising the seller's profit. Sometimes, these two goals are in conflict, in which case a trade-off between them can be considered. Second, *resource-effectiveness*: The auction takes a small number of human-interacted rounds. The auction rules are sufficiently simple so that the seller and the bidders can follow them in "reasonable" time. Third, *security against the (malicious) seller*: The seller cannot increase the final price or change the winner without being caught. Fourth, *privacy*: No information about the bids of honest bidders is revealed, except the information that can be derived from the winner's identity and the contract price. Fifth, *minimal cognitive cost*: The cognitive cost of computing the valuation is small. Other properties are security against shills, collusive bids, jump bidding, etc [Kri02].

The cognitive cost of strategy planning is especially important in online auctions [UPF98,PUF98]. Since other participation costs decrease considerably due to the use of software agents, cognitive costs of computing one's valuation start to dominate. Therefore, it becomes important to decrease cognitive costs by devising an auction mechanism that neither requires the bidders to do an elaborated homework to compute their precise valuation and strategy, nor requires them to do extensive online calculations to react properly to the bidding strategies of other bidders. Such an auction mechanism should still have other desirable properties.

One must trade off between some of the mentioned properties. Clearly, the more information is leaked during an auction, the smaller is the cognitive cost. In most cases, this results in a higher seller's revenue [MW82] and possibly more efficient allocation in the presence of bounded-rational bidders. Usually, this means that multi-round actions with gradual information leakage are therefore revenue-maximising and also guarantee the best results for bounded-rational bidders. An interesting alternative approach was presented in [PWZ00], who constructed a two-round second-price sealed-bid auction PWZ mechanism with the same seller revenues as the English auctions, but with the drawback (from the privacy standpoint) that the two highest bidders of the first round—who continue in the second round—obtained the the distribution of first round losers' bids. The PWZ mechanism is resource-effective, and also slightly better than the Vickrey mechanism in cognitive cost. However, if the bidders are bounded-rational, then the PWZ mechanism is not Pareto-efficient, since all but 2 players do not get a second chance to revise their bids, and the remaining 2 players only get one more chance for it.

Proxy bidding is a designated online auction mechanism that assumes that all bidders use software agents 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 does not rise anymore. Clearly, proxy bidding has smaller cognitive costs than one-shot auctions, and on the other hand, has smaller participation costs (due to the smaller number of human-interactive rounds) than English auctions. Hence, proxy bidding offers a balance between the cognitive cost and the resource-effectiveness of the English and Vickrey auction mechanisms. This may explain the dominance of proxy bidding 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 proxy bidding has its downsides. In particular, it does not solve the problem of revealed statistics even when cryptographically secured. (E.g., identities of persons who participate at every time moment, and therefore also partial information about their valuations, are revealed to the seller.) Moreover, if many bounded-rational people participate, proxy bidding can have a large number of rounds. 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 ideal.

Bid Privacy and Security against Sellers. Clearly, a malicious seller can change the results of an auction to his benefit when it is not possible to verify his actions or when he obtains too much information about bidders' valuations. This is commonly seen as a reason why Vickrey auctions are not employed in practice [RTK90,RH95]. This observation has 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 intentionally reveal the bid statistics to alleviate the cognitive cost.

We believe that a good auction mechanism should emphasise privacy and security against the seller over the cognitive costs. Our (informal) reasoning behind this belief is that it is easier to define what is the privacy (and what is a privacy leak) than to model the cognitive costs, as the latter vary widely from one bidder to another. For example, if instead of a single bid, information about two competing bids will be leaked, then this is certainly a privacy leak, but can bidders use this additional information to adjust their estimate of their own values? Probably yes, but how much exactly do they gain? If one cannot guarantee that a 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.)

Cryptographic auction schemes. *Cryptographic auction schemes* are cryptographic algorithms to support specific mechanisms, that, when correctly followed by an honest party, ensure that certain well-defined privacy/security-against-the-seller properties will

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

Table 1. Comparison of different existing auction mechanisms and the new 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. The first column refers to efficiency in the private value model.

be held w.r.t. her. In particular, a good auction scheme must ensure that neither a cheating seller nor cheating bidders can affect the allocation. Andrew Yao [Yao82] was the first to consider cryptographic (English) auctions. Cryptographic auction schemes for different auction mechanisms have been designed since then. (See [NPS99,LAN02] for some examples and an overview of the related literature.) In particular, cryptographic Vickrey auction schemes satisfy all desired properties that were described in the beginning of this chapter, except that they do not minimise the cognitive costs. The best cryptographic auction schemes guarantee security against the seller and privacy, to the extent required by the auction mechanism.

Summary of auction mechanisms. There are many well-known auction mechanisms, like 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 summarised in Table 1. (Note that we do not consider revenue maximisation: generally speaking, it is not achieved by any of the standard mechanisms, and also it requires more information about the bidders' valuations that we are willing to assume.) We do not know of any mechanism-scheme combination that satisfies all the previously described auction desiderata. Note that not all five desiderata, as described in the beginning of the current section, are equally important in all situations. Traditionally, one has mainly been stressing the first two properties. We will concentrate on online auctions, where, as we will see, the last three properties will gain in importance.

4 New Mechanism

4.1 High-Level Description

In this section, we describe the new cryptographically secured multi-round sealed-bid auction mechanism. Discussion and explanation will follow.

Notation. Let $P = \{v_1, \dots, v_V\}$ be the set of possible valuations, e.g., $\{0.01, 0.03, \dots, 0.90, 0.94, 1.00\}$; in practice, this means that some valuations are rounded off. The auction consists of the setup phase, rounds $1, \dots, R$ (where R is not fixed in advance), and the closing phase.

Setup phase. Assume that ϕ is a monotonic bijective function from P to the set of actual bids $[0, V-1]$ that is sent—in a signed form—to the bidders by the seller during the auction setup. (ϕ may be unknown by the auction authority A.) The mechanism is parameterised by public values $m \geq 2$ and $\varepsilon < 1$, selected by the seller S and announced to everybody before the auction. Intuitively, ε specifies to what degree the auction takes on the character of an English auction ($\varepsilon \to 1$) or a Vickrey auction ($\varepsilon = 0$), and m specifies the amount of deliberately leaked information. There are B bidders $1, \dots, B$, one seller S and the auction authority A. Anybody 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. The participants obtain a committing key, an encryption key and a signature key of the parties, with whom they will start to communicate. Otherwise, auctions are set up as usual, in particular by publicly announcing details such as the closing date.

Let $(X_1^r, \dots, X_B^r)$ be the list of bids made in the rth round, $r \geq 1$, in non-increasing order, and let Y_i^r be the bidder who made the bid X_i^r. Note that $X_i^r \in [0, V-1]$ for all r and i. We assume that $b_i = X_i^0 = 0$ and let (Y_i^0) be an arbitrary permutation of all bidders.

Auction round $r \geq 1$. Before the first round, all bidders receive a signal s_i about their true values. At the beginning of the rth round, $r \geq 1$, the bidders compute an estimate $e_i^r \in P$ of their true private values that depends on their initial signal and on the public information obtained in the previous rounds. Intuitively, for rational agents it should be the case that $(1-\varepsilon)v_i \leq e_i^r \leq v_i$. Bidders enter $b_i^r = \phi(\beta_i(e_i^r))$ into their software agent, where β is the ith strategy function. After that, the agents participate in a cryptographically secured sealed-bid auction protocol between bidders, the seller and the authority. Every bidder i submits an encryption of b_i^r, and argues in zero-knowledge that

$$\phi\left(\frac{1}{1-\varepsilon}\phi^{-1}(b_i^1)\right) \geq b_i^r \geq b_i^{r-1} \quad . \tag{1}$$

At the end of rth round, the authority outputs a signed tuple $\mathsf{view}(r) := (X_2^r, \dots, X_m^r; C_K(X_1^r; \rho^r))$ for a new random value ρ^r. The authority accompanies this with a non-interactive zero-knowledge argument that $\mathsf{view}(r)$ is correctly computed. All this is published in an authenticated manner to all bidders, who can do independent verifications.

Closing phase. The auction lasts $R \geq 2$ rounds and stops iff $X_2^R = X_2^{R-1}$. (This is verified by all bidders by using the published zero-knowledge arguments.) The contract price will be X_2^R. Then Y_1^R is established by using another (interactive) cryptographic protocol. If there is a tie, one of the winners is selected by using, e.g., the equal probability rule.

4.2 Cryptographic implementation

Every round of the new mechanism is a cryptographically secured 2nd-price auction where instead of only the second highest bid, $m-1$ bids are revealed. Next, we outline some cryptographic implementation details. We will base our implementation on the LAN mth-price auction scheme [LAN02], although we stress that this is just an example cryptographic implementation. Additional tools [LAN02] can be employed to make the implementation secure against replaying attacks.

To simplify the zero-knowledge arguments, we will assume that bid 0 corresponds to some absolute minimal price $p_{\min}$, and that $\phi(x) = d \cdot \log_{1/(1-\varepsilon)} \frac{x}{p_{\min}}$ for some fixed d. (This same assumption makes also sense from the psychological and auction-theoretic viewpoints, see Section 4.3.) In this case, $\phi^{-1}(b) = p_{\min}(\frac{1}{1-\varepsilon})^{b/d}$, and the left side of (1) simplifies to the requirement that $(\frac{1}{1-\varepsilon})^{1+(b_i^1/d)} \geq (\frac{1}{1-\varepsilon})^{b_i^r/d}$, or $b_i^r \leq d+b_i^1$.

As in the LAN scheme, we will accompany all encryptions with corresponding commitments. Assume K is A's public key. In every round r, the ith bidder sends an encryption of $B^{b_i^r}$ to the seller S, by using an authenticated channel. This is accompanied by an efficient non-interactive statistical zero-knowledge (NISZK) argument that the bid was correctly formed [LAN02] (this is a range argument in exponents), and that (1) holds (this consists of two range arguments). These arguments can be shortened by using a different encoding function $Z_B(b_i)$ instead of B^{b_i} [Lip03a]. Both the bids and the NISZK arguments are stored on a cryptographic bulletin-board that is made publicly available to all bidders. (They can also simply be sent to all bidders.)

Next, the seller forwards the product of encrypted bids to the authority, who decrypts the bids, finds out the m highest bids and sends $\mathsf{view}(r)$ back to the seller over an authenticated channel; note that X_1^r is not revealed to the seller. This is accompanied with an NISZK argument that $C_K(X_1^r; \rho^r)$ commits to the highest bid, and that $(X_2^r, \ldots, X_m^r)$ are the next $m-1$ highest bids (this can be done as a straightforward extension of the protocol from [LAN02] for proving that $\tilde{X}$is the mth highest bid), and an NISZK range argument for either $X_2^r = X_2^{r-1}$ or $X_2^r > X_2^{r-1}$. After verifying the NISZK arguments, the seller posts $\mathsf{view}(r)$ 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 $(C_K(b_1^r), \ldots, C_K(b_B^r), \mathsf{view}(r))$ together with the signatures and NISZK arguments) is stored by all bidders.

In the closing phase, all bidders verify the correctness of closing and that the winning price was determined correctly (another range argument). Y_1^R can be established by using a method proposed in [Lip03b]: namely, all bidders and the seller participate in a proxy verifiable private equality test, after what the seller gets to know which bidder bid X_1^R without getting to know the value of X_1^R.

Alternative cryptographic implementations. Alternatively, one can implement the described auction mechanism by using Yao's model of general two-party computation [Yao82]. This would involve the design of a specific circuit that is suitable for the described mechanism, as successfully done by Naor, Pinkas and Sumner [NPS99] for mth-price auctions, although in the case of the new mechanism, the circuit will be considerably larger. It is also not immediately clear how to extend the Naor-Pinkas-

Sumner scheme efficiently to a multi-round scheme, where the number of rounds is not bounded. The LAN auction scheme is more communication-efficient (especially when the number of bidders is large), while the Naor-Pinkas-Sumner scheme, as corrected by [JS02], will not reveal any unwarranted information to A. Also, one can use any of the available mth-price cryptographic auction schemes that rely on threshold trust, although not all of them might be flexible enough to be used with the new mechanism. Also, we share the viewpoint of [NPS99,LAN02] that threshold trust between > 2 machines, possibly operated by the (occasional and thus untrusted) seller himself is not sensible in most of the auction scenarios.

4.3 Discussion

The rôle of ε. We call the bidders who are able to ε-approximate their true valuation *ε-rational*. Intuitively, one may assume that it is common knowledge that non-ε-rational rational bidders will not participate. A value of ε relevant in practice 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. However, if the seller wants to have a greater participation at the expense of risking to have longer auctions and jump bidding, she might even set $\varepsilon \leftarrow 0.999$.

The function ϕ. As we already saw in Section 4.2, a suitable function ϕ can simplify the cryptographic implementation. The specific choice of ϕ proposed in Section 4.2 makes sense from both psychological and auction-theoretic viewpoint. Really, people are often thinking about the object's value on the logarithmic scale ("the first item is worth 3 times more than the second item") rather than on the linear scale. One should note, however, that this choice of function ϕ requires the seller to set a lower bound $p_{\min} = \phi^{-1}(0)$ and an upper bound $p_{\max} = \phi^{-1}(V-1)$ on the selling price, although the difference between these two values can be made almost arbitrarily large, since $\phi^{-1}(V-1)/\phi^{-1}(0) = (\frac{1}{1-\varepsilon})^{(V-1)/d}$. Assuming, say, that $\varepsilon = 0.95$, $V = 201$ and $d = 100$, this would make the price increase by 3% when bid is increased by 1, and we would have $p_{\max} \approx 400 p_{\min}$. This setting seems to be perfect for most auctions.

Equilibria. Setting $b_i^r > \phi(e_i^r)$ can occasionally result in negative payoffs. If the bidders are conservative then $\beta(e_i^r) \leq v_i$, so $b_i^r \leq \phi(v_i)$. Moreover, in many practically relevant cases the strategy of bidding strictly less than $\phi(e_i^r)$ is weakly dominated, so truth-telling results in a non-dominated equilibrium. For instance, if the bidders' values are private, i.e., the current price does not affect a bidder's estimate of the value, the usual argument for Vickrey auctions can be used to show that bidding $\phi(e_i^r)$ is a dominant strategy.

Moreover, truth-telling can be dominant in certain special cases of the common value model as well. In particular, we can show that this is the case for the "experts vs. amateurs" model. In this model, the valuation of the bidder i is of the form $v_i = w_i + T z_i$, where v_i, z_i are independently but not necessarily identically distributed random variables, and T is a random variable (same for all bidders) that can be equal to 0 or 1. Some bidders (let us call them experts) know the actual value of T, while others do not.

This model captures the markets in which some users (e.g., art dealers in an art auction) can determine whether the object being sold has some desired properties (e.g., whether a coin is fake or authentic), while others do not have this ability. In these markets, proxy bidding with a fixed deadline is susceptible to "sniping", i.e., experts bidding in the very last minute to prevent others from observing change in the posted price and adjusting their values (and hence their bids). This may result in inefficient allocation, and thus it is desirable to have a mechanism that does not encourage sniping.

Note that, intuitively, when $T = 1$, the experts may want to shade their bid to conceal this fact: it might be the case that when non-experts bid just w_i, the expert gets the object even though his own value is not particularly high, while if the others knew that $T = 1$, they would outbid him (in some sense, this is similar to sniping). Fortunately, we can show that because of our choice of finishing criteria our scheme does not have this problem, and, in fact, always achieves efficient allocation assuming conservative bidders.

Theorem 1. *Assume that all bidders are conservative, i.e., they avoid strategies that may lead to negative payoffs. Then conservative truthful bidding (experts bid $w_i + Tz_i$, others bid w_i if they cannot determine T from the outcomes of the previous rounds, and $w_i + Tz_i$ otherwise) is a Nash equilibrium, that is, no single bidder can gain by cheating.*

Proof. Consider the behaviour of bidder 1 in the first round assuming that everyone else bids truthfully. If bidder 1 is an amateur, or $T = 0$, the usual argument for Vickrey auctions applies. Now, suppose that $T = 1$. If bidder 1's truthful bid would not be the highest bid (assuming everyone else bids truthfully). Then bidder 1 cannot win the auction at a price that is lower than his value, so he might as well bid truthfully and lose. Hence, let us assume that bidder 1's truthful bid is higher that all other bids. Let $b = \max\{b_2^1, \ldots, b_n^1\}$. Suppose that bidder 1 decides to shade his bid. If he bids more than b (but less than his true value), the public information will be the same as in the case of truthful bidding, so this will not help. Alternatively, he can bid less than b, which means that he does not win the current round. Then, it might still be possible for everyone to derive that $T = 1$ (for instance, there may be several other experts who bid truthfully), so in the next round everyone will bid $w_i + Tz_i$, and the setting is that of ordinary Vickrey auction. Finally, it might be the case that when bidder 1 cheats, others cannot be sure that $T = 1$. Then they will not change their bids, and unless bidder 1 bids more than b, the auction ends and he loses. To avoid that, he himself has to bid more than b in the second round, so we are back to square one. □

Cognitive cost. Our mechanism becomes Pareto-efficient as soon as all bidders are able to calculate their valuations with an arbitrary high but *a priori* known accuracy, given that the bidders are sufficiently rational to avoid a limited number of well-specified "bad" strategies. More precisely, one can easily prove the next theorem:

Theorem 2. *Assume that the underlying cryptographic implementation is secure. The described auction mechanism is Pareto-efficient with overwhelming probability if (a) The highest valuator is honest and in particular double-checks all zero-knowledge arguments and signatures, (b) The ith bidder never bids more than $\phi(v_i)$; and (c) The* *highest valuator does not set $b_i^r \leq X_2^{r-1}$ if $v_i > \phi^{-1}(X_2^{r-1})$,*[1] *(d) The highest valuator is ε-rational.*

Proof. Assume that bidder 1 had the highest valuation, and assume that (a) holds. Then it is known that after the closing $X_2^R = X_2^{R-1}$, and by (c), no bidder but 1 has a valuation higher than $\phi^{-1}(X_2^{R-1})$. By (a), the highest valuator still participates in the round R, and by (d), he is allowed to place a high enough bid. Finally, by (b), $v_{Y_1^R} > \phi^{-1}(X_2^{R-1})$ and thus Y_1^R is the highest valuator. □

The virtue of this result is to make it precise when exactly the highest valuator will not obtain the item. In particular, it happens if his behaviour is in some sense quite irrational, the cryptographic implementation is insecure or other bidders are not conservative. In cryptography, it is important to give the security proofs under minimal assumptions, and our approach is the same. Moreover, all our assumptions are feasible.

Importantly, one can trade off cognitive cost and privacy by publishing the tuple $(X_2^r, \dots, X_m^r)$, $m > 2$ instead of just X_2^r. Moreover, the mechanism can be generalised to reveal some other function of the bid vector, e.g., the number of bids exceeding a given threshold, the number of bidders who increased their bids compared to the previous round, etc., provided that this function can be efficiently computed in a secure manner. Depending on the structure of bidder's valuations, this can decrease the cognitive costs significantly, while having a negligible effect on privacy. This allows for an almost continuous tradeoff between the cognitive costs and the privacy. However, whenever a privacy leak can be quantified much more easily than the possible win in cognitive costs (and this usually the case), we would recommend to use the value $m = 2$.

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 $b_i^r = b_i^{r-1}$ intuitively equals to passing: by doing so, one is guaranteed not to win at round r, unless his bid in round $r - 1$ was the highest one. The chosen solution is superior to the one where the bidders can pass if their bids are not high enough, 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 signalling each other.)

One can additionally decrease the number of expected rounds by requiring that if b_i^r increases, then $\phi^{-1}(b_i^r) > (1 + \delta)\phi^{-1}(b_i^{r-1})$ for some public value 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 the described mechanism to achieve additional effectiveness. However, since we assume that the bidders are conservative, it also has the potential to decrease the revenues of the seller by a factor of $(1 + \delta)$.

[1] We can make this assumption weaker, by assuming that he does not set $b_i^r \leq X_2^{r-1}$ if $v_i > (1+\delta)\phi^{-1}(X_2^{r-1})$ for some δ. Then the scheme will be δ-efficient, i.e., the value of the bidder who gets the object is within a factor of $1/(1 + \delta)$ from the highest value.

4.4 Security Analysis

When a secure cryptographic implementation is used, 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, maybe knowing that he cannot win, he bids more than his value solely to increase the contract price of other players [FB01]. That is, an antisocial bidder acts not to maximise his utility, but to minimise the utility of other players. We assume that antisocial bidders are conservative: that is, they will not bid more than the maximum of X_2^{r-1} and their own valuation. (They do not risk to come out with a negative payoff.) A *shill* is an antisocial bidder that is manipulated by the seller so as to drive up the price.

Theorem 3. *Suppose that the bidders' signals are sufficiently independent, namely, that from observing his own signal and the public information* $(X_2^{r-1}, \ldots, X_m^{r-1})$*, a bidder* j *cannot conclude with certainty that another bidder* i *has a value* v_i *such that* $\phi(v_i) > X_2^{r-1} + \delta$ *for a fixed value of* δ*. Then the proposed auction mechanism is secure against shills and antisocial bidders, as soon as all signatures and zero-knowledge arguments are verified.*

Proof. In the round r, knowing the value X_2^{r-1}, a shill j will make some bid $b_j^r > \phi(e_j^r)$. If $b_j^r \leq X_2^{r-1}$ then the second highest bid will not increase. Assume $b_j^r > X_2^{r-1}$. According to our assumption, j cannot be sure that his bid is lower than the highest bid, or that in the next round someone will be willing to bid more than b_r^j. So, there is a chance that he will have to pay the price himself, and, being conservative, he will refrain from submitting a bid that is higher than his value. □

Security against collusive bids. For $m = 2$, the proposed auction mechanism is secure against collusive bids by the same reasons why it is secure against shills' bids: namely, the collusive bidders must bid more than the current highest bid to get their signal through. 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 previous argumentation that in the first-price auctions the bidders are not motivated to jump-bid does not clearly apply always—for example when Y_1 knows the approximate value of X_2.

While the described auction mechanism does not feature complete security against the jump bidding, it provides an approximate protection. First, being a second-price auction, it is secure against the case when one bidder jump bids, since only a (relatively moderate) X_2 would be published and other bidders would still have a chance to over-bid it. Now, assume that at least two bidders jump bid, say bid within a fraction $1-\delta \gg 1-\varepsilon$ of their real valuations. In this case, $\phi^{-1}(X_2^1) \geq (1-\delta)V_2$, and the minimum price Y_1^1 has to pay is $\phi^{-1}(X_2^1) \geq (1-\delta)V_2$ instead of V_2. This "worst" case would only happen in the case $\phi^{-1}(X_1^1) > V_2$, assuming that Y_2^1 would over-bid X_1^1 otherwise and that Y_1^1 and Y_2^1 do not collude.

Thus, in this case the highest valuator can get the item $1/(1-\delta) \ll 1/(1-\varepsilon)$ times cheaper than in the case when somebody else would also be doing the homework. The smaller is ε, the less can be gained by jump bidding. A cautious seller might have ε to be relatively low 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 are to be expected, a larger ε will be more beneficial to the seller.

Adding another finishing criterion to the described auction mechanism makes it secure against nonconservative shills but insecure against jump bidding. Namely, if we say that the auction is finished if either $X_1^R = X_1^{R-1}$ or $X_2^R = X_2^{R-1}$, then a shill has an effect on the auction only when he bids more than X_1^{R-1} (and thus wins the auction). On the other hand, in this case a jump-bidder would be guaranteed to win the auction with his bid X_1^{R-1} unless a higher valuator will bid more than X_1^{R-1} (without knowing this value!) during the next round.

Security against premature finishing. A possible alternative to requiring everybody not to decrease their bids over time is to instead have the same scheme where this requirement is replaced by declaring Y_1^{R-1} 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. In the case when 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 $\phi^{-1}(X_2^{R-1})$, which might be less than the valuation of Y_2^{R-1}. The mechanism devised in this paper does not have this problem.

5 Comparison with Related Work and Conclusions

To our knowledge, the first paper to emphasise the cognitive costs in online auctions is by Parker, Ungar, and Foster [PUF98]. Their paper analysed the existing mechanisms from this perspective and concluded that English auctions are the best in the context of bounded rationality. A large body of research has followed. However, it mostly consisted of papers that did not actually propose new mechanisms, but instead suggested criteria for choosing between already existing and well-known mechanisms. Moreover, the emphasis of the above-mentioned papers is on fully autonomous agents, and it is assumed 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 security against sellers and privacy properties of the online auctions. Various authors have been proposing a wide range of cryptographic schemes that guarantee security against sellers and privacy of various auction mechanisms under various assumptions, including and excluding threshold trust. Again, the focus has been on the existing mechanisms.

Our approach is different. We first asked what is relevant in online auctions. Our conclusion was that security against sellers and privacy are more important than cognitive costs (since those are hard to model precisely), 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 security over cognitive convenience, and on cognitive convenience over computational convenience. Moreover, the described 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 $(X_2^R, \ldots, X_m^R)$). It has long been argued that security issues and huge cognitive cost are two main reasons why non-manipulable auction mechanisms like the Vickrey auction are not widely used in practice. The scheme described in this paper mitigates both concerns and is non-manipulable whenever Vickrey auctions are.

The described 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 trust is more important than its bid statistics leakage to an established authority. Moreover, the scheme of [LAN02] is very efficient and easy to understand. A full description of, say, the Naor-Pinkas-Sumner [NPS99] auction scheme would have made the paper less modular. However, several other cryptographic schemes can be used here.

Finally, one can simplify the proposed mechanism-protocol interleaving in a straightforward way to obtain a secure proxy bidding protocol. To our knowledge, no cryptographic protocol to securify proxy bidding has been proposed before.

Acknowledgements. The second author would like to thank N. Asokan and Valtteri Niemi for fruitful discussions while writing the first version of this paper in 2001. A part of this work was done while the first author was visiting Helsinki University of Technology. This work was partially supported by Nokia research and the Finnish Defence Forces Research Institute of Technology.

References

[Bla02] Matt Blaze, editor. *Financial Cryptography — Sixth International Conference*, volume 2357 of *Lecture Notes in Computer Science*, Southhampton Beach, Bermuda, March 11–14 2002. Springer-Verlag.

[CoI03] National White Collar Crime Center and Federal Bureau of Investigation. IFCC 2002 Internet Fraud Report. Available at `http://www1.ifccfbi.gov/strategy/IFCC_2002_IFCCReport.pdf`, as of April 2003, 2003.

[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.

[FB01] Gerhard Weiß Felix Brandt. Antisocial Agents and Vickrey Auctions. In *Intelligent Agents VIII*, pages 335–347, Seattle, WA, USA, August 1–3 2001. Revised papers.

[JS02] Ari Juels and Michael Szydlo. A Two-Server, Sealed-Bid Auction Protocol. In Blaze [Bla02], pages 72–86.

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

[Lai03] Chi Sung Laih, editor. *Advances on Cryptology — ASIACRYPT 2003*, volume 2894 of *Lecture Notes in Computer Science*, Taipei, Taiwan, November 30–December 4 2003. Springer-Verlag.

[LAN02] Helger Lipmaa, N. Asokan, and Valtteri Niemi. Secure Vickrey Auctions without Threshold Trust. In Blaze [Bla02], pages 87–101.

[Lip03a] Helger Lipmaa. On Diophantine Complexity and Statistical Zero-Knowledge Arguments. In Laih [Lai03], pages 398–415.

[Lip03b] Helger Lipmaa. Verifiable Homomorphic Oblivious Transfer and Private Equality Test. In Laih [Lai03], pages 416–433.

[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.

[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.