

(Im)Possibility of Unconditionally Privacy-Preserving Auctions

Felix Brandt
Computer Science Department
Carnegie Mellon University
Pittsburgh PA 15213
brandtf@cs.cmu.edu

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

Abstract

We investigate how to obtain bid privacy in sealed-bid auctions. In particular, this paper focuses on unconditional full privacy, i.e., privacy that relies neither on trusted third parties (like auctioneers) or trusted fractions of bidders, nor on computational intractability assumptions (like the hardness of factoring). These constraints imply a scenario in which bidders exchange messages according to some predefined protocol in order to jointly determine the auction outcome without revealing any additional information. It turns out that the first-price sealed-bid auction can be emulated by an unconditionally fully private protocol. However, the protocol's round complexity is exponential in the number of bits that represent a bid, and we show there is no more efficient protocol. On the other hand, we prove the impossibility of fully privately emulating the second-price sealed-bid (Vickrey) auction for more than two bidders. This impossibility holds even when relaxing various privacy constraints such as protecting just a single losing bid (while maintaining anonymity) or tolerating the revelation of complete information to a coalition of at least half of the bidders.

1. Introduction

Auctions are key mechanisms for allocating goods, services, tasks, and resources among multiple agents (e.g., [26, 24, 20, 8]). At the same time, privacy is a crucial issue in multiagent systems. A major reason why people may be hesitant to use software agents, or to participate in Internet commerce themselves, is the worry that too much of their private information is revealed. Furthermore, in the modern electronic society, the information might get propagated to large numbers of parties, stored in permanent databases, and automatically used in undesirable ways. In this paper, we will study the possibility of executing the most common sealed-bid auction protocols in a way that preserves the bidders' privacy to a maximal extent.

Sealed-bid auctions are not only widely used for the selling of goods, they also have been shown to be applicable to task assignment, scheduling, and finding the shortest path in a network with selfish nodes. Bid privacy is of increasing importance in such auctions, and various schemes that avoid blind trust in a single auctioneer have been proposed recently. In contrast to existing work, this paper deals with *unconditional full privacy*, i.e., privacy that relies neither on trusted third parties (like auctioneers) or trusted fractions of bidders, nor on computational intractability assumptions (like the hardness of factoring). We derive several impossibility and possibility results in this domain.

Our setting consists of one seller and n bidders that intend to come to an agreement on the selling of a good.¹ The two major (sealed-bid) mechanisms that yield such an agreement are the *first-price* and *second-price* (Vickrey) [25] auctions. In both mechanisms, each bidder submits a sealed bid to a trusted-third party called the auctioneer² expressing how much he is willing to pay. The auctioneer declares the bidder who submitted the highest bid as the winner of the auction. In the first-price auction, the winning bidder pays the amount that he bid, whereas in the second-price auction, he has to pay the amount of the second-highest bid. Both auction formats have their strengths and weaknesses. For example, the first-price auction yields more revenue when bidders are risk-averse. The second-price auction, on the other hand, is strategy-proof, which means that bidders are best off bidding their true valuation of the good to be sold (when valuations are independent). This eliminates an agent's need to counterspeculate on the other agents' valuations. Interestingly, the side-effects of this striking advantage are said to contribute to the sparseness of the second-price auction in the real world for two reasons [22, 21, 23]: Bidders are reluctant to reveal their true valuations to the auctioneer, and bidders doubt the correctness of the result as

- ¹ All the presented results also hold for similar auctions for other areas of application, in particular procurement auctions where there is one buyer and multiple sellers.
- ² Sometimes the auctioneer and the seller are the same person.

they do not pay what they bid (unlike in the first-price auction). For example, the auctioneer might create a second-highest bid slightly below the highest bid in order to increase his revenue. Both issues are based on a lack of trust in the auctioneer. For this reason, it would be desirable to somehow “force” the auctioneer to always select the right outcome (*correctness*) and “prohibit” the propagation of private bid information (*privacy*). Various schemes for satisfying these desiderata (for first-price as well as second-price auctions) have been proposed in recent years (e.g., [18, 15, 1, 7, 6]).³ Virtually all of them rely on at least some of the following three assumptions.

1. A certain fraction of third parties (auctioneers) is trustworthy.
2. The adversary, *i.e.*, parties that intend to violate correctness and privacy, is limited to polynomially-bounded computational power.
3. One-way functions exist.

Regarding assumption (1): privacy is usually obtained by distributing the trust onto several auctioneers and using various forms of secure multiparty computation (MPC). However, a coalition of *all* auctioneers can always breach privacy. For this reason, we distribute the computation of the auction outcome on bidders themselves. We say that an auction is *fully private* [7] if it is distributed on bidders and privacy can only be breached by a coalition of *all* bidders.⁴

Assumptions (2) and (3) are based on computational intractability. When relying on intractability assumptions (e.g., the hardness of factoring), it has been shown that MPC allows the computation of arbitrary functions so that no private input can be uncovered by a polynomially-bounded adversary [14]. Unfortunately, assumption (3) not only relies on the unproven assumption $\mathcal{P} \neq \mathcal{NP}$ but also on the widely unknown field of average-case complexity and further, more specific assumptions. Moreover, even when these conjectures are true, it may be possible to breach privacy in the future when sufficient computational power becomes available⁵, violating assumption (2). The results in this paper do *not* rely on intractability. This is called *unconditional privacy* (aka. *non-cryptographic* or *information-theoretic privacy*) as the adversary’s computational power is unlimited. It is known that only a restricted class of functions can be fully privately computed in this model.⁶ Section 2

presents some known results about this class of functions. As is standard in unconditional MPC [4, 9], we assume a complete network of private channels between agents. However, sometimes, a given protocol can also be implemented by just providing a broadcast channel (see Theorem 3).

In order to simplify the presentation, we will focus on what are known as *passive* (aka. *honest-but-curious*) adversaries in the cryptographic literature, *i.e.*, we assume that participants follow the prescribed protocol. Active adversaries, on the other, hand may arbitrarily deviate from the protocol by sending manipulated messages. This assumption does not restrict the applicability of our results because there are standard cryptographic techniques (zero-knowledge arguments in our case) that force active adversaries to act according to a protocol (see *e.g.*, [13]). However, using these techniques will incur overhead. After all, *negative* results in the passive adversary model also hold in a model that allows active adversaries.

In a nutshell, this paper investigates the availability of distributed protocols that allow n bidders to jointly determine the outcome of first-price or second-price auctions by exchanging messages according to some predefined rules and without revealing unnecessary information. In the rest of this paper, this is called *emulation* of an auction.

The results on second-price auctions also hold for a generalization called uniform-price or $(M + 1)$ st-price auction. In an $(M + 1)$ st-price auction, the seller offers M identical items and each bidder desires to buy *one* of them. It has been proven that it is a strategy-proof mechanism to sell those items to the M highest bidders for the uniform price given by the $(M + 1)$ st highest bid [25]. The Vickrey auction is just a special case of this mechanism for the selling of single goods ($M = 1$).

In the case of ties, we deliberately leave the outcome undefined. As a consequence, the impossibility results of this paper hold regardless of what is done in case of a tie: picking the auction winner at random, using bidder priorities, or even revealing the identities of tied bidders.

The remainder of this paper is structured as follows. Section 2 presents some known theoretic results that we will leverage in our proofs. In Sections 3 and 4, we study the existence of fully private protocols that emulate first-price and second-price auctions, respectively. In both sections, we consider public outcome functions in which all bidders learn the auction outcome as well as private outcome functions in which only the winning bidder learns the outcome. In Section 5, we propose several relaxations of our strict privacy model and investigate the possibility of fully private auction protocols under these loosened restrictions. The paper concludes with an overview of obtained results and a brief outlook in Section 6.

³ [27] even deals with combinatorial auctions.

⁴ In cryptographic terms, this is $(n - 1)$ -privacy.

⁵ This does not require super-polynomial computational power. The security parameter used for a protocol might be too low considering future computational power. *E.g.*, 512-bit RSA keys were considered secure a while ago but are not secure anymore.

⁶ When assuming that a majority of the agents is trustworthy (recall that this is not *full* privacy), *all* functions can be jointly computed in the unconditional model [4, 9].

2. Preliminaries

In this section we review some key results which we will use as building blocks in our proofs. We say that function f is *privately computable* if it can be jointly computed by agents using a distributed, randomized protocol consisting of several rounds. In each round, each agent may send a message to any other agent. Each message an agent sends is a function of his input (*i.e.*, his bid), his independent random input, the messages he received so far, and the recipient. When the protocol is finished, all agents know the value of f . No subset of agents is capable of uncovering any information besides what can be inferred from the function outcome and the coalition's inputs.

A complete characterization of all privately computable Boolean functions has been given:

Theorem 1 [11] *A Boolean function is privately computable if and only if it is of the form $f(x_1, x_2, \dots, x_n) = B_1(x_1) \oplus B_2(x_2) \oplus \dots \oplus B_n(x_n)$, where $B_i(x_i)$ are Boolean predicates and \oplus is the Boolean exclusive-or operator.*

Such a complete characterization for general (non-Boolean) functions is not yet known (except for only two agents [16]). However, there are necessary conditions for the private computability of a function.

Lemma 1 (Corners Lemma) [16]⁷ *Let $f : X \times Y \rightarrow Z$ be a privately computable 2-ary function. For every $x_1, x_2 \in X$ and $y_1, y_2 \in Y$, if $f(x_1, y_1) = f(x_1, y_2) = f(x_2, y_1) = a$, then $f(x_2, y_2) = a$.*

Lemma 2 (Partition Lemma) [11] *Let $f : X_1 \times X_2 \times \dots \times X_n \rightarrow Z$ be a privately computable n -ary function. Then, for each $i \in \{1, 2, \dots, n\}$ the 2-ary function $f_2(x_i, (x_1, x_2, \dots, x_{i-1}, x_{i+1}, x_{i+2}, \dots, x_n)) \stackrel{\text{def}}{=} f(x_1, x_2, \dots, x_n)$ is privately computable.⁸*

By combining Lemma 1 and Lemma 2, we can obtain a necessary condition for the possibility of privately computing an n -ary function. This can be used to prove that the outcome of an auction with n bidders is *not* privately computable (as in Theorems 2, 4, 5, and 6).

Lemma 3 *Let \vec{x} and \vec{y} be vectors of $n - 1$ bids and x and y single bids. It is impossible to fully privately emulate an auction if (\vec{x}, x) , (\vec{x}, y) , and (\vec{y}, x) all yield outcome a and (\vec{y}, y) yields not a . In this case, \vec{x}, \vec{y}, x, y are called an “em-bedded OR”.*

⁷ The Corners Lemma was also implicitly used in [11]. It was referred to as “Corners Lemma” for the first time in [10].

⁸ This is a special case of the Partition Lemma as defined in [10] for $t = n - 1$.

Due to a lack of a more detailed characterization of n -ary privately computable functions, the only way to show that a function is privately computable is to give a concrete protocol that fulfills this task (as in Theorems 3 and 6). As first observed by Benaloh [5], there is a simple protocol to privately compute modular sums.

Lemma 4 [5] $f(x_1, x_2, \dots, x_n) = \sum_{i=1}^n x_i \bmod p$ is privately computable.

Proof: Each agent i chooses n random values $x_{ij} \in \mathbb{Z}_p$ so that the modular sum $\sum_{j=1}^n x_{ij} \bmod p = x_i$. He then sends each addend x_{ij} to agent j and keeps x_{ii} . After all agents have done this, each agent i publishes $s_i = \sum_{j=1}^n x_{ji} \bmod p$, *i.e.*, the modular sum of his remaining x_{ii} and the $n - 1$ addends he received. $f(x_1, x_2, \dots, x_n) = \sum_{i=1}^n s_i \bmod p$ can be computed by each participant. \square

For example, this protocol could be used to allow a group of participants to privately compute their average salary without revealing any individual salary.

3. First-Price Auctions

Every social-welfare-maximizing auction assigns the item for sale to the bidder who values it most (when bidders' valuations are positive and the seller's valuation is zero). In other words, the *arg max* function yields the auction winner. For the case of first-price auctions, the *max* function yields the selling price. In the following theorem, we prove that neither function can be computed fully privately.

Theorem 2 *The max and arg max functions cannot be emulated by fully private protocols.*

Proof: Let $f_{\max}(x_1, x_2, \dots, x_n) = \max\{x_1, x_2, \dots, x_n\}$ and $\vec{x} = (2, \underbrace{1, \dots, 1}_{n-2})$, $\vec{y} = (1, \underbrace{1, \dots, 1}_{n-2})$, $x = 2$, and $y = 1$.

Then $f_{\max}(\vec{x}, x) = f_{\max}(\vec{x}, y) = f_{\max}(\vec{y}, x) = 2$. However, $f_{\max}(\vec{y}, y) = 1$.

f_{\max}	1	2	...
$1, 1, \dots, 1$	1	2	
$2, 1, \dots, 1$	2	2	

It follows from Lemma 3 that f_{\max} is *not* privately computable.

Let $f_{\arg \max}(x_1, x_2, \dots, x_n) = \arg \max_{i=1}^n \{x_i\}$ be a function that yields the index of the greatest argument. Furthermore, let $\vec{x} = (2, \underbrace{1, \dots, 1}_{n-2})$, $\vec{y} = (4, \underbrace{1, \dots, 1}_{n-2})$, $x = 5$, and $y = 3$. Then $f_{\arg \max}(\vec{x}, x) = f_{\arg \max}(\vec{x}, y) = f_{\arg \max}(\vec{y}, x) = n$. However, $f_{\arg \max}(\vec{y}, y) = 1$.

$f_{\arg \max}$	3	5	...
2, 1, ..., 1	n	n	
4, 1, ..., 1	1	n	

It follows from Lemma 3 that $f_{\arg \max}$ is *not* privately computable. \square

Even though the winner and selling price functions cannot be computed separately, it turns out that it is possible to privately compute both at the same time. Let $\vec{b} = (b_1, b_2, \dots, b_n)$ be the vector of submitted bids.

Definition 1 *The first-price sealed-bid auction's public outcome is defined by the following function.*

$$f^1(\vec{b}) = (\max(\vec{b}), \arg \max(\vec{b}))$$

Theorem 3 *The first-price sealed-bid auction can be emulated by a fully private protocol that requires $2^v - 1$ rounds of interaction in the worst-case where v is the number of bits used to represent a bid. There is no more efficient private protocol for this task.*

Proof: When examining f^1 for just two bidders, it turns out that the Corners Lemma is not applicable (irrelevant of tie resolution). Bold numbers denote the winner's index.

f^1	1	2	3	4	5	...
1		(2, 2)	(3, 2)	(4, 2)	(5, 2)	
2	(2, 1)		(3, 2)	(4, 2)	(5, 2)	
3	(3, 1)	(3, 1)		(4, 2)	(5, 2)	
4	(4, 1)	(4, 1)	(4, 1)		(5, 2)	
5	(5, 1)	(5, 1)	(5, 1)	(5, 1)		
\vdots						

It is important to note that the lack of an embedded OR is not sufficient to show that a function is privately computable (not even for only two agents). The Corners Lemma can only be used to prove that a function is *not* privately computable. However, Kushilevitz has given a complete characterization of privately computable functions for *two* agents [16] and it turns out that $f^1(b_1, b_2)$ is indeed privately computable. Even more interestingly, we can show that f^1 is privately computable for *any* number of agents which is beyond the capabilities of existing theory. Nevertheless, such a result can always be proven constructively by providing a specific protocol that privately computes the desired function. Consider the following protocol.

1. $j = 2^v$
2. Each agent broadcasts either 1 or 0 depending on whether he is willing to pay price j or not.⁹

⁹ Different tie resolution policies can be implemented by prescribing the order of each agent's broadcast, e.g., random order or priority order. In any case, a tie will result in some (tiny) amount of additional information to be revealed (if there are more than two bidders).

3. If all agents broadcasted 0, set $j = j - 1$ and proceed to step 2. Otherwise, j is the selling price and the bidder(s) who submitted 1 wins the auction.

This protocol builds on the same principle as the Dutch (or descending) auction in which an auctioneer gradually (or continuously) lowers the selling price until a bidder expresses his willingness to buy (e.g., [19]).

Even though the theoretic results in [16] cannot decide the private computability of general n -ary functions, they give a lower bound for the number of rounds needed to compute a given n -ary function (if it can be computed privately). In the interest of space, we do not go into the details of [16], but merely state that f^1 's so-called decomposition tree has depth $2^v - 1$ when bids consist of v bits. This implies exponential round complexity (in fact, $2^v - 1$ rounds) and thus the optimality of the proposed protocol. \square

As mentioned in Section 1, unconditionally private protocols require a complete network of private channels. An outstanding property of the proposed protocol is that the availability of a *broadcast channel* replaces the need for private channels as there is no interaction between bidders. In practice, it is usually much easier to establish a broadcast channel than private channels between agents. This can be seen by the popularity of real-world Dutch auctions in flower of fish markets. Also, in reality, the physical presence of bidders allows for very efficient synchronization via a common timer.

As a pleasant side effect, the availability of a secure broadcast channel¹⁰ provides security against active adversaries, i.e., privacy is guaranteed even in the presence of bidders that deviate from the protocol specification. Generally, security against active adversaries requires the extensive use of costly zero-knowledge proofs/arguments.

It might seem that the outcome function defined in Definition 1 reveals the minimal amount of information needed to perform the required transaction (e.g., selling of a good). However, the notion of minimal revelation can be refined even further by moving to *asymmetric* information revelation. It is not necessary that losing bidders learn who won the auction and which price this agent has to pay. Since a protocol that is secure against active adversaries is *provably correct*, there is no need to publicly announce the outcome for reasons of transparency. On account of this, we will now consider the joint computation of n functions $f_i^1(\vec{b})$ so that bidder i only learns the result of his private outcome function.

¹⁰ Byzantine agreement [17] is not feasible in this context as it either requires intractability assumptions or the trustworthiness of two thirds of the agents.

Definition 2 The first-price sealed-bid auction's private outcome function is

$$f_i^1(\vec{b}) = \begin{cases} b_i & \text{if } i = \arg \max(\vec{b}) \\ 0 & \text{otherwise} \end{cases}.$$

In practice, it might be desirable to include the seller in the protocol and compute all f_i^1 functions for him as well. This prevents a bidder from aborting the protocol if he is unsatisfied with the auction outcome, leaving the seller uninformed about that outcome.

Theorem 4 There is no fully private protocol that computes the private outcome $f_i^1(\vec{b})$ of a first-price sealed bid auction.

Proof: With the notable exception of [3] which only addresses the two-party case, the theory on privately computable functions only deals with the case where *all* agents get to know the function value. The results in this setting cannot be directly transferred to a setting where only *one* agent learns the function value. However, the following lemma is sufficient to show the impossibility of private computation in the latter case.

Lemma 5 If a function $f(x_1, x_2, \dots, x_n)$ cannot be privately computed so that all agents learn the function value, it cannot be computed for a single agent (or any subset of agents).

Proof of Lemma 5: Indirect proof. If function f can be computed so that a single agent learns the output, then it can also be computed so that all agents receive the function value by simple adding a protocol step in which the designated agent sends the output to all remaining agents. \square

Now, we continue the proof of Theorem 4. With the help of Lemma 5, we can prove the impossibility of computing f_i^1 by using a chain of necessary conditions. It suffices to use Lemma 3 to show the impossibility of a private protocol for any n .

Let $\vec{x} = (\underbrace{4, 1, \dots, 1}_{n-2}, \underbrace{2, 1, \dots, 1}_{n-2})$, $x = 1$, and $y = 3$,

and consider the outcome function of bidder i : $f_n^1(\vec{x}, x) = f_n^1(\vec{x}, y) = f_n^1(\vec{y}, x) = 0$. However, $f_n^1(\vec{y}, y) = 3$.

f_n^1	1	3	...
2, 1, ..., 1	0	3	
4, 1, ..., 1	0	0	

It follows from Lemma 3 that f_i^1 is *not* privately computable for any i . Lemma 5 implies that there is no protocol to compute f_i^1 privately so that only bidder i learns the outcome. \square

4. Second-Price Auctions

In this section, we investigate the existence of fully private protocols that emulate second-price sealed-bid (Vickrey) auctions.

Definition 3 The second-price sealed-bid (Vickrey) auction's public outcome is defined by the following function.¹¹

$$f^2(\vec{b}) = (\max(\vec{b}_{-\arg \max(\vec{b})}), \arg \max(\vec{b}))$$

Proposition 1 There is a fully private protocol that emulates the second-price sealed-bid auction for two bidders.

Proof: When there are just two bidders, the Dutch auction style protocol proposed in the proof of Theorem 3 can be applied in reverse to find the *lowest* instead of the highest bid. This is equivalent to a two-bidder English (ascending) auction. Beginning at the lowest possible price, the price rises incrementally until one of the bidders is *not* willing to pay the given price. This does reveal the identity of the second-highest bidder, but this information can always be inferred from the outcome if there are only two bidders. \square

Unlike the first-price auction, the second-price auction's outcome can *not* be computed fully privately if there are more than two bidders.

Theorem 5 There is no fully private protocol that emulates the second-price sealed-bid auction for more than two bidders.

Proof: We construct a general counter-example for any $n > 2$. Let $\vec{x} = (3, 2, \underbrace{1, \dots, 1}_{n-3})$, $\vec{y} = (3, \underbrace{1, \dots, 1}_{n-2})$, $x = 2$, and $y = 1$. Then $f^2(\vec{x}, x) = f^2(\vec{x}, y) = f^2(\vec{y}, x) = (2, 1)$ (Bidder 1 wins the auction at price 2). However, $f^2(\vec{y}, y) = (1, 1)$ (Bidder 1 wins at price 1).

f^2	1	2	...
3, 1, 1, ..., 1	(1, 1)	(2, 1)	
3, 2, 1, ..., 1	(2, 1)	(2, 1)	

It follows from Lemma 3 that f^2 is *not* privately computable. \square

The positive impact of such an impossibility result is that, in the future, no efforts need to be wasted in trying to find a protocol with the claimed properties. This effect is enhanced in Section 5.1 where it is shown that even the search for a Vickrey auction protocol that only hides *some* specific information is futile.

¹¹ \vec{b}_{-i} denotes vector \vec{b} with component i removed.

Unfortunately, there is also no fully private second-price auction protocol in which only the winner learns the outcome, *i.e.*, the second-highest bid.

Definition 4 *The second-price sealed-bid auction's private outcome function is*

$$f_i^2(\vec{b}) = \begin{cases} \max(\vec{b}_{-i}) & \text{if } i = \arg \max(\vec{b}) \\ 0 & \text{otherwise} \end{cases}$$

Corollary 1 *There is no protocol that computes the private outcome of a second-price sealed bid auction.*

Proof: The counter-example given in the proof of Theorem 5 (\vec{x} , \vec{y} , x , and y) also yields an embedded OR when considering f_1^2 , the outcome function for bidder 1. \square

5. Security Model Relaxations

The unconditional full privacy model considered in this paper is very strict. It is natural to ask for relaxations that may invalidate the impossibility of private second-price auctions. The following options come to mind.

- Allow partial revelation of bids, *e.g.*, the highest bid, by modifying the outcome function
- Allow coalitions of bidders to uncover information by relaxing full privacy
- Guarantee high probability of correctness instead of correctness for sure

In the following, we will investigate the private emulation of second-price auctions under these weakened assumptions.

5.1. Partial Revelation

The more information an outcome function reveals about the bids, the more likely it can be privately computed. In this section, we study whether the revelation of a limited amount of information enables the private computation of second-price auctions. One of the weakest privacy requirements is anonymity.

Definition 5 (Anonymity) *An auction protocol is anonymous if the outcome does not change when the bids of two losing bidders are exchanged.*

Even under this weak requirement, there is no second-price auction protocol that protects just a single losing bid.¹² On

the positive side, there is an anonymous protocol in which the highest bid remains private but all other bid amounts are revealed (but not who submitted which bid).

Theorem 6 *A fully private protocol that anonymously emulates the second-price sealed-bid auction reveals information about all losing bids (in the worst case).*

Proof: By contradiction. In an anonymous auction, the bids can only be distinguished by their numerical order. Assume that the k th highest bid is not revealed ($k > 2$ because the second-highest bid has to be revealed in a Vickrey auction). Let $b^{(i)}$ be the i th order statistic of \vec{b} , *i.e.*, the i th highest bid. Then $g_k(\vec{b}) = (b^{(1)}, b^{(2)}, \dots, b^{(k-1)}, b^{(k+1)}, b^{(k+2)}, \dots, b^{(n)}, \arg \max(\vec{b}))$ defines the modified second-price outcome function that only hides bid $b^{(k)}$. We will now apply Lemma 3 to this general case.

Let $\vec{x} = (n, n-1, \dots, n-k+2, n-k, n-k-1, \dots, 1)$, $\vec{y} = (n, n-1, \dots, n-k+3, n-k+1, n-k, \dots, 1)$, $x = n-k+2$, $y = n-k+1$. Then $g_k(\vec{x}, x) = g_k(\vec{x}, y) = g_k(\vec{y}, x) = (\vec{x}, 1)$. However, $g_k(\vec{y}, y) = (\vec{y}, 1)$ which proves the impossibility of fully privately computing g_k according to Lemma 3. The following table shows an example for four bidders when only the third highest bid should be kept private ($n = 4, k = 3$).

g_k	2	3	...
4, 2, 1	(4, 2, 1, 1)	(4, 3, 1, 1)	
4, 3, 1	(4, 3, 1, 1)	(4, 3, 1, 1)	

It remains to be shown that it is possible to privately compute function $g(\vec{b}) = (b^{(2)}, b^{(3)}, \dots, b^{(n)}, \arg \max(\vec{b}))$ which reveals the winner and all losing bid amounts (independently of bidders' identities). Interestingly, this task can be fulfilled by a protocol similar to an anonymized English (ascending) auction.

1. $j = 1$
2. Each agent i sets $x_i = \begin{cases} 1 & \text{if } b_i \leq j \\ 0 & \text{otherwise} \end{cases}$.
3. Agents jointly compute $s = \sum_{i=1}^n x_i \bmod (n+1)$ according to the protocol defined in Lemma 4.
4. If $s > 1$, set $j = j + 1$, and proceed to step 2.
5. If $b_i \geq j$, agent i broadcasts his identity and wins the auction.

As mentioned in Section 1, there are standard cryptographic means to ensure that agents follow the protocol truthfully and do not manipulate, for example by wrongfully broadcasting their identity in step 5. \square

¹² Interestingly, revealing the identity of the second-highest bidder (in addition to the Vickrey auction outcome) does not help either. Since proving this fact requires a stronger tool (the undecomposability of matrices [16]) than the Corners Lemma, we omit the proof due to limited space.

5.2. Uncovering by Coalitions

So far, we required that no coalition consisting of less than n agents may be able to uncover private information (full privacy). In this section, we examine whether loosening this restriction will enable the private emulation of Vickrey auctions. It turns out that the argument used in the proof of Theorem 5 also works with a version of the Partition Lemma (Lemma 2) for a model in which complete information can be uncovered by coalitions including at least half of the bidders (in contrast to coalitions of n bidders as in full privacy). As mentioned in Section 1, assuming that a majority of participants is trustworthy allows the private computation of *any* function (including f^2). As a consequence, the private computation of the Vickrey auction's outcome is only possible when assuming that coalitions consist of strictly less than $\frac{n}{2}$ bidders. A higher threshold cannot be obtained.

5.3. Correctness with High Probability

In this section, we review whether allowing an error probability enables the private computation of the second-price auction. It has been shown that allowing error probability ε (where $\varepsilon < \frac{1}{2}$) does not enable the private computation of functions that cannot be computed with perfect correctness in (i) the Boolean n -party case [11] and (ii) the general 2-party case [16]. The auction setting we consider belongs to the general n -party case for which such a result is not known. However, it seems likely that the equivalence of error-free and mostly-correct private computation also holds for this setting.

6. Conclusions and Future Work

Sealed-bid auctions are not only widely used for the selling of goods, they also have been shown to be applicable to task assignment, scheduling, and finding the shortest path in a network with selfish nodes. Bid privacy is of increasing importance in such auctions, and various schemes that avoid blind trust in a single auctioneer have been proposed recently. In contrast to existing work, this paper deals with *unconditional full privacy*, i.e., privacy that relies neither on trusted third parties (like auctioneers) or trusted fractions of bidders, nor on computational intractability assumptions (like the hardness of factoring). We investigated the availability of distributed protocols that allow a group of bidders to jointly determine the outcome of first-price and second-price auctions by exchanging messages according to some predefined rules and without revealing unnecessary information. We derived several impossibility and possibility results in this domain:

The first-price auction can be emulated by fully private protocols. However, such a protocol will always have exponential round complexity. When modifying the specification so that only the winning bidder learns the outcome, the first-price auction cannot be emulated fully privately.

There is a fully private protocol that emulates the second-price auction for *two* bidders. However, the second-price auction cannot be emulated by a private protocol for more than two bidders even when

- just the auction winner learns the outcome,
- just protecting a single losing bid (but maintaining anonymity), or
- tolerating the revelation of complete information to a coalition of at least half of the bidders.

On the positive side, we proposed a fully private second-price auction protocol that is anonymous and only hides the highest bid.¹³

Future work includes the investigation of unconditionally fully private auction protocols that only reveal partial information on each bid, e.g. “the lowest bid is greater than 10”. Theorem 6 states that some information on *all* losing bids has to be revealed in the worst case. It seems worthwhile to minimize this amount of information for practical instances. So far, theoretic results on minimum revelation protocols are only known for two agents [2]. A related field of study is that of using an elicitor that incrementally asks questions from the bidders about their bids on an as-needed basis until the elicitor has enough information to determine the auction winner [6, 12]. This approach also provides partial unconditional privacy and it might be possible to transfer results from one setting to the other.

Acknowledgements

This material is based upon work supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/1-1 and by the National Science Foundation under grants IIS-9800994, ITR IIS-0081246, and ITR IIS-0121678.

References

- [1] M. Abe and K. Suzuki. M+1-st price auction using homomorphic encryption. In *Proceedings of the 5th International Conference on Public Key Cryptography (PKC)*, volume 2274 of *Lecture Notes in Computer Science (LNCS)*, pages 115–224. Springer, 2002.
- [2] R. Bar-Yehuda, B. Chor, and E. Kushilevitz. Privacy, additional information, and communication. In *Proceedings of*

13 Much better results can be obtained in a setting that relies on computational intractability. We recently proposed a fully private 3-round Vickrey auction protocol in this setting [7].

the 5th IEEE Conference on Structure in Complexity Theory, pages 55–65, 1990.

- [3] A. Beimel, T. Malkin, and S. Micali. The all-or-nothing nature of two-party secure computation. In *Advances in Cryptology - Proceedings of the 19th Annual International Cryptology Conference (CRYPTO)*, volume 1666 of *Lecture Notes in Computer Science (LNCS)*, pages 80–97. Springer, 1999.
- [4] M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the 20th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 1–10. ACM Press, 1988.
- [5] J. Benaloh. Secret sharing homomorphisms: Keeping shares of a secret secret. In *Advances in Cryptology - Proceedings of the 13th Annual International Cryptology Conference (CRYPTO)*, volume 263 of *Lecture Notes in Computer Science (LNCS)*, pages 251–260. Springer, 1987.
- [6] F. Brandt. Cryptographic protocols for secure second-price auctions. In M. Klusch and F. Zambonelli, editors, *Cooperative Information Agents V*, volume 2182 of *Lecture Notes in Artificial Intelligence (LNAI)*, pages 154–165. Springer, 2001.
- [7] F. Brandt. Fully private auctions in a constant number of rounds. In R. N. Wright, editor, *Proceedings of the 7th Annual Conference on Financial Cryptography (FC)*, volume 2742 of *Lecture Notes in Computer Science (LNCS)*, pages 223–238. Springer, 2003.
- [8] F. Brandt, W. Brauer, and G. Weiß. Task assignment in multiagent systems based on Vickrey-type auctioning and leveled commitment contracting. In M. Klusch and L. Kerschberg, editors, *Cooperative Information Agents IV*, volume 1860 of *Lecture Notes in Artificial Intelligence (LNAI)*, pages 95–106. Springer, 2000.
- [9] D. Chaum, C. Crépeau, and I. Damgård. Multi-party unconditionally secure protocols. In *Proceedings of the 20th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 11–19. ACM Press, 1988.
- [10] B. Chor, M. Geréb-Graus, and E. Kushilevitz. On the structure of the privacy hierarchy. *Journal of Cryptology*, 7(1):53–60, 1994.
- [11] B. Chor and E. Kushilevitz. A zero-one law for Boolean privacy. In *Proceedings of the 21st Annual ACM Symposium on the Theory of Computing (STOC)*, pages 36–47. ACM Press, 1989.
- [12] W. Conen and T. Sandholm. Preference elicitation in combinatorial auctions. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, pages 256–259. ACM Press, 2001.
- [13] O. Goldreich. *Foundations of Cryptography: Volume 1, Basic Tools*. Cambridge University Press, 2001.
- [14] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *Proceedings of the 19th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 218–229. ACM Press, 1987.
- [15] H. Kikuchi. (M+1)-st-price auction protocol. In *Proceedings of the 5th Annual Conference on Financial Cryptography (FC)*, volume 2339 of *Lecture Notes in Computer Science (LNCS)*, pages 351–363. Springer, 2001.
- [16] E. Kushilevitz. Privacy and communication complexity. In *Proceedings of the 30th Symposium on Foundations of Computer Science (FOCS)*, pages 416–421. IEEE Computer Society Press, 1989.
- [17] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.
- [18] M. Naor, B. Pinkas, and R. Sumner. Privacy preserving auctions and mechanism design. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, pages 129–139. ACM Press, 1999.
- [19] J. A. Rodríguez-Aguilar, F. J. Martín, P. N., P. García, and C. Sierra. Competitive scenarios for heterogeneous trading agents. In *Proceedings of the 2nd International Conference on Autonomous Agents*, pages 293–300. ACM Press, 1998.
- [20] J. A. Rodríguez-Aguilar, P. Noriega, C. Sierra, and J. Padget. FM96.5: A Java-based electronic auction house. In *Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM-97)*, 1997.
- [21] M. H. Rothkopf and R. M. Harstad. Two models of bid-taker cheating in Vickrey auctions. *Journal of Business*, 68(2):257–267, 1995.
- [22] M. H. Rothkopf, T. J. Teisberg, and E. P. Kahn. Why are Vickrey auctions rare? *Journal of Political Economy*, 98(1):94–109, 1990.
- [23] T. Sandholm. Issues in computational Vickrey auctions. *International Journal of Electronic Commerce, Special issue on Intelligent Agents for Electronic Commerce*, 4(3):107–129, 2000. Special Issue on Applying Intelligent Agents for Electronic Commerce. A short, early version appeared at the 2nd International Conference on Multi-Agent Systems (ICMAS), pp. 299–306, 1996.
- [24] T. Sandholm. eMediator: A next generation electronic commerce server. *Computational Intelligence*, 18(4):656–676, 2002. Special issue on Agent Technology for Electronic Commerce. Early versions appeared in the Conference on Autonomous Agents (AGENTS-00), pp. 73–96, 2000; and AAAI-99 Workshop on AI in Electronic Commerce, pp. 46–55, 1999.
- [25] W. Vickrey. Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.
- [26] P. R. Wurman, M. P. Wellman, and W. E. Walsh. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *Proceedings of the 2nd International Conference, on Autonomous Agents*, pages 301–308, 1998.
- [27] M. Yokoo and K. Suzuki. Secure multi-agent dynamic programming based on homomorphic encryption and its application to combinatorial auctions. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 112–119. ACM Press, 2002.