FELIX BRANDT University of Munich and TUOMAS SANDHOLM Carnegie Mellon University

We investigate whether it is possible to preserve privacy in sealed-bid auctions to a maximal extent. In particular, this paper focuses on *unconditional full privacy*, *i.e.*, privacy that relies neither on trusted third parties (like auctioneers), 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 bid size, and there is no more efficient protocol. On the other hand, we prove the impossibility of privately emulating the second-price sealed-bid auction for more than two bidders. This impossibility holds even when relaxing various privacy constraints such as allowing the revelation of all but one losing bid (while maintaining anonymity) or allowing the revelation of the second highest bidder's identity.

Categories and Subject Descriptors: E.4 [**Data**]: Coding and Information Theory; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*; K.4.4 [**Computing Milieux**]: Electronic Commerce—*Security*

General Terms: Economics, Security, Theory

Additional Key Words and Phrases: Auctions, Multiparty Computation

1. INTRODUCTION

Auctions are key mechanisms for allocating scarce resources among multiple agents. 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

ACM Journal Name, Vol. V, No. N, Month 20YY, Pages 1–20.

An early version of this paper appeared in the Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS, July 19–23, 2004), pages 810–817, ACM Press.

Authors' addresses: Felix Brandt, Computer Science Department, University of Munich, Oettingenstr. 67, 80538 Munich, Germany, email: brandtf@tcs.ifi.lmu.de. Tuomas Sandholm, Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA, email: sandholm@cs.cmu.edu.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee. © 20YY ACM 0000-0000/20YY/0000-0001 \$5.00

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 study the possibility of executing the most common types of sealed-bid auctions in a way that preserves the bidders' privacy to a maximal extent.

Our setting consists of one seller and n bidders who intend to come to an agreement on the selling of a single good.¹ The two major sealed-bid mechanisms that yield such an agreement are the first-price auction and the second-price auction (aka. Vickrey auction, named after Nobel Laureate William Vickrey who first proposed it [Vickrey 1961]). In both mechanisms, each bidder submits a sealed bid to a trusted-third party called the auctioneer 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). Thus, in contrast to the first-price auction, bidders need not estimate other bidders' valuations. Interestingly, the side-effects of this striking advantage are said to contribute to the fact that second-price auctions are not commonly used in practice, for two reasons [Rothkopf et al. 1990; Rothkopf and Harstad 1995; Sandholm 2000]:

- (1) Bidders are reluctant to reveal their true valuations to the auctioneer since the auctioneer can exploit this information during and after the auction, or spread it to others in ways that adversely affect the bidder.
- (2) Bidders doubt the correctness of the result as they do not pay what they bid. For example, the auctioneer might create a fake second highest bid slightly below the highest bid in order to increase his revenue (see, *e.g.*, [Porter and Shoham 2005]).

Both issues mentioned above are rooted in 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 (see Section 2). 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 computational power of the adversary (*i.e.*, a virtual entity that corrupts parties in order to breach privacy) is polynomially-bounded in a security parameter.

¹All 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. Our results on second-price auctions extend to a generalization of second-price auctions called *uniform-price* auctions which are used to sell multiple units of the same item in a single auction.

ACM Journal Name, Vol. V, No. N, Month 20YY.

(3) One-way functions exist.

Assumption (1) has its origin in secure multiparty computation (MPC) which is commonly used to distribute trust onto several auctioneers. However, a coalition of *all* auctioneers can never be prevented from breaching privacy. For this reason, we advocate a security model in which the computation of the auction outcome is distributed *on the bidders themselves*. We say that an auction protocol is *fully private* if it is distributed on bidders and no coalition of bidders is capable of revealing private information about any of the remaining bidders. This is sometimes also called (n - 1)-privacy.

Assumptions (2) and (3) deal with computational intractability. When relying on intractability assumptions (e.g., the hardness of factoring), it is known that MPC allows the computation of arbitrary functions so that no private information can be uncovered by a polynomially-bounded adversary [Goldreich et al. 1987]. While intractability assumptions are well-established in practice, they are somewhat unsatisfactory from a theoretical point of view. Unfortunately, assumption (3) not only relies on the unproven assumption $P \neq NP$ but also on the *average-case* hardness of computational problems about which even less is known. Moreover, even when assumption (3) is true, it may be possible to breach privacy in the future when sufficient computational power becomes available, thus violating assumption (2).² The results in this paper do *not* rely on intractability. This is called *unconditional* or *information-theoretic* privacy as the adversary's computational power is unlimited.

It is known that only a restricted class of functions can be computed while maintaining unconditional full privacy (see Section 3 for further details).³ In the unconditional privacy model, computational intractability assumptions are replaced with a complete network of private channels between agents. However, sometimes, a particular protocol can also be implemented by just providing a broadcast channel (see, *e.g.*, Theorem 4.3) which is usually easier to establish in practice.

Regarding the adversary, we focus on what is known as a *passive* (aka. *honest-but-curious*) adversary in the cryptographic literature, *i.e.*, we assume that participants follow the prescribed protocol honestly. This assumption does not restrict the applicability of our results because there are standard cryptographic techniques that force agents to act according to a protocol (see, *e.g.*, [Goldreich 2001]).⁴ However, using these techniques will incur overhead. Clearly, *negative* results in the passive adversary model also hold in stronger adversarial models, *e.g.*, models with active or adaptive adversaries.

In a nutshell, this paper investigates the possibility of distributed protocols that allow n bidders to jointly determine the outcome of first-price or second-price auc-

²This does not require super-polynomial computational power. The security parameter used for a protocol might be too low with respect to future computational power. *E.g.*, 512-bit RSA keys were considered secure some years ago but are not secure anymore.

 $^{^{3}}$ 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 passive adversary model [Ben-Or et al. 1988; Chaum et al. 1988].

⁴For example, it would be possible to use zero-knowledge *arguments* in a model where the adversary's computational power is only assumed to be polynomially bounded *during* the protocol.

tions by exchanging messages without revealing any additional information beyond the auction outcome. In the rest of this paper, this is called *emulation* of an auction. In the case of ties, we deliberately leave the outcome undefined. As a consequence, the impossibility results of this paper hold regardless of tie resolution.

The remainder of this paper is organized as follows. Related work on cryptographic auction protocols is reviewed in Section 2. Section 3 presents some known theoretical results that we will leverage in our proofs. In Section 4, we study the existence of private protocols that emulate first-price and second-price auctions, respectively. 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 private auction protocols under those loosened restrictions. The paper concludes with a summary of the results and a brief outlook in Section 6. The Appendix contains an example run of the auction protocol proposed in Section 4.

2. RELATED RESEARCH

The interest in cryptographic protocols for auctions has been dramatically increasing. Starting with the work by Nurmi and Salomaa [1993] and Franklin and Reiter [1996], numerous secure sealed-bid auction schemes have been proposed in recent years, *e.g.*, [Abe and Suzuki 2002; Baudron and Stern 2001; Brandt 2003; Cachin 1999; Harkavy et al. 1998; Juels and Szydlo 2002; Kikuchi et al. 1998; Kikuchi 2001; Lipmaa et al. 2002; Naor et al. 1999; Sako 2000].

The proposed protocols essentially fall into three categories depending on the underlying security model. First, there are protocols where the trust is distributed on multiple, symmetric auctioneers who jointly determine the outcome using some form of threshold MPC (e.g., [Harkavy et al. 1998; Kikuchi et al. 1998; Kikuchi 2001; Sako 2000]). Other protocols introduce a second party, for example an "auction issuer" or "auction authority", in addition to the auctioneer, and employ asymmetric MPC, for instance Yao's garbled circuit technique (e.g., Abe and Suzuki 2002; Baudron and Stern 2001; Cachin 1999; Juels and Szydlo 2002; Lipmaa et al. 2002; Naor et al. 1999]). Finally, there are protocols where bidders themselves jointly compute the auction outcome without relying on trusted third parties at all [Brandt 2002; 2003; Brandt and Sandholm 2005]. The main advantage of these protocols is that they are *fully private*, *i.e.*, when relying on computational intractability assumptions, no coalition of parties is capable of breaching privacy. The drawbacks implied by such a model are low robustness and relatively high computational and communication complexity (although round complexity is low and constant).

All of the above protocols rely on diverse intractability assumptions. Rather than designing practical auction protocols that make trade-offs between efficiency and privacy, the goal of this paper is to investigate "what can be achieved at all" while maintaining privacy to a maximal extent, namely unconditional full privacy. We believe that such an approach sheds light on the general feasibility of private auction protocols and may deter others from trying to devise protocols for settings in which no protocol can possibly exist. For example, two recent second-price protocols [Peng

et al. 2002; López et al. 2004] were incorrectly claimed to be unconditionally secure (which they cannot be according to Theorem 4.9). While the protocol by Peng et al. [2002] clearly relies on the Diffie-Hellman assumption, the protocol by López et al. [2004] reveals a substantial amount of bid statistics to colluding bidders. The latter shortcoming has been acknowledged in a subsequent publication by the same authors [Rodríguez and López 2006].

3. PRELIMINARIES

In this section we review some key results which we will use as building blocks in our proofs. When we refer to "privacy" in the following, we always mean unconditional full privacy. The computational model we employ is the standard information-theoretic private-channels model introduced independently by Ben-Or et al. [1988] and Chaum et al. [1988], inspired by earlier work of Yao [1979]. Function $f(x_1, x_2, \ldots, x_n)$ is jointly computed by agents using a distributed, randomized protocol consisting of several rounds. In order to enable secure message exchange, we assume a complete synchronous network of private channels between agents. In each round, each agent may send a message to any other agent. Each message an agent sends is a function of his input x_i , his independent random input r_i , the messages he received so far, and the recipient. When the protocol is finished, all agents know the value of $f(x_1, x_2, \ldots, x_n)$.

Intuitively, a distributed protocol for computing a function $f(x_1, x_2, \ldots, x_n)$ is private if any coalition of agents is incapable of learning any information besides what can be inferred from $f(x_1, x_2, \ldots, x_n)$ and the coalition's inputs. More precisely, for any pair of input vectors the messages seen by agents belonging to a curious coalition $T \subseteq N$ should be identically distributed if the agents in T supply identical inputs and the function yields the same output in both cases. In order to state this more formally, we need to introduce two notational conventions. Let VIEW_T be a function that, given a vector of individual inputs \vec{x} and random values \vec{r} , yields the concatenation of all (prefix-free) messages exchanged between members of T and $\bar{T} = N \setminus T$ in the protocol. Further, let $\langle t \mid V \rangle$ be the probability distribution of term t with the probability taken over all random variables in V.

Definition 3.1. A protocol for computing f is private if for all $T \subseteq N$, any pair of input vectors $\vec{x}, \vec{y} \in X^n$ that satisfy $\forall i \in T : x_i = y_i$ and $f(\vec{x}) = f(\vec{y})$, and any choice of random inputs $\{r_i\}_{i \in T}$

$$\langle \operatorname{VIEW}_T(\vec{x}, \{r_i\}_{i \in T}) \mid \{r_i\}_{i \in \bar{T}} \rangle = \langle \operatorname{VIEW}_T(\vec{y}, \{r_i\}_{i \in T}) \mid \{r_i\}_{i \in \bar{T}} \rangle.$$

A complete characterization of all privately computable *Boolean* functions has been given by Chor and Kushilevitz [1989].

THEOREM 3.2. [Chor and Kushilevitz 1989] A Boolean function is privately computable if and only if it is of the form $f(x_1, x_2, ..., x_n) = B_1(x_1) \oplus B_2(x_2) \oplus \cdots \oplus B_n(x_n)$, 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 parties.

3.1 2-Party Computation

In the following, we present the complete characterization of privately computable 2-ary functions proposed by Kushilevitz [1989]. The characterization is based on the representation of any 2-ary function f as a function table (or matrix) M_f .

Definition 3.3. An $m \times m$ matrix is called rows-decomposable (or columnsdecomposable) if there is a partitioning of rows (or columns) $P = \biguplus_i P_i$, so that there is no pair of rows (or columns) from two different partitions which share an identical entry at the same position, *i.e.*, there are no $r \in P_i$, $s \in P_j$ with $i \neq j$, and $k \in \{1, 2, \ldots, m\}$ so that $r_k = s_k$.

Definition 3.4. A matrix is *decomposable* if it is monochromatic, *i.e.*, all of its elements are identical, or it can be rows- or columns-decomposed into decomposable submatrices.

Tables I and II show examples of a decomposable and a non-decomposable matrix, respectively.

Table I.	A de	ecom	posable matrix
	Α	С	В
	\mathbf{C}	Α	В
	D	D	В

Table II. A no	n-decomposable matrix	
----------------	-----------------------	--

Α	В	В	
Α	\mathbf{C}	D	
\mathbf{E}	Е	D	

THEOREM 3.5. [Kushilevitz 1989] A 2-ary function f is privately computable if and only if the function's matrix M_f is decomposable.

Kushilevitz also showed that decomposing matrices is equivalent to privately computing the corresponding function where each decomposition step relates to one round in the protocol.

THEOREM 3.6. [Kushilevitz 1989] The minimal number of rounds needed to privately compute function f is given by the depth of M_f 's minimal decomposition tree, i.e., the minimal number of consecutive rows- or columns-decompositions until all resulting submatrices are monochromatic.

In many cases, it is *sufficient* to test a simple condition in order to show the *impossibility* of computing certain functions. Such a condition is given by the so-called Corners Lemma.⁵ As a matter of fact, the more complex notion of decomposability will only be used in Theorem 5.3.

⁵The Corners Lemma was implicitly used by Chor and Kushilevitz [1989] and Kushilevitz [1989]. It was referred to as "Corners Lemma" for the first time by Chor et al. [1994].

ACM Journal Name, Vol. V, No. N, Month 20YY.

LEMMA 3.7 CORNERS LEMMA. [Kushilevitz 1989] Let $f : X \times Y \to Z$ 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 $f(x_2, y_2) = a$.

Clearly, a function that violates the Corners Lemma is not decomposable. The reverse, however, does not hold (see, e.g., the function matrix given in Table II).

3.2 *n*-Party Computation

As mentioned above, a complete characterization of general n-ary functions that can be privately computed is not known. However, there is a necessary condition for the private computability of such functions.⁶

LEMMA 3.8. [Chor and Kushilevitz 1989] Let $f : X_1 \times X_2 \times \cdots \times X_n \to Z$ be a privately computable n-ary function. Then, for each $i \in \{1, 2, ..., n\}$ the 2-ary function $f_2(x_i, (x_1, x_2, ..., x_{i-1}, x_{i+1}, x_{i+2}, ..., x_n)) = f(x_1, x_2, ..., x_n)$ is privately computable. Furthermore, a lower bound for the round complexity of computing f is given by the maximum number of rounds required to compute $f_2(x_i, (x_1, x_2, ..., x_{i-1}, x_{i+1}, x_{i+2}, ..., x_n))$ for any i.

By combining Lemma 3.7 and Lemma 3.8, 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 4.1, 4.6, 4.9, and 5.2).

LEMMA 3.9. Let \vec{x} and \vec{y} be vectors of n-1 bids and x and y single bids. It is impossible to privately emulate an auction if (\vec{x}, x) , (\vec{x}, y) , and (\vec{y}, x) all yield outcome a and (\vec{y}, y) does not yield a.

If the antecedent of the previous lemma is satisfied, \vec{x}, \vec{y}, x, y are called an "embedded OR" because the corresponding 2×2 submatrix resembles the Boolean OR function.

Due to the 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 4.3 and 5.2). As first observed by Benaloh [1987], there is a simple protocol to privately compute modular sums based on the homomorphicity of secret sharing.

LEMMA 3.10. [Benaloh 1987] $f(x_1, x_2, \ldots, x_n) = \sum_{i=1}^n x_i \mod p$ is privately computable.

PROOF. Each agent *i* chooses *n* random values $x_{ij} \in \mathbb{Z}_p$ so that $\sum_{j=1}^n x_{ij} \mod p = x_i$. He then sends each addend x_{ij} to agent *j* and keeps x_{ii} . After all agents have done this, each agent *i* publishes $s_i = \sum_{j=1}^n x_{ji} \mod p$, *i.e.*, the modular sum of his remaining x_{ii} and the n-1 addends he received. The function value $f(x_1, x_2, \ldots, x_n) = \sum_{i=1}^n s_i \mod p$ can be computed by each participant. It is easily verified that the described protocol indeed satisfies privacy. \Box

In contrast to the 2-party case, randomization is *necessary* to privately compute any non-degenerate function with n parties.

⁶This is a special case of the Partition Lemma as defined by Chor et al. [1994] for t = n - 1.

4. PRIVACY-PRESERVING AUCTION PROTOCOLS

Before looking at the emulation of first- and second-price auctions, we will give an introductory positive result and a general negative result. First, we propose a simple auction protocol that does not meet all our desiderata but raises hope that appropriate protocols might indeed exist. Then we show the impossibility of only determining the auction winner while preserving privacy.

It is very tempting to construct an auction protocol based on the homomorphicity of secret sharing (Lemma 3.10). For example, consider the following first-price auction scheme based on the unary representation of bids.⁷ Let v be the number of bits necessary to represent a bid (in binary) and a be the index of the bidder executing the protocol steps based on his bid b_a .

(1) Choose Y_{aj} for each $j \in \{1, 2, ..., 2^v\}$ and b_{aj}^{+i} for each j and i so that

$$\sum_{i=1}^{n} b_{aj}^{+i} = \begin{cases} Y_{aj} & \text{if } j \le b_a \\ 0 & \text{otherwise.} \end{cases}$$

- (2) Send b_{aj}^{+i} for each j to bidder i for each $i \neq a$.
- (3) After having received all b_{ij}^{+a} , publish $b_j^{+a} = \sum_{i=1}^n b_{ij}^{+a}$ for each j.
- (4) Compute $B_j = \sum_{a=1}^n b_j^{+a}$ for each j by using the published b_j^{+a} .
- (5) If $B_j = Y_{aj}$ for any j, then bidder a won the auction. The selling price, $\max\{j \mid B_j \neq 0\}$, is visible to all bidders.

The Appendix contains an example run of this protocol. While the protocol is unconditionally private (obviously no intractability assumptions are being made), it is *not* fully private. The winning bidder *i* can identify the second highest bid by looking for the greatest *j* so that $B_j \neq Y_{ij}$. More generally, the *k* highest bidders can jointly reveal the bid of the (k + 1)st-highest bidder. Nevertheless, the protocol raises hope that unconditionally fully private auction protocols actually exist. The proposed protocol is particularly interesting because it requires just two rounds of interaction.

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 most types of auctions. Perhaps surprisingly, this rather simple function cannot be computed privately.

THEOREM 4.1. The arg max function cannot be computed privately.

PROOF. Let $f_{\arg \max}(x_1, x_2, \ldots, x_n) = \arg \max_{i \in \{1, 2, \ldots, n\}}(x_i)$ be a function that yields the index of the greatest argument. Furthermore, let

$$\vec{x} = (2, \underbrace{1, \dots, 1}_{n-2}), \ \vec{y} = (4, \underbrace{1, \dots, 1}_{n-2}), \ x = 5, \ \text{and} \ y = 3.$$

Table II	Table III. Embedded OR in $f_{\arg \max}$				
	$f_{ m argmax}$	3	5	_	
	$2, 1, \dots, 1$	n	n		
	$4, 1, \ldots, 1$	1	n		

Then,

$$f_{rg \max}(\vec{x}, x) = f_{rg \max}(\vec{x}, y) = f_{rg \max}(\vec{y}, x) = n.$$

However, $f_{\arg \max}(\vec{y}, y) = 1$ (see Table III). It follows from Lemma 3.9 that $f_{\arg \max}$ is *not* privately computable. \Box

4.1 First-Price Auctions

J

Even though the arg max function cannot be computed privately, it turns out that it is possible to privately compute a function that reveals the winner while at the same time revealing the highest bid. Let $\vec{b} = (b_1, b_2, \ldots, b_n)$ be the vector of submitted bids, $\max(\vec{b}) = \max_{i \in \{1,2,\ldots,n\}}(b_i)$, and $\arg\max(\vec{b}) = \arg\max_{i \in \{1,2,\ldots,n\}}(b_i)$.

Definition 4.2. The first-price auction's public outcome is defined by the function

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

When examining f^1 for just two bidders (Table IV), it turns out that, in contrast to the proof of Theorem 4.1, the Corners Lemma is not applicable (when disregarding ties). Bold numbers denote the winner's index.

Table IV. f^1 for two bidders						
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)		
:						

Clearly, 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, according to the complete characterization of privately computable functions for *two* agents given in Theorem 3.5, $f^1(b_1, b_2)$ is indeed privately computable. Matrix M_{f^1} can be decomposed by alternatingly cutting off the columns and rows to the far right and bottom, respectively (see Table IV). Moreover, it can be shown that f^1 is privately computable for *any* number of agents.

THEOREM 4.3. The first-price auction can be emulated by a 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 (in binary). There is no more efficient private protocol for this task.

⁷This protocol was introduced by Brandt [2002], based on an idea by Kikuchi et al. [1998].

PROOF. Consider the following protocol:

- (1) Let $j = 2^v$.
- (2) Each agent broadcasts either 1 or 0 depending on whether he is willing to pay price j or not.⁸
- (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 win(s) the auction.

This protocol strongly resembles the Dutch (or descending) auction in which an auctioneer gradually (or continuously) lowers the selling price until the first bidder expresses his willingness to buy (see, *e.g.*, [Krishna 2002]).⁹ Clearly, no information beyond the auction outcome is revealed, even when bidders collude.

Lemma 3.8 gives a lower bound for the number of rounds needed to compute f^1 by looking at the minimal decomposition tree of the 2-ary function $f^1(b_1, b_2)$. This decomposition tree has depth $2 \cdot (2^v - 1)$ (if we allow for the simultaneous broadcasting of messages, the depth is $2^v - 1$) which implies exponential round complexity for computing f^1 (in fact, $2^v - 1$ rounds) and thus the optimality of the proposed protocol. \Box

As mentioned in Section 3, unconditionally private protocols generally require a complete network of private channels. An outstanding property of the protocol proposed above 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 and 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.

It might seem that the outcome function given in Definition 4.2 reveals the minimal amount of information required to perform the required transaction (*e.g.*, selling a good). However, the notion of minimal revelation can be refined even further by moving to *asymmetric* information revelation since it is unnecessary that losing bidders learn who won the auction and which price the winner has to pay. On this account, 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.

⁸Different 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 (small) amount of additional information to be revealed (if there are more than two bidders). As stated in Section 1, we disregard ties in this paper.

⁹The equivalence of first-price and Dutch auctions was first observed by Vickrey [1961].

¹⁰The availability of a broadcast channel is crucial since Byzantine agreement [Lamport et al. 1982] is not feasible in our setting as it either requires intractability assumptions or the trustworthiness of two thirds of the agents.

ACM Journal Name, Vol. V, No. N, Month 20YY.

Definition 4.4. The first-price auction's private outcome function is defined as

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

In practical applications (where security against active adversaries is required), it might be desirable to include the seller in the protocol and compute all f_i^1 functions so that only bidder *i* and the seller learn the outcome. This prevents a bidder from aborting the protocol when he is unsatisfied with the auction outcome, leaving the seller uninformed (see [Brandt 2003; Brandt and Sandholm 2005]).

With the notable exception of work by Beimel et al. [1999] who 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. Results from that setting cannot be directly transferred to a setting where only *one* agent learns the function value. However, the following simple lemma is sufficient to show the impossibility of private computation in the latter case.

LEMMA 4.5. If a function $f(x_1, x_2, ..., x_n)$ cannot be privately computed so that all agents learn the function value, it cannot be computed so that only a single agent (or any subset of agents) learns the function value.

PROOF. The statement can easily be shown using an indirect argument. 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. \Box

We are now ready to show the impossibility of privately computing the private outcome function of a first-price auction.

THEOREM 4.6. There is no private protocol that computes the private outcome $f_i^1(\vec{b})$ of a first-price auction.

PROOF. With the help of Lemma 4.5, we can prove the impossibility of computing f_i^1 by using a chain of necessary conditions. It suffices to use Lemma 3.9 to show the impossibility of a private protocol for any n. Let

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

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, it turns out that $f_n^1(\vec{y}, y) = 3$ (see Table V). Thus, according to

Table V	. Embed	lded	OR in f_i	1
	f_n^1	1	3	
2,	$1,\ldots,1$	0	3	
4,	$1,\ldots,1$	0	0	

Lemma 3.9, f_i^1 is not privately computable for any *i*. Lemma 4.5 implies that there is no protocol to compute f_i^1 privately so that only bidder *i* learns the outcome. \Box

4.2 Second-Price Auctions

In this section, we investigate the existence of private protocols that emulate secondprice sealed-bid auctions. Recall that in second-price auctions, the bidder who submitted the highest bid wins the auction and is required to pay the amount of the *second*-highest bid.

 $Definition\,$ 4.7. The second-price auction's public outcome is defined by the function 11

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

PROPOSITION 4.8. There is a private protocol that emulates the second-price auction for two bidders.

PROOF. When there are just two bidders, the Dutch auction style protocol proposed in the proof of Theorem 4.3 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 the same information can always be inferred from the outcome if there are only two bidders. The highest bid remains private. \Box

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

THEOREM 4.9. There is no private protocol that emulates the second-price auction for more than two bidders.

PROOF. The following counter-example shows the impossibility of privately computing f^2 for n > 2. Let

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

Table VI. Em	bedded	OR in f^2	
f^2	1	2	
$3, 1, 1, \ldots, 1$	(1, 1)	(2, 1)	
$3, 2, 1, \ldots, 1$	(2, 1)	(2, 1)	

Then,

$$f^{2}(\vec{x}, x) = f^{2}(\vec{x}, y) = f^{2}(\vec{y}, x) = (2, 1)$$

(*i.e.*, bidder 1 wins the auction at price 2). However, $f^2(\vec{y}, y) = (1, 1)$ (*i.e.*, bidder 1 wins at price 1, see Table VI). It follows from Lemma 3.9 that f^2 is not privately computable. \Box

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.

 $^{{}^{11}\}vec{b}_{-i}$ denotes vector \vec{b} without component *i*.

ACM Journal Name, Vol. V, No. N, Month 20YY.

This effect is enhanced in Section 5.1 where it is shown that even the search for a second-price auction protocol that hides a limited amount of information is futile.

Clearly, the impossibility of computing the second-price public outcome function also implies the impossibility of computing the private outcome function where only the auction winner learns the auction result.

Definition 4.10. The second-price 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 4.11. There is no protocol that privately computes the private outcome of a second-price auction.

PROOF. If f_i^2 could be computed privately, then f^2 would be privately computable as well by letting the winning bidder *i* broadcast f_i^2 and his identity. \Box

5. SECURITY MODEL RELAXATIONS

Given the impossibility result of Theorem 4.9, a natural follow-up problem is to investigate how far this impossibility reaches, *i.e.*, to investigate whether the impossibility still holds under slightly loosened assumptions. By following this path, we either obtain feasibility results for settings that are "almost" as strict as the original setting, or negative results that highlight the robustness of the impossibility. In this section, we will see that the impossibility of Theorem 4.9 turns out to be rather robust. Three security model relaxations that we consider in this paper are:

- —allowing partial revelation of bids, e.g., the highest bid, by modifying the outcome function,
- —allowing coalitions of bidders to uncover information (*i.e.*, relaxing full privacy), and
- -guaranteeing high probability of correctness instead of correctness for sure.

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

5.1 Partial Revelation

The more information an outcome function reveals about the bids, the more likely becomes the existence of a private protocol for computing it. In this section, we study whether the revelation of a limited amount of information enables the private computation of second-price auctions. A substantial yet reasonable weakening of privacy is *anonymity*, where we restrict indistinguishability to input vectors that are permutations of one another.

Definition 5.1. A protocol for computing f is anonymous if for all $T \subseteq N$, any pair of input vectors $\vec{x}, \vec{y} \in X^n$ so that \vec{x} is a permutation of $\vec{y}, \forall i \in T : x_i = y_i$, and $f(\vec{x}) = f(\vec{y})$, and any choice of random inputs $\{r_i\}_{i \in T}$

$$\langle \operatorname{VIEW}_T(\vec{x}, \{r_i\}_{i \in T}) \mid \{r_i\}_{i \in \bar{T}} \rangle = \langle \operatorname{VIEW}_T(\vec{y}, \{r_i\}_{i \in T}) \mid \{r_i\}_{i \in \bar{T}} \rangle.$$

Loosely speaking, 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 when allowing the revelation of all but one losing bid. 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 5.2. There is no anonymous second-price auction protocol that only hides a single losing bid. There exists an anonymous second-price auction protocol that does not reveal the winning bid.

PROOF. In an anonymous auction, the bids can only be distinguished by their numerical order. Assume, for a contradiction, 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 3.9. Let

$$\vec{x} = (n, n - 1, \dots, n - k + 2, n - k, n - k - 1, \dots, 1),$$

$$\vec{y} = (n, n - 1, \dots, n - k + 3, n - k + 1, n - k, \dots, 1),$$

$$x = n - k + 2, \text{ and}$$

$$y = n - k + 1.$$

Then,

$$g_k(\vec{x}, x) = g_k(\vec{x}, y) = g_k(\vec{y}, x) = (\vec{x}, \mathbf{1}).$$

However, $g_k(\vec{y}, y) = (\vec{y}, \mathbf{1})$ which proves the impossibility of privately computing g_k according to Lemma 3.9. Table VII shows an example for four bidders when only the third highest bid should be kept private (n = 4, k = 3).

Tal	ole VII.	Embedd	ed OR in g_3	
9	13	2	3	
4, 5	2,1 (4	(4, 2, 1, 1)	(4, 3, 1, 1)	
4, 3	3,1 (4	4, 3, 1, 1)	(4, 3, 1, 1)	

It remains to be shown that it is indeed *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 (but not the corresponding bidders' identities). Interestingly, this task can be accomplished by a protocol similar to an anonymized English (*i.e.*, ascending) auction.¹²

(1) Let j = 1.

 $^{^{12}}$ As mentioned in Section 1, standard cryptographic techniques have to be employed to ensure that agents follow the protocol truthfully and do not manipulate, for example, by wrongfully broadcasting their identity in Step (5).

ACM Journal Name, Vol. V, No. N, Month 20YY.

- (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 $s = \sum_{i=1}^{n} x_i \mod (n+1)$ according to the protocol defined in Lemma 3.10.
- (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.

It is easily verified that this protocol satisfies privacy. \Box

The constructions used in the proofs of Theorem 4.9 and Theorem 5.2 rely on some bidders submitting identical bids in order to yield an embedded OR (nevertheless, the winning bid in these constructions is always unique). If we somehow prohibit identical bids (e.g., by letting each bidder *i* bid $b_i \cdot n + i$ instead of b_i), we essentially reveal the identity of the second highest bidder and avoid the occurrence of embedded ORs. Thus, it seems as if moving away from anonymity by disclosing the identity of the second highest bidder (in addition to the second-price auction outcome) might enable the private emulation of second-price auctions. However, the following theorem shows that this is not the case.

THEOREM 5.3. There is no private second-price auction protocol that only reveals the identity of the second highest bidder, in addition to the auction outcome.

PROOF. Let

 $h(\vec{b}) = ((\max(\vec{b}_{-\arg\max(\vec{b})}, \arg\max(\vec{b}_{-\arg\max(\vec{b})})), \arg\max(\vec{b}))$

be the modified outcome function that also reveals the identity of the second highest bidder. It turns out that this function contains no embedded OR. Nevertheless, h's matrix is not decomposable, which can be seen by looking at the submatrix given in Table VIII. The first pair of numbers represents the selling price and the identity of the bidder who submitted the second highest bid. The last number is the identity of the winning bidder.

Table VIII. M_h is non-decomposable					
h	1	4	7		
$2, 3, 1, \ldots, 1$	((2,1),2)	((3, 2), n)	((3, 2), n)		
$2, 6, 1, \ldots, 1$	((2, 1), 2)	((4, n), 2)	$((6, 2), \mathbf{n})$		
$5, 6, 1, \ldots, 1$	((5, 1), 2)	((5, 1), 2)	$((6, 2), \mathbf{n})$		

In fact, the matrix given in Table VIII is isomorphic to the matrix given in Table II as an example of a non-decomposable matrix. It follows from Theorem 3.5 and Lemma 3.8 that h is not privately computable. \Box

Theorem 5.3 significantly strengthens the impossibility result of Theorem 4.9 because its proof does not rely on bidders submitting identical bids.

5.2 Uncovering by Coalitions

So far, we required that no coalition consisting of up to n-1 agents should be able to uncover private information (full privacy). However, there are functions that cannot be computed fully privately but can be computed privately when only allowing curious coalitions of size n-2 [Chor et al. 1994]. In this section, we examine whether loosening full privacy enables the private emulation of second-price auctions.

The proofs of Theorem 4.9 and Theorem 5.3 can easily be modified to be used with a version of the Partition Lemma (Lemma 3.8) where the inputs are partitioned into two sets of equal size rather than into a set of size n-1 and a singleton, because both proofs only rely on a *constant* number of relevant bids (no matter how large n is). This implies the impossibility of privately emulating second-price auctions even when only up to $\lceil \frac{n}{2} \rceil$ bidders are allowed to collude. This bound is tight because, as mentioned in Section 1, assuming that a strict majority of participants (*i.e.*, more than $\frac{n}{2}$ bidders) is trustworthy allows the private computation of *any* function (including f^2) [Ben-Or et al. 1988; Chaum et al. 1988].

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 $\epsilon < \frac{1}{2}$ does not enable the private computation of functions that cannot be computed with perfect correctness in (i) the Boolean *n*-party case [Chor and Kushilevitz 1989] and (ii) the general 2-party case [Kushilevitz 1989]. 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. We leave this as an open question for future research.

6. CONCLUSIONS AND FUTURE RESEARCH

Sealed-bid auctions are not only widely used for the selling of goods, they have also been applied to task assignment, scheduling, or finding the shortest path in a network with selfish nodes. Bid privacy is of increasing importance in such auctions, and various schemes that distribute trust onto several auctioneers have been proposed recently. In contrast to existing work, this paper dealt with *unconditional full privacy*, *i.e.*, privacy that relies neither on trusted third parties (like auctioneers), 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 without revealing any additional information beyond the outcome. We derived several impossibility and possibility results in this domain (see Table IX).

The first-price auction can be emulated by a private protocol. 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 privately.

There is a 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 hiding a single losing bid (but maintaining anonymity),
- -revealing the identity of the second highest bidder, or
- -tolerating the revelation of complete information to any coalition consisting of at least half of the bidders.

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

Т	able IX. Existence of Private Auction F	Protocols
	Public Outcome	Private Outcome
First-Price	YES (exponential number of rounds)	NO
Second-Price	NO (unless $n = 2$)	NO (unless $n = 2$)

Future work includes the design of auction protocols that only reveal partial information on each bid such as "the lowest bid is greater than 10". Theorem 5.2 states that some information on *all* losing bids has to be revealed. It would be interesting to minimize this amount of information for practical instances. So far, theoretic results on minimum revelation protocols are only known for two parties [Bar-Yehuda et al. 1990].

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 he has enough information to determine the auction winner (see, *e.g.*, [Sandholm and Boutilier 2006]). This approach also provides partial unconditional privacy and it might be possible to transfer results from one setting to the other.

APPENDIX

This section contains an example run of the protocol given in Section 4. All computations take place in the additive group \mathbb{Z}_{11} . There are three bidders (n = 3)and two bits for each bid (v = 2) resulting in four possible valuations. Let $b_1 = 1$, $b_2 = 3$, and $b_3 = 1$. Each bidder chooses vector $\vec{y}_i = (Y_{i1}, Y_{i2}, \ldots, Y_{ik})$ at random. Let $\vec{y}_1 = (4, 10, 3, 5), \ \vec{y}_2 = (8, 1, 5, 9), \ \text{and} \ \vec{y}_3 = (2, 8, 10, 7)$. Then, each bidder generates his bid vector $\vec{b}_i = (b_{i1}, b_{i2}, \ldots, b_{ik})$ according to \vec{y} and b_i , and creates a

¹³Much better results can be obtained when relying on computational intractability. For example, there are computationally fully private first-price and second-price auction protocols that only require a constant number of rounds [Brandt 2006].

3-partition of \vec{b}_i .

$$\vec{b}_{1} = \begin{pmatrix} 0\\0\\0\\4 \end{pmatrix} = \vec{b}_{1}^{+1} + \vec{b}_{1}^{+2} + \vec{b}_{1}^{+3} = \begin{pmatrix} 5\\1\\7\\8 \end{pmatrix} + \begin{pmatrix} 2\\8\\6\\0 \end{pmatrix} + \begin{pmatrix} 4\\2\\9\\7 \end{pmatrix}$$
$$\vec{b}_{2} = \begin{pmatrix} 0\\5\\1\\8 \end{pmatrix} = \vec{b}_{2}^{+1} + \vec{b}_{2}^{+2} + \vec{b}_{2}^{+3} = \begin{pmatrix} 9\\1\\2\\6 \end{pmatrix} + \begin{pmatrix} 3\\7\\5\\3 \end{pmatrix} + \begin{pmatrix} 10\\8\\5\\10 \end{pmatrix}$$
$$\vec{b}_{3} = \begin{pmatrix} 0\\0\\2 \end{pmatrix} = \vec{b}_{3}^{+1} + \vec{b}_{3}^{+2} + \vec{b}_{3}^{+3} = \begin{pmatrix} 4\\10\\0\\3 \end{pmatrix} + \begin{pmatrix} 6\\8\\2\\2 \end{pmatrix} + \begin{pmatrix} 1\\4\\9\\8 \end{pmatrix}$$

Bidder 1 keeps b_1^{+1} , sends b_1^{+2} to bidder 2 and b_1^{+3} to bidder 3. Bidder 2 and 3 do likewise. After that, each bidder sums up the two shares he received and his own remaining share, and publishes the resulting vector.

$$\vec{b}^{+1} = \sum_{i=1}^{n} \vec{b}^{+1}_{i} = \begin{pmatrix} 7\\1\\9\\6 \end{pmatrix}, \ \vec{b}^{+2} = \sum_{i=1}^{n} \vec{b}^{+2}_{i} = \begin{pmatrix} 0\\1\\2\\5 \end{pmatrix}, \ \vec{b}^{+3} = \sum_{i=1}^{n} \vec{b}^{+3}_{i} = \begin{pmatrix} 4\\3\\1\\3 \end{pmatrix}$$

All bidders can derive the result by summing up the published vectors.

$$\vec{B} = \sum_{i=1}^{n} \vec{b}_i = \sum_{i=1}^{n} \vec{b}^{+i} = \begin{pmatrix} 0\\5\\1\\3 \end{pmatrix}$$

The selling price, 3, (the lowest price at which nobody bid) is visible to all bidders. Bidder 2 can tell that he won the auction because the second and the third component of \vec{B} are equal to the corresponding components of \vec{y}_2 .¹⁴ The two losing bidders cannot identify the winner (without colluding) or reveal each other's bid.

ACKNOWLEDGMENTS

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, ITR IIS-0121678, and ITR IIS-0427858 and a Sloan Fellowship.

REFERENCES

ABE, M. AND SUZUKI, K. 2002. M+1-st price auction using homomorphic encryption. In Proceedings of the 5th International Conference on Public Key Cryptography (PKC). Lecture Notes in Computer Science (LNCS), vol. 2274. Springer-Verlag, 115–224.

 $^{^{14}}$ This only works when there is no tie for the highest bid. There are various possibilities to proceed in the case of a tie.

ACM Journal Name, Vol. V, No. N, Month 20YY.

- BAR-YEHUDA, R., CHOR, B., AND KUSHILEVITZ, E. 1990. Privacy, additional information, and communication. In Proceedings of the 5th IEEE Conference on Structure in Complexity Theory. 55–65.
- BAUDRON, O. AND STERN, J. 2001. Non-interactive private auctions. In Proceedings of the 5th Annual Conference on Financial Cryptography (FC). Lecture Notes in Computer Science (LNCS), vol. 2339. Springer-Verlag, 300–313.
- BEIMEL, A., MALKIN, T., AND MICALI, S. 1999. The all-or-nothing nature of two-party secure computation. In Advances in Cryptology - Proceedings of the 19th Annual International Cryptology Conference (CRYPTO). Lecture Notes in Computer Science (LNCS), vol. 1666. Springer-Verlag, 80–97.
- BEN-OR, M., GOLDWASSER, S., AND WIGDERSON, A. 1988. Completeness theorems for noncryptographic fault-tolerant distributed computation. In *Proceedings of the 20th Annual ACM* Symposium on the Theory of Computing (STOC). ACM Press, 1–10.
- BENALOH, J. 1987. Secret sharing homomorphisms: Keeping shares of a secret secret. In Advances in Cryptology Proceedings of the 13th Annual International Cryptology Conference (CRYPTO). Lecture Notes in Computer Science (LNCS), vol. 263. Springer-Verlag, 251–260.
- BRANDT, F. 2002. Secure and private auctions without auctioneers. Tech. Rep. FKI-245-02, Department for Computer Science, Technical University of Munich. ISSN 0941-6358.
- BRANDT, F. 2003. Fully private auctions in a constant number of rounds. In Proceedings of the 7th Annual Conference on Financial Cryptography (FC), R. N. Wright, Ed. Lecture Notes in Computer Science (LNCS), vol. 2742. Springer-Verlag, 223–238.
- BRANDT, F. 2006. How to obtain full privacy in auctions. International Journal of Information Security 5, 4, 201–216.
- BRANDT, F. AND SANDHOLM, T. 2005. Efficient privacy-preserving protocols for multi-unit auctions. In Proceedings of the 9th International Conference on Financial Cryptography and Data Security (FC), A. Patrick and M. Yung, Eds. Lecture Notes in Computer Science (LNCS), vol. 3570. Springer-Verlag, 298–312.
- CACHIN, C. 1999. Efficient private bidding and auctions with an oblivious third party. In Proceedings of the 6th ACM Conference on Computer and Communications Security. ACM Press, 120–127.
- CHAUM, D., CRÉPEAU, C., AND DAMGÅRD, I. 1988. Multi-party unconditionally secure protocols. In Proceedings of the 20th Annual ACM Symposium on the Theory of Computing (STOC). ACM Press, 11–19.
- CHOR, B., GERÉB-GRAUS, M., AND KUSHILEVITZ, E. 1994. On the structure of the privacy hierarchy. Journal of Cryptology 7, 1, 53–60.
- CHOR, B. AND KUSHILEVITZ, E. 1989. A zero-one law for Boolean privacy. In *Proceedings of the* 21st Annual ACM Symposium on the Theory of Computing (STOC). ACM Press, 62–72.
- FRANKLIN, M. K. AND REITER, M. K. 1996. The design and implementation of a secure auction service. *IEEE Transactions on Software Engineering* 22, 5, 302–312.
- GOLDREICH, O. 2001. Foundations of Cryptography: Volume 1, Basic Tools. Cambridge University Press.
- GOLDREICH, O., MICALI, S., AND WIGDERSON, A. 1987. 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)*. ACM Press, 218–229.
- HARKAVY, M., TYGAR, J. D., AND KIKUCHI, H. 1998. Electronic auctions with private bids. In Proceedings of the 3rd USENIX Workshop on Electronic Commerce. 61–74.
- JUELS, A. AND SZYDLO, M. 2002. A two-server, sealed-bid auction protocol. In Proceedings of the 6th Annual Conference on Financial Cryptography (FC), M. Blaze, Ed. Lecture Notes in Computer Science (LNCS), vol. 2357. Springer-Verlag, 72–86.
- KIKUCHI, H. 2001. (M+1)st-price auction protocol. In Proceedings of the 5th Annual Conference on Financial Cryptography (FC). Lecture Notes in Computer Science (LNCS), vol. 2339. Springer-Verlag, 351–363.

- KIKUCHI, H., HARKAVY, M., AND TYGAR, J. D. 1998. Multi-round anonymous auction protocols. In Proceedings of the 1st IEEE Workshop on Dependable and Real-Time E-Commerce Systems. 62–69.
- KRISHNA, V. 2002. Auction Theory. Academic Press.
- KUSHILEVITZ, E. 1989. Privacy and communication complexity. In Proceedings of the 30th Symposium on Foundations of Computer Science (FOCS). IEEE Computer Society Press, 416–421.
- LAMPORT, L., SHOSTAK, R., AND PEASE, M. 1982. The Byzantine generals problem. ACM Transactions on Programming Languages and Systems 4, 3, 382–401.
- LIPMAA, H., ASOKAN, N., AND NIEMI, V. 2002. Secure Vickrey auctions without threshold trust. In Proceedings of the 6th Annual Conference on Financial Cryptography (FC), M. Blaze, Ed. Lecture Notes in Computer Science (LNCS), vol. 2357. Springer-Verlag, 87–101.
- LÓPEZ, N., NÚÑEZ, M., RODRÍGUEZ, I., AND RUBIO, F. 2004. Improving privacy in Vickrey auctions. ACM SIGEcom Exchanges 5, 1, 1–12.
- NAOR, M., PINKAS, B., AND SUMNER, R. 1999. Privacy preserving auctions and mechanism design. In Proceedings of the 1st ACM Conference on Electronic Commerce (ACM-EC). ACM Press, 129–139.
- NURMI, H. AND SALOMAA, A. 1993. Cryptographic protocols for Vickrey auctions. Group Decision and Negotiation 2, 363–373.
- PENG, K., BOYD, C., DAWSON, E., AND VISWANATHAN, K. 2002. Non-interactive auction scheme with strong privacy. In *Proceedings of the 5th International Conference on Information Security* and Cryptology (ICISC). Lecture Notes in Computer Science (LNCS), vol. 2587. Springer-Verlag, 407–420.
- PORTER, R. AND SHOHAM, Y. 2005. On cheating in sealed-bid auctions. Decision Support Systems 39, 1, 41–54.
- RODRÍGUEZ, I. AND LÓPEZ, N. 2006. Analyzing the privacy of a Vickrey auction mechanism. International Journal of E-Business Research 2, 3, 17–27.
- ROTHKOPF, M. H. AND HARSTAD, R. M. 1995. Two models of bid-taker cheating in Vickrey auctions. *Journal of Business* 68, 2, 257–267.
- ROTHKOPF, M. H., TEISBERG, T. J., AND KAHN, E. P. 1990. Why are Vickrey auctions rare? Journal of Political Economy 98, 1, 94–109.
- SAKO, K. 2000. An auction protocol which hides bids of losers. In Proceedings of the 3rd International Conference on Public Key Cryptography (PKC). Lecture Notes in Computer Science (LNCS), vol. 1751. Springer-Verlag, 422–432.
- SANDHOLM, T. 2000. Issues in computational Vickrey auctions. International Journal of Electronic Commerce, Special issue on Intelligent Agents for Electronic Commerce 4, 3, 107–129.
- SANDHOLM, T. AND BOUTILIER, C. 2006. Preference elicitation in combinatorial auctions. In Combinatorial Auctions, P. Cramton, Y. Shoham, and R. Steinberg, Eds. MIT Press, Chapter 10, 233–263.
- VICKREY, W. 1961. Counter speculation, auctions, and competitive sealed tenders. Journal of Finance 16, 1, 8–37.
- YAO, A. C. 1979. Some complexity questions related to distributed computing. In Proceedings of the 11th Annual ACM Symposium on the Theory of Computing (STOC). ACM Press, 209–213.