

**DIMACS Technical Report 2009-14**  
**Version of 2009-05-16**

# Approximate Privacy: Foundations and Quantification

by

Joan Feigenbaum<sup>1,2</sup>  
Department of Computer Science  
Yale University  
New Haven, CT 06520, USA  
joan.feigenbaum@yale.edu

Aaron D. Jaggard<sup>3,4</sup>  
DIMACS  
Rutgers University  
Piscataway, NJ 08854, USA  
adj@dimacs.rutgers.edu

Michael Schapira<sup>5</sup>  
Department of Computer Science  
Yale University  
New Haven, CT 06520, USA  
and  
Computer Science Division  
University of California  
Berkeley, CA 94720, USA  
michael.schapira@yale.edu

<sup>1</sup>DIMACS Permanent Member

<sup>2</sup>Supported in part by NSF grants 0331548 and 0534052 and IARPA grant FA8750-07-0031.

<sup>3</sup>DIMACS Assistant Research Professor

<sup>4</sup>Supported in part by NSF grants 0751674 and 0753492.

<sup>5</sup>Supported by NSF grant 0331548.

---

DIMACS is a collaborative project of Rutgers University, Princeton University, AT&T Labs–Research, Bell Labs, NEC Laboratories America and Telcordia Technologies, as well as affiliate members Avaya Labs, HP Labs, IBM Research, Microsoft Research, Stevens Institute of Technology, Georgia Institute of Technology, Rensselaer Polytechnic Institute and The Cancer Institute of New Jersey. DIMACS was founded as an NSF Science and Technology Center.

## ABSTRACT

Increasing use of computers and networks in business, government, recreation, and almost all aspects of daily life has led to a proliferation of online sensitive data about individuals and organizations. Consequently, concern about the *privacy* of these data has become a top priority, particularly those data that are created and used in electronic commerce. There have been many formulations of privacy and, unfortunately, many negative results about the feasibility of maintaining privacy of sensitive data in realistic networked environments. We formulate communication-complexity-based definitions, both worst-case and average-case, of a problem's *privacy-approximation ratio*. We use our definitions to investigate the extent to which approximate privacy is achievable in two standard problems: the  $2^{\text{nd}}$ -price *Vickrey auction* [18] and the *millionaires problem* of Yao [20].

For both the  $2^{\text{nd}}$ -price Vickrey auction and the millionaires problem, we show that not only is perfect privacy impossible or infeasibly costly to achieve, but even *close approximations* of perfect privacy suffer from the same lower bounds. By contrast, we show that, if the values of the parties are drawn uniformly at random from  $\{0, \dots, 2^k - 1\}$ , then, for both problems, simple and natural communication protocols have privacy-approximation ratios that are linear in  $k$  (*i.e.*, logarithmic in the size of the space of possible inputs). We conjecture that this improved privacy-approximation ratio is achievable for *any* probability distribution.

# 1 Introduction

Increasing use of computers and networks in business, government, recreation, and almost all aspects of daily life has led to a proliferation of online sensitive data about individuals and organizations. Consequently, the study of privacy has become a top priority in many disciplines. Computer scientists have contributed many formulations of the notion of *privacy-preserving computation* that have opened new avenues of investigation and shed new light on some well studied problems.

One good example of a new avenue of investigation opened by concern about privacy can be found in auction design, which was our original motivation for this work. Traditional auction theory is a central research area in Economics, and one of its main questions is how to incent bidders to behave truthfully, *i.e.*, to reveal private information that auctioneers need in order to compute optimal outcomes. More recently, attention has turned to the complementary goal of enabling bidders *not* to reveal private information that auctioneers do *not* need in order to compute optimal outcomes. The importance of bidders' privacy, like that of algorithmic efficiency, has become clear now that many auctions are conducted online, and Computer Science has become at least as relevant as Economics.

Our approach to privacy is based on communication complexity. Although originally motivated by agents' privacy in mechanism design, our definitions and tools can be applied to distributed function computation in general. Because perfect privacy can be impossible or infeasibly costly to achieve, we investigate approximate privacy. Specifically, we formulate both worst-case and average-case versions of the *privacy-approximation ratio* of a function  $f$  in order to quantify the amount of privacy that can be maintained by parties who supply sensitive inputs to a distributed computation of  $f$ . We also study the tradeoff between privacy preservation and communication complexity.

Our points of departure are the work of Chor and Kushilevitz [6] on characterization of privately computable functions and that of Kushilevitz [15] on the communication complexity of private computation. Starting from the same place, Bar-Yehuda *et al.* [1] also provided a framework in which to quantify the amount of privacy that can be maintained in the computation of a function and the communication cost of achieving it. Their definitions and results are different from the ones we present here; as explained in Section 5 below, a precise characterization of the relationship between their formulation and ours is an interesting direction for future work.

## 1.1 Our Approach

Consider an auction of a Bluetooth headset with 2 bidders, 1 and 2, in which the auctioneer accepts bids ranging from \$0 to \$7 in \$1 increments. Each bidder  $i$  has a private value  $x_i \in \{0, \dots, 7\}$  that is the maximum he is willing to pay for the headset. The item is sold in a 2<sup>nd</sup>-price Vickrey auction, *i.e.*, the *highest* bidder gets the item (with ties broken in favor of bidder 1), and the price he pays is the *second highest* bid. The demand for privacy arises naturally in such scenarios [17]: In a straightforward protocol, the auctioneer receives sealed bids from both bidders and computes the outcome based on this information. Say, *e.g.*, that

bidder 1 bids \$3, and bidder 2 bids \$6. The auctioneer sells the headset to bidder 2 for \$3. It would not be at all surprising however if, in subsequent auctions of headsets in which bidder 2 participates, the same auctioneer set a reservation price of \$5. This could be avoided if the auction protocol allowed the auctioneer to learn the fact that bidder 2 was the highest bidder (something he needs to know in order to determine the outcome) but did not entail the full revelation of 2’s private value for the headset.




Figure 1: The minimal knowledge requirements for 2<sup>nd</sup>-price auctions

Observe that, in some cases, revelation of the *exact* private information of the highest bidder is *necessary*. For example, if  $x_1 = 6$ , then bidder 2 will win only if  $x_2 = 7$ . In other cases, the revelation of *a lot* of information is necessary, *e.g.*, if bidder 1’s bid is 5, and bidder 2 outbids him, then  $x_2$  must be either 6 or 7. An auction protocol is said to achieve *perfect objective privacy* if the auctioneer learns nothing about the private information of the bidders that is not needed in order to compute the result of the auction. Figure 1 describes the information the auctioneer *must* learn in order to determine the outcome of a 2<sup>nd</sup>-price auction. Observe that the auctioneer’s failure to distinguish between two potential pairs of inputs that belong to *different* rectangles in Fig. 1 implies his inability to determine the winner or the price the winner must pay. Also observe, however, that the auctioneer need not be able to distinguish between two pairs of inputs that belong to *the same* rectangle.

Using the “minimal knowledge requirements” described in Fig. 1, we can now characterize a perfectly (objective) privacy-preserving auction protocol as one that induces this *exact* partition of the space of possible inputs into subspaces in which the inputs are indistinguishable to the auctioneer. Unfortunately, perfect privacy is often hard or even impossible to

achieve. For  $2^{nd}$ -price auctions, Brandt and Sandholm [4] show that *every* perfectly private auction protocol has exponential communication complexity. This provides the motivation for our definition of *privacy-approximation ratio*: We are interested in whether there is an auction protocol that achieves “good” privacy guarantees without paying such a high price in computational efficiency. We no longer insist that the auction protocol induce a partition of inputs *exactly* as in Fig. 1 but rather that it “approximate” the optimal partition well. We define two kinds of privacy-approximation ratio (PAR): *worst case* PAR and *average-case* PAR.

The worst-case PAR of a protocol  $P$  for the  $2^{nd}$ -price auction is defined as the maximum ratio between the size of a set  $S$  of indistinguishable inputs in Fig. 1 and the size of a set of indistinguishable inputs induced by  $P$  that is contained in  $S$ . If a protocol is perfectly privacy preserving, these sets are always the same size, and so the worst-case PAR is 1. If, however, a protocol fails to achieve perfect privacy, then at least one “ideal” set of indistinguishable inputs *strictly* contains a set of indistinguishable inputs induced by the protocol. In such cases, the worst-case PAR will be strictly higher than 1.

Consider, *e.g.*, the sealed-bid auction protocol in which both bidders reveal their private information to the auctioneer, who then computes the outcome. Obviously, this naive protocol enables the auctioneer to distinguish between *every* two pairs of private inputs, and so each set of indistinguishable inputs induced by the protocol contains *exactly one* element. The worst-case PAR of this protocol is therefore  $\frac{8}{1} = 8$ . (If bidder 2’s value is 0, then in Fig. 1 the auctioneer is unable to determine which value in  $\{0, \dots, 7\}$  is  $x_1$ . In the sealed bid auction protocol, however, the auctioneer learns the exact value of  $x_1$ .) The *average-case* PAR is a natural Bayesian variant of this definition: We now assume that the auctioneer has knowledge of some *market statistics*, in the form of a probability distribution over the possible private information of the bidders. PAR in this case is defined as the *average* ratio and not as the maximum ratio as before.

To formalize and generalize the above intuitive definitions, we make use of machinery from communication-complexity theory. Specifically, we use the concepts of *monochromaticity* and *tilings* to make formal the notions of sets of indistinguishable inputs and of the approximability of privacy.

## 1.2 Our Findings

We present both upper and lower bounds on the privacy-approximation ratio for both the millionaires problem and  $2^{nd}$ -price auctions with 2 bidders. Our analysis of these two environments takes place within Yao’s 2-party communication model [19], in which the private information of each party is a  $k$ -bit string, representing a value in  $\{0, \dots, 2^k - 1\}$ . In the millionaires problem, the two parties (the millionaires) wish to keep their private information hidden *from each other*. Hence, their goal is to preserve *subjective* privacy. In electronic-commerce environments, each party (bidder) often communicates with the auctioneer via a secure channel, and so the objective in the  $2^{nd}$ -price auction is to prevent *a third party* (the auctioneer) from learning “too much” about the bidders. This requirement is equivalent to what we refer to as “*objective* privacy preservation”, in a setting in which the two parties

communicate directly with each other (simulating the auctioneer).

Informally, for both the  $2^{nd}$ -price Vickrey auction and the millionaires problem, we obtain the following results: We show that not only is perfect privacy impossible or infeasibly costly to achieve, but even close approximations of perfect privacy suffer from the same lower bounds. By contrast, we show that, if the values of the parties are drawn uniformly at random from  $\{0, \dots, 2^k - 1\}$ , then, for both problems, simple and natural communication protocols have privacy-approximation ratios that are linear in  $k$  (*i.e.*, logarithmic in the size of the space of possible inputs). We conjecture that this improved PAR is achievable for *any* probability distribution. The correctness of this conjecture would imply that, no matter what *beliefs* the protocol designer may have about the parties' private values, a protocol that achieves reasonable privacy guarantees exists.

### 1.3 Related Work: Defining Privacy-Preserving Computation

As explained above, the privacy work of Bar-Yehuda *et al.* [1] and the work presented in this paper have common ancestors in [6, 15]. Similarly, the work of Brandt and Sandholm [4] uses Kushilevitz's formulation to prove an exponential lower bound on the communication complexity of privacy-preserving  $2^{nd}$ -price Vickrey auctions.

There are many formulations of privacy-preserving computation, both exact and approximate, that are not based on the definitions and tools in [6, 15]. We now briefly review some of them and explain how they differ from ours.

The most extensively developed approach to privacy in distributed computation is that of *secure, multiparty function evaluation* (SMFE). Indeed, to achieve agent privacy in algorithmic mechanism design, which was our original motivation, one could, in principle, simply start with a strategyproof mechanism and then have the agents themselves compute the outcome and payments using an SMFE protocol. However, as observed by Brandt and Sandholm [4], these protocols fall into two main categories, and both have inherent disadvantages from the point of view of mechanism design. *Information-theoretically* private protocols, the study of which was initiated by Ben-Or, Goldwasser, and Wigderson [3] and Chaum, Crépeau, and Damgaard [5], rely on the assumption that a constant fraction of the agents are “honest” (or “obedient” in the terminology of distributed algorithmic mechanism design [10]), *i.e.*, that they follow the protocol perfectly even if they know that doing so will lead to an outcome that is not as desirable to them as one that would result from their deviating from the protocol; clearly, this assumption is antithetical to the main premise of mechanism design, which is that all agents will behave strategically, deviating from protocols when and only when doing so will improve the outcome from their points of view. Multiparty protocols that use *cryptography* to achieve privacy, the study of which was initiated by Yao [20, 21], rely on (plausible but currently unprovable) complexity-theoretic assumptions. Often, they are also very communication-intensive; for some mechanisms of interest, efficient cryptographic protocols have been obtained (see, for example, [7, 17]).

In this paper, we consider protocols that compute exact results but preserve privacy only approximately. One can also ask what it means for a protocol to compute approximate results in a privacy-preserving manner; indeed, this question has also been studied [2, 9, 14],

but it is unrelated to the questions we ask here. Similarly, definitions and techniques from *differential privacy* [8] (and its mechanism-design extensions [11]), in which the goal is to add noise to the result of a database query in such a way as to preserve the privacy of the individual database records (and hence protect the data subjects) but still have the result convey nontrivial information, are inapplicable to the problems that we study here.

## 1.4 Paper Outline

In the next section, we review and expand upon the connection between perfect privacy and communication complexity. We present our formulations of approximate privacy, both worst-case and average-case, in Section 3 and our main results in Section 4. Discussion and future directions can be found in Section 5.

# 2 Perfect Privacy and Communication Complexity

We now briefly review Yao’s model of two-party communication and notions of objective and subjective perfect privacy; see Kushilevitz and Nisan [16] for a comprehensive overview of communication complexity theory. Note that we only deal with *deterministic* communication protocols. Our definitions can be extended to randomized protocols.

## 2.1 Two-Party Communication Model

There are two parties, 1 and 2, each holding a  $k$ -bit *input string*. The input of party  $i$ ,  $x_i \in \{0, 1\}^k$ , is the *private information* of  $i$ . The parties communicate with each other in order to compute the value of a function  $f : \{0, 1\}^k \times \{0, 1\}^k \rightarrow \{0, 1\}^t$ . The two parties alternately send messages to each other. In communication round  $j$ , one of the parties sends a bit  $q_j$  that is a function of that party’s input and the history  $(q_1, \dots, q_{j-1})$  of previously sent messages. We say that a bit is *meaningful* if it is not a constant function of this input and history and if, for every meaningful bit transmitted previously, there some combination of input and history for which the bit differs from the earlier meaningful bit. Non-meaningful bits (*e.g.*, those sent as part of protocol-message headers) are irrelevant to our work here and will be ignored. A *communication protocol* dictates, for each party, when it is that party’s turn to transmit a message and what message he should transmit, based on the history of messages and his value.

A communication protocol  $P$  is said to compute  $f$  if, for every pair of inputs  $(x_1, x_2)$ , it holds that  $P(x_1, x_2) = f(x_1, x_2)$ . As in [15], the last message sent in a protocol  $P$  is assumed to contain the value  $f(x_1, x_2)$  and therefore may require up to  $t$  bits. The *communication complexity* of a protocol  $P$  is the maximum, over all input pairs, of the number of bits transmitted during the execution of  $P$ .

Any function  $f : \{0, 1\}^k \times \{0, 1\}^k \rightarrow \{0, 1\}^t$  can be visualized as a  $2^k \times 2^k$  matrix with entries in  $\{0, 1\}^t$ , in which the rows represent the possible inputs of party 1, the columns

represent the possible inputs of party 2, and each entry contains the value of  $f$  associated with its row and column inputs. This matrix is denoted by  $A(f)$ .

**Definition 1 (Regions, partitions)** *A region in a matrix  $A$  is any subset of entries in  $A$  (not necessarily a submatrix). A partition of  $A$  is a collection of disjoint regions in  $A$  whose union equals  $A$ .*

**Definition 2 (Monochromaticity)** *A region  $R$  in a matrix  $A$  is called monochromatic if all entries in  $R$  contain the same value. A monochromatic partition of  $A$  is a partition all of whose regions are monochromatic.*

Of special interest in communication complexity are specific kinds of regions and partitions called rectangles, and tilings, respectively:

**Definition 3 (Rectangles, Tilings)** *A rectangle in a matrix  $A$  is a submatrix of  $A$ . A tiling of a matrix  $A$  is a partition of  $A$  into rectangles.*

**Definition 4 (Refinements)** *A tiling  $T_1(f)$  of a matrix  $A(f)$  is said to be a refinement of another tiling  $T_2(f)$  of  $A(f)$  if every rectangle in  $T_1(f)$  is contained in some rectangle in  $T_2(f)$ .*

Monochromatic rectangles and tilings are an important concept in communication-complexity theory, because they are linked to the execution of communication protocols. Every communication protocol  $P$  for a function  $f$  can be thought of as follows:

1. Let  $R$  and  $C$  be the sets of row and column indices of  $A(f)$ , respectively. For  $R' \subseteq R$  and  $C' \subseteq C$ , we will abuse notation and write  $R' \times C'$  to denote the submatrix of  $A(f)$  obtained by deleting the rows not in  $R'$  and the columns not in  $C'$ .
2. While  $R \times C$  is not monochromatic:
  - One party  $i \in \{0, 1\}$  sends a single bit  $q$  (whose value is based on  $x_i$  and the history of communication).
  - If  $i = 1$ ,  $q$  indicates whether 1's value is in one of two disjoint sets  $R_1, R_2$  whose union equals  $R$ . If  $x_1 \in R_1$ , both parties set  $R = R_1$ . If  $x_1 \in R_2$ , both parties set  $R = R_2$ .
  - If  $i = 2$ ,  $q$  indicates whether 2's value is in one of two disjoint sets  $C_1, C_2$  whose union equals  $C$ . If  $x_2 \in C_1$ , both parties set  $C = C_1$ . If  $x_2 \in C_2$ , both parties set  $C = C_2$ .
3. One of the parties sends a last message (consisting of up to  $t$  bits) containing the value in all entries of the monochromatic rectangle  $R \times C$ .






Figure 2: A tiling that is not induced by any communication protocol [15]

Observe that, for every pair of private inputs  $(x_1, x_2)$ ,  $P$  terminates at some monochromatic rectangle in  $A(f)$  that contains  $(x_1, x_2)$ . We refer to this rectangle as “*the monochromatic rectangle induced by  $P$  for  $(x_1, x_2)$* ”. We refer to the tiling that consists of all rectangles induced by  $P$  (for all pairs of inputs) as “*the monochromatic tiling induced by  $P$* ”.

**Remark 2.1** *There are monochromatic tilings that cannot be induced by communication protocols. For example, observe that the tiling in Fig. 2 (which is essentially an example from [15]) has this property.*

## 2.2 Perfect Privacy

Informally, we say that a two-party protocol is *perfectly privacy-preserving* if the two parties (or a third party observing the communication between them) cannot learn more from the execution of the protocol than the value of the function the protocol computes. (These definition can be extended naturally to protocols involving more than two participants.)

Formally, let  $P$  be a communication protocol for a function  $f$ . The *communication string* passed in  $P$  is the concatenation of all the messages  $(q_1, q_2 \dots)$  sent in the course of the execution of  $P$ . Let  $s_{(x_1, x_2)}$  denote the communication string passed in  $P$  if the inputs of the parties are  $(x_1, x_2)$ . We are now ready to define perfect privacy. The following two definitions handle privacy from the point of view of a party  $i$  that does not want the other party (that is, of course, familiar not only with the communication string, but also with *his own* value) to learn more than necessary about  $i$ ’s private information. We say that a protocol is perfectly private with respect to party 1 if 1 never learns more about party 2’s private information than necessary to compute the outcome.

**Definition 5 (Perfect privacy with respect to 1)** [6, 15]  *$P$  is perfectly private with respect to party 1 if, for every  $x_2, x'_2$  such that  $f(x_1, x_2) = f(x_1, x'_2)$ , it holds that  $s_{(x_1, x_2)} = s_{(x_1, x'_2)}$ .*

Informally, Def. 5 says that party 1’s knowledge of the communication string passed in the protocol and his knowledge of  $x_1$  do not aid him in distinguishing between two possible inputs of 2. Similarly:

**Definition 6 (Perfect privacy with respect to 2)** [6, 15]  *$P$  is perfectly private with respect to party 2 if, for every  $x_1, x'_1$  such that  $f(x_1, x_2) = f(x'_1, x_2)$ , it holds that  $s_{(x_1, x_2)} = s_{(x'_1, x_2)}$ .*

**Observation 2.2** For any function  $f$ , the protocol in which party  $i$  reveals  $x_i$  and the other party computes the outcome of the function is perfectly private with respect to  $i$ .

**Definition 7 (Perfect subjective privacy)**  $P$  achieves perfect subjective privacy if it is perfectly private with respect to both parties.

The following definition considers a different form of privacy—privacy from a *third party* that observes the communication string but has no *a priori* knowledge about the private information of the two communicating parties. We refer to this notion as “*objective privacy*”.

**Definition 8 (Perfect objective privacy)**  $P$  achieves perfect objective privacy if, for every two pairs of inputs  $(x_1, x_2)$  and  $(x'_1, x'_2)$  such that  $f(x_1, x_2) = f(x'_1, x'_2)$ , it holds that  $s_{(x_1, x_2)} = s_{(x'_1, x'_2)}$ .

Kushilevitz [15] was the first to point out the interesting connections between perfect privacy and communication-complexity theory. Intuitively, we can think of any monochromatic rectangle  $R$  in the tiling induced by a protocol  $P$  as a set of inputs that are *indistinguishable* to a third party. This is because, by definition of  $R$ , for any two pairs of inputs in  $R$ , the communication string passed in  $P$  must be the same. Hence we can think of the privacy of the protocol in terms of the tiling induced by that protocol.

Ideally, every two pairs of inputs that are assigned the same outcome by a function  $f$  will belong to the same monochromatic rectangle in the tiling induced by a protocol for  $f$ . This observation enables a simple characterization of perfect privacy-preserving mechanisms.

**Definition 9 (Ideal monochromatic partitions)** A monochromatic region in a matrix  $A$  is said to be a maximal monochromatic region if no monochromatic region in  $A$  properly contains it. The ideal monochromatic partition of  $A$  is made up of the maximal monochromatic regions.

**Observation 2.3** For every possible value in a matrix  $A$ , the maximal monochromatic region that corresponds to this value is unique. This implies the uniqueness of the ideal monochromatic partition for  $A$ .

**Observation 2.4 (A characterization of perfectly privacy-preserving protocols)** A communication protocol  $P$  for  $f$  is perfectly privacy-preserving iff the monochromatic tiling induced by  $P$  is the ideal monochromatic partition of  $A(f)$ . This holds for all of the above notions of privacy.

### 3 Privacy-Approximation Ratios

Unfortunately, perfect privacy should not be taken for granted. As we shall see in Sec. 4, in many environments, perfect privacy can be either impossible or very costly (in terms of communication complexity) to obtain. To measure the *level* of privacy provided by a communication protocol we introduce the notion of *privacy-approximation ratios* (PARs).

### 3.1 Worst-Case PARs

For any communication protocol  $P$  for a function  $f$ , we denote by  $R^P(x_1, x_2)$  the monochromatic rectangle induced by  $P$  for  $(x_1, x_2)$ . We denote by  $R^I(x_1, x_2)$  the monochromatic region containing  $A(f)_{(x_1, x_2)}$  in the ideal monochromatic partition of  $A(f)$ . Intuitively,  $R^P(x_1, x_2)$  is the set of inputs that are indistinguishable from  $(x_1, x_2)$  to  $P$ .  $R^I(x_1, x_2)$  is the set of inputs that *would be* indistinguishable from  $(x_1, x_2)$  if perfect privacy were preserved. We wish to assess how far one is from the other. The size of a region  $R$ , denoted by  $|R|$ , is the cardinality of  $R$ , *i.e.*, the number of inputs in  $R$ .

We can now define worst-case *objective* PAR as follows:

**Definition 10 (Worst-case objective PAR of  $P$ )** *The worst-case objective privacy-approximation ratio of communication protocol  $P$  for function  $f$  is*

$$\alpha = \max_{(x_1, x_2)} \frac{|R^I(x_1, x_2)|}{|R^P(x_1, x_2)|}.$$

*We say that  $P$  is  $\alpha$ -objective-privacy-preserving in the worst case.*

**Definition 11 ( $i$ -partitions)** *The 1-partition of a region  $R$  in a matrix  $A$  is the set of disjoint rectangles  $R_{x_1} = \{x_1\} \times \{x_2 \text{ s.t. } (x_1, x_2) \in R\}$  (over all possible inputs  $x_1$ ). 2-partitions are defined analogously.*

Intuitively, given any region  $R$  in the matrix  $A(f)$ , if party  $i$ 's actual private information is  $x_i$ , then  $i$  can use this knowledge to eliminate all the parts of  $R$  other than  $R_{x_i}$ . Hence, the other party should be concerned not with  $R$  but rather with the  $i$ -partition of  $R$ .

**Definition 12 ( $i$ -induced tilings)** *The  $i$ -induced tiling of a protocol  $P$  is the refinement of the tiling induced by  $P$  obtained by  $i$ -partitioning each rectangle in it.*

**Definition 13 ( $i$ -ideal monochromatic partitions)** *The  $i$ -ideal monochromatic partition is the refinement of the ideal monochromatic partition obtained by  $i$ -partitioning each region in it.*

For any communication protocol  $P$  for a function  $f$ , we use  $R_i^P(x_1, x_2)$  to denote the monochromatic rectangle containing  $A(f)_{(x_1, x_2)}$  in the  $i$ -induced tiling for  $P$ . We denote by  $R_i^I(x_1, x_2)$  the monochromatic rectangle containing  $A(f)_{(x_1, x_2)}$  in the  $i$ -ideal monochromatic partition of  $A(f)$ .

**Definition 14 (Worst-case PAR of  $P$  with respect to  $i$ )** *The worst-case privacy-approximation ratio with respect to  $i$  of communication protocol  $P$  for function  $f$  is*

$$\alpha = \max_{(x_1, x_2)} \frac{|R_i^I(x_1, x_2)|}{|R_i^P(x_1, x_2)|}.$$

*We say that  $P$  is  $\alpha$ -privacy-preserving with respect to  $i$  in the worst case.*

**Definition 15 (Worst-case subjective PAR of  $P$ )** *The worst-case subjective privacy-approximation ratio of communication protocol  $P$  for function  $f$  is the maximum of the worst-case privacy-approximation ratio with respect to each party.*

**Definition 16 (Worst-case PAR)** *The worst-case objective (subjective) PAR for a function  $f$  is the minimum, over all protocols  $P$  for  $f$ , of the worst-case objective (subjective) PAR of  $P$ .*

### 3.2 Average-Case PARs

As we shall see below, it is also useful to define an average-case version of PAR. As the name suggests, the average-case objective PAR is the *average* ratio between the size of the monochromatic rectangle containing the private inputs and the corresponding region in the ideal monochromatic partition.

**Definition 17 (Average-case objective PAR of  $P$ )** *Let  $D$  be a probability distribution over the space of inputs. The average-case objective privacy-approximation ratio of communication protocol  $P$  for function  $f$  is*

$$\alpha = E_D \left[ \frac{|R^I(x_1, x_2)|}{|R^P(x_1, x_2)|} \right].$$

*We say that  $P$  is  $\alpha$ -objective privacy-preserving in the average case with distribution  $D$  (or with respect to  $D$ ).*

We define average-case PAR with respect to  $i$  analogously, and average-case subjective PAR as the maximum over all players  $i$  of the average-case PAR with respect to  $i$ . We define the *average-case objective (subjective) PAR for a function  $f$*  as the minimum, over all protocols  $P$  for  $f$ , of the average-case objective (subjective) PAR of  $P$ .

## 4 Bounds on Privacy-Approximation Ratios

In this section, we prove upper and lower bounds on the privacy-approximation ratios for two classic problems: Yao’s millionaires problem and the  $2^{nd}$ -price Vickrey auction.

### 4.1 Problem Specifications

**The millionaires problem.** Two millionaires want to know which one is richer. Each millionaire’s wealth is private information known only to him, and the millionaire wishes to keep it that way. The goal is to discover the identity of the richer millionaire while preserving the (subjective) privacy of both parties.

**Definition 18** (THE MILLIONAIRES PROBLEM<sub>k</sub>)

*Input:*  $x_1, x_2 \in \{0, \dots, 2^k - 1\}$  (each represented by a  $k$ -bit string)

*Output:* the identity of the party with the higher value, i.e.,  $\arg \max_{i \in \{0,1\}} x_i$  (breaking ties lexicographically).

No perfectly privacy-preserving communication protocol for THE MILLIONAIRES PROBLEM<sub>k</sub> exists [15]. Hence, we are interested in the PARs for this well studied problem.

**2<sup>nd</sup>-price Vickrey auction.** A single item is offered to 2 bidders, each with a private value for the item. The auctioneer’s goal is to allocate the item to the bidder with the highest value. The fundamental technique in mechanism design for inducing truthful behavior in single-item auctions is Vickrey’s 2<sup>nd</sup>-price auction [18]: Allocate the item to the highest bidder, and charge him the second-highest bid.

**Definition 19** (2ND-PRICE AUCTION<sub>k</sub>)

*Input:*  $x_1, x_2 \in \{0, \dots, 2^k - 1\}$  (each represented by a  $k$ -bit string)

*Output:* the identity of the party with the higher value, i.e.,  $\arg \max_{i \in \{0,1\}} x_i$  (breaking ties lexicographically), and the private information of the of the other party.

Brandt and Sandholm [4] show that a perfectly privacy-preserving communication protocol exists for 2ND-PRICE AUCTION<sub>k</sub>. Specifically, perfect privacy is obtained via the *ascending-price English auction*: Start with a price of  $p = 0$  for the item. In each time step, increase  $p$  by 1 until one of the bidders indicates that his value for the item is less than  $p$ . At that point, allocate the item to the other bidder for a price of  $p - 1$ . If  $p$  reaches a value of  $2^k - 1$  (that is, the values of both bidders are  $2^k - 1$ ) allocate the item to bidder 1 for a price of  $2^k - 1$ .

Moreover, it is shown in [4] that the English auction is essentially *the only* perfectly privacy-preserving protocol for 2ND-PRICE AUCTION<sub>k</sub>. Thus, perfect privacy requires, in the worst-case, the transmission of  $\Omega(2^k)$  bits.  $2k$  bits suffice, because bidders can simply reveal their inputs. Can we obtain “good” privacy without paying such a high price in communication?

## 4.2 The Millionaires Problem

The following theorem shows that not only is perfect subjective privacy unattainable for THE MILLIONAIRES PROBLEM<sub>k</sub>, but a stronger result holds:

**Theorem 4.1 (A worst-case lower bound on subjective PAR)** *No communication protocol for THE MILLIONAIRES PROBLEM<sub>k</sub> obtains a subjective PAR less than  $2^{\frac{k}{2}}$  in the worst-case.*

**Proof:** Consider a communication protocol  $P$  for THE MILLIONAIRES PROBLEM<sub>k</sub>. Let  $R$  represent the space of possible inputs of millionaire 1, and let  $C$  represent the space of possible inputs of millionaire 2. In the beginning,  $R = C = \{0, \dots, 2^k - 1\}$ . Consider the

first (meaningful) bit  $q$  transmitted in course of  $P$ 's execution. Let us assume that this bit is transmitted by millionaire 1. This bit indicates whether 1's value belongs to one of two disjoint subsets of  $R$ ,  $R_1$  and  $R_2$ , whose union equals  $R$ . Because we are interested in the worst case, we can choose adversarially to which of these subsets 1's input belongs. Without loss of generality, let  $0 \in R_1$ . We decide adversarially that 1's value is in  $R_1$  and set  $R = R_1$ . Similarly, if  $q$  is transmitted by millionaire 2, then we set  $C$  to be the subset of  $C$  containing 0 in the partition of 2's inputs induced by  $q$ . We continue this process recursively for each bit transmitted in  $P$ .

Observe that, as long as both  $R$  and  $C$  contain at least two values,  $P$  is incapable of computing THE MILLIONAIRES PROBLEM $_k$ . This is because 0 belongs to both  $R$  and  $C$ , and so  $P$  cannot eliminate, for either of the millionaires, the possibility that that millionaire has a value of 0 and the other millionaire has a positive value. Hence, this process will go on until  $P$  determines that the value of one of the millionaires is exactly 0, *i.e.*, until either  $R = \{0\}$  or  $C = \{0\}$ . Let us examine these two cases:

- **Case I:**  $R = \{0\}$ . Consider the subcase in which  $x_2$  equals 0. Recall that  $0 \in C$ , and so this is possible. Observe that, in this case,  $P$  determines the exact value of  $x_1$ , despite the fact that, in the 2-ideal-monochromatic partition, all  $2^k$  possible values of  $x_1$  are in the same monochromatic rectangle when  $x_2 = 0$  (because for all these values 1 wins). Hence, we get a lower bound of  $2^k$  on the subjective privacy-approximation ratio.

- **Case II:**  $C = \{0\}$ . Let  $m$  denote the highest input in  $R$ . We consider two subcases. If  $m \leq 2^{\frac{k}{2}}$ , then observe that the worst-case subjective privacy-approximation ratio is at least  $2^{\frac{k}{2}}$ . In the 2-ideal-monochromatic partition, all  $2^k$  possible values of  $x_1$  are in the same monochromatic rectangle if  $x_2 = 0$ , and the fact that  $m \leq 2^{\frac{k}{2}}$  implies that  $|R| \leq 2^{\frac{k}{2}}$ .

If, on the other hand,  $m > 2^{\frac{k}{2}}$ , then consider the case in which  $x_1 = m$  and  $x_2 = 0$ . Observe that, in the 1-ideal-monochromatic partition, all values of millionaire 2 in  $\{0, \dots, m - 1\}$  are in the same monochromatic rectangle if  $x_1 = m$ . However,  $P$  will enable millionaire 1 to determine that millionaire 2's value is exactly 0. This implies a lower bound of  $m$  on the subjective privacy-approximation. We now use the fact that  $m > 2^{\frac{k}{2}}$  to conclude the proof. □

By contrast, we show that fairly good privacy guarantees can be obtained in the average case. We define the BISECTION PROTOCOL for THE MILLIONAIRES PROBLEM $_k$  as follows: Ask each millionaire whether his value lies in  $[0, 2^{k-1})$  or in  $[2^{k-1}, 2^k)$ ; continue this binary search until the millionaires' answers differ, at which point we know which millionaire has the higher value. If the answers never differ the tie is broken in favor of millionaire 1.

We may exactly compute the average-case subjective PAR with respect to the uniform distribution for the BISECTION PROTOCOL applied to THE MILLIONAIRES PROBLEM $_k$ . Figure 3 illustrates the approach. The far left of the figure shows the ideal partition (for

$k = 3$ ) of the value space for THE MILLIONAIRES PROBLEM $_k$ ; these regions are indicated with heavy lines in all parts of the figure. The center-left shows the 1-partition of the regions in the ideal partition; the center-right shows the 1-induced tiling that is induced by the BISECTION PROTOCOL. The far right illustrates how we may rearrange the tiles that partition the bottom-left region in the ideal partition (by reflecting them across the dashed line) to obtain a tiling of the value space that is the same as the tiling induced by applying the BISECTION AUCTION to 2ND-PRICE AUCTION $_k$ .




Figure 3: Left to right: The ideal partition (for  $k = 3$ ) for THE MILLIONAIRES PROBLEM $_k$ ; the 1-partition of the ideal regions; the 1-induced tiling induced by the BISECTION PROTOCOL; the rearrangement used in the proof of Thm. 4.2

**Theorem 4.2 (The average-case subjective PAR of the bisection protocol)** *The average-case subjective PAR with respect to the uniform distribution for the BISECTION PROTOCOL applied to THE MILLIONAIRES PROBLEM $_k$  is  $\frac{k}{2} + 1$ .*

**Proof:** Given a value of  $i$ , consider the  $i$ -induced-tiling obtained by running the BISECTION PROTOCOL for THE MILLIONAIRES PROBLEM $_k$  (as in the center-right of Fig. 3 for  $i = 1$ ). Rearrange the rectangles in which player  $i$  wins by reflecting them across the line running from the bottom-left corner to the top-right corner (the dashed line in the far right of Fig. 3). This produces a tiling of the value space in which the region in which player 1 wins is tiled by tiles of width 1, and the region in which player 2 wins is tiled by tiles of height 1; in computing the average-case-approximate-privacy with respect to  $i$ , the tile-size ratios that we use are the heights (widths) of the tiles to the height (width) of the tile containing all values in that column (row) for which player 1 (2) wins. This tiling and the tile-size ratios in question are exactly as in the computation of the average-case *objective* privacy for 2ND-PRICE AUCTION $_k$ ; the argument used in Thm. 4.12 (for  $g(k) = k$ ) below completes the proof.  $\square$

Consider the case in which a third party is observing the interaction of the two millionaires. How much can this observer learn about the private information of the two millionaires? We show that, unlike the case of subjective privacy, good PARs are unattainable even in the average case.

Because the values  $(i, i)$  (in which case player 1 wins) and the values  $(i, i + 1)$  (in which player 2 wins) must all appear in different tiles in any tiling that refines the ideal partition of the value space for THE MILLIONAIRES PROBLEM $_k$ , any such tiling must include at least  $2^k$  tiles in which player 1 wins and  $2^k - 1$  tiles in which player 2 wins. The total contribution of a tile in which player 1 wins is the number of values in that tile times the ratio of the ideal region containing the tile to the size of the tile, divided by the total number ( $2^{2k}$ ) of values in the space. Each tile in which player 1 wins thus contributes  $\frac{(1+2^k)2^k}{2^{2k+1}}$  to the average-case privacy-approximation ratio under the uniform distribution; similarly, each tile in which player 2 wins contributes  $\frac{2^k(2^k-1)}{2^{2k+1}}$  to this ratio. This leads directly to the following result.

**Proposition 4.3 (A lower bound on average-case objective PAR)** *The average-case objective PAR for THE MILLIONAIRES PROBLEM $_k$  with respect to the uniform distribution is at least  $2^k - \frac{1}{2} + 2^{-(k+1)}$ .*

There are numerous different tilings of the value space that achieve this ratio and that can be realized by communication protocols. For the BISECTION PROTOCOL, we obtain the same exponential (in  $k$ ) growth rate but with a larger constant factor.

**Proposition 4.4 (The average-case objective PAR of the bisection protocol)** *The BISECTION PROTOCOL for THE MILLIONAIRES PROBLEM $_k$  obtains an average-case objective PAR of  $3 \cdot 2^{k-1} - \frac{1}{2}$  with respect to the uniform distribution.*

**Proof:** The bisection mechanism induces a tiling that refines the ideal partition and that has  $2^{k+1} - 1$  tiles in which the player 1 wins and  $2^k - 1$  tiles in which the player 2 wins. The contributions of each of these tiles is as noted above, from which the result follows.  $\square$

Finally, Table 1 summarizes our average-case PAR results (with respect to the uniform distribution) for THE MILLIONAIRES PROBLEM $_k$ .

	Average-Case Obj. PAR	Average-Case Subj. PAR
Any Protocol	$\geq 2^k - \frac{1}{2} + 2^{-(k+1)}$	
Bisection Protocol	$\frac{3}{2}2^k - \frac{1}{2}$	$\frac{k}{2} + 1$

Table 1: Average-case PARs for THE MILLIONAIRES PROBLEM $_k$

### 4.3 $2^{nd}$ -Price Auctions

We now consider *objective privacy* for 2ND-PRICE AUCTION $_k$  (*i.e.*, privacy with respect to the auctioneer). Bisection auctions [12, 13] for 2ND-PRICE AUCTION $_k$  are defined similarly to the BISECTION PROTOCOL for THE MILLIONAIRES PROBLEM $_k$ : Use binary search to find a value  $c$  that lies between the two bidders' values, and let the bidder with the higher value be bidder  $j$ . (If the values do not differ, we will also discover this; in this case, award



the item to bidder 1, who must pay the common value.) Use binary search on the interval that contains the value of the lower bidder in order to find the value of the lower bidder.

More generally, we refer to an auction protocol as a  $c$ -bisection auction, for a constant  $c \in (0, 1)$ , if in each step the interval  $R$  is partitioned into two disjoint subintervals: a lower subinterval of size  $c|R|$  and an upper subinterval of size  $(1 - c)|R|$ . Hence, the BISECTION AUCTION is a  $c$ -bisection auction with  $c = \frac{1}{2}$ . We prove that no  $c$ -bisection auction for THE MILLIONAIRES PROBLEM $_k$  obtains a subexponential objective PAR:

**Theorem 4.5 (A worst-case lower bound for  $c$ -bisection auctions)** *For any constant  $c > \frac{1}{2^k}$ , the  $c$ -bisection auction for 2ND-PRICE AUCTION $_k$  has a worst-case PAR of at least  $2^{\frac{k}{2}}$ .*

**Proof:** Consider the ideal monochromatic partition of 2ND-PRICE AUCTION $_k$  depicted in Fig. 1. Observe that, for perfect privacy to be preserved, it must be that bidder 2 transmits the first (meaningful) bit, and that this bit partitions the space of inputs into the leftmost shaded rectangle (the set  $\{0, \dots, 7\} \times \{0\}$ ) and the rest of the value space (ignoring the rectangles depicted that further refine  $\{0, \dots, 7\} \times \{1, \dots, 7\}$ ). What if the first bit is transmitted by player 2 and does not partition the space into rectangles in that way? We observe that any other partition of the space into two rectangles is such that, in the worst case, the privacy-approximation ratio is at least  $2^{\frac{k}{2}}$  (for any value of  $c$ ): If  $c \leq 1 - 2^{-\frac{k}{2}}$ , then the case in which  $x_1 = c2^k - 1$  gives us the lower bound. If, on the other hand,  $c > 1 - 2^{-\frac{k}{2}}$ , then the case that  $x_1 = 0$  gives us the lower bound. Observe that such a bad PAR is also the result of bidder 1's transmitting the first (meaningful) bit. □

By contrast, as for THE MILLIONAIRES PROBLEM $_k$ , reasonable privacy guarantees are achievable in the average case:

**Theorem 4.6 (The average-case objective PAR of the bisection auction)** *The average-case objective PAR of the BISECTION AUCTION is  $\frac{k}{2} + 1$  with respect to the uniform distribution.*

**Proof:** This follows by taking  $g(k) = k$  in Thm. 4.12. □

We note that the worst-possible approximation of objective privacy comes when the each value in the space is in a distinct tile; this is the tiling induced by the sealed-bid auction. The resulting average-case privacy-approximation ratio is exponential in  $k$ .

**Proposition 4.7 (Largest possible objective PAR)** *The largest possible (for any protocol) average-case objective PAR with respect to the uniform distribution for 2ND-PRICE AUCTION $_k$  is*

$$\frac{1}{2^{2k}} \left[ \sum_{j=0}^{2^k-1} j^2 + \sum_{j=1}^{2^k-1} j^2 \right] = \frac{2}{3}2^k + \frac{1}{3}2^{-k}$$

### 4.3.1 Subjective privacy

We now look briefly at subjective privacy for 2ND-PRICE AUCTION<sub>k</sub>. For subjective privacy with respect to 1, we start with the 1-partition for 2ND-PRICE AUCTION<sub>k</sub>; Fig. 4 shows the refinement of the 1-partition induced by the BISECTION AUCTION for  $k = 4$ . Separately considering the refinement of the 2-partition for 2ND-PRICE AUCTION<sub>k</sub> by the BISECTION AUCTION, we have the following results.




Figure 4: The BISECTION-AUCTION-induced refinement of the 1-partition for 2ND-PRICE AUCTION<sub>k</sub> ( $k = 4$ )

#### Theorem 4.8 (The average-case PAR with respect to 1 of the bisection auction)

The average-case PAR with respect to 1 of the BISECTION AUCTION is

$$\frac{k+3}{4} - \frac{k-1}{2^{k+2}}$$

with respect to the uniform distribution.

**Proof:** The approach is similar to that in the proof of Thm. 4.6. Let  $T_k^1$  be the refinement of the 2ND-PRICE-AUCTION<sub>k</sub> 1-partition of the  $k$ -bit value space induced by the BISECTION AUCTION. Let  $x_k$  be the number of rectangles in  $T_k^1$  in which bidder 2 (the column player) wins; let  $z_k$  be the number in which bidder 1 (the row player) wins. Then we have  $x_{k+1} = x_k + x_k + 2^k$ ; with  $x_0 = 0$ , this gives  $x_k = k2^{k-1}$ . We also have  $z_{k+1} = 2z_k + 2^{2k}$ ; with  $z_0 = 1$ , this gives  $z_k = 2^{k-1}(1 + 2^k)$ .

Let  $y_k$  equal the sum over all rectangles  $R \in T_k^1$  in which bidder 2 wins of the quantity  $2^k - |R|$ . Then the PAR in question equals  $\frac{2^k x_k - y_k + z_k}{2^{2k}}$ . We have that  $y_{k+1} = 2y_k + 2^k x_k + \sum_{i=1}^{2^k} i$ ; with  $y_0 = 0$ , this gives  $y_k = 2^{k-2}(1 - 2^k + k + k2^k)$ . Using the expressions for  $x_k$ ,  $y_k$ , and  $z_k$  that we have derived, we see that the PAR in question equals  $\frac{k+3}{4} - \frac{k-1}{2^{k+2}}$ .  $\square$

#### Theorem 4.9 (The average-case PAR with respect to 2 of the bisection auction)

The average-case PAR with respect to 2 of the BISECTION AUCTION is

$$\frac{k+5}{4} + \frac{k-1}{2^{k+2}}$$

with respect to the uniform distribution.

**Proof:** Let  $T_k^2$  be the refinement of the  $\text{2ND-PRICE-AUCTION}_k$  2-partition of the  $k$ -bit value space induced by the  $\text{BISECTION AUCTION}$ . Let  $c_k$  be the total number of rectangles in  $T_k^2$  in which bidder 1 (the row player) wins; let  $e_k$  be the number of rectangles in which bidder 2 (the column player) wins. Then we have  $c_{k+1} = 2c_k + 2^k$ ; with  $c_0 = 1$ , this gives  $c_k = (k + 2)2^{k-1}$ . We also have  $e_{k+1} = 2e_k + 2^{2k}$ ; with  $e_0 = 0$ , this gives  $e_k = 2^{k-1}(2^k - 1)$ .

Let  $d_k$  equal the sum over all rectangles  $R \in T_k^2$  in which bidder 1 wins of the quantity  $2^k - |R|$ . Then the PAR in question equals  $\frac{2^k c_k - d_k + e_k}{2^{2k}}$ . We have that  $d_{k+1} = 2d_k + 2^k e_k + \sum_{i=0}^{2^k-1} i$ ; with  $d_0 = 0$ , this gives  $d_k = (k + 1)(2^k - 1)2^{k-2}$ . Using the expressions for  $c_k$ ,  $d_k$ , and  $e_k$  that we have derived, we see that the PAR in question equals  $\frac{k+5}{4} + \frac{k-1}{2^{k+2}}$ .  $\square$

**Corollary 4.10 (The average-case subjective PAR of the bisection auction)** *The average-case subjective PAR of the  $\text{BISECTION AUCTION}$  with respect to the uniform distribution is*

$$\frac{k+5}{4} + \frac{k-1}{2^{k+2}}.$$

As for objective privacy, the sealed-bid auction gives the largest possible average-case subjective PAR.

**Proposition 4.11 (Largest possible subjective PAR)** *The largest possible (for any protocol) average-case subjective PAR with respect to the uniform distribution for  $\text{2ND-PRICE AUCTION}_k$  is*

$$\frac{2^k}{3} + 1 - \frac{1}{3 \cdot 2^k}.$$

**Proof:** For the sealed-bid auction, the average-case PAR with respect to 1 is

$$\frac{1}{2^{2k}} \left[ \sum_{j=1}^{2^k} j + \sum_{j=1}^{2^k-1} j^2 \right] = \frac{2^k}{3} + \frac{1}{3 \cdot 2^{k-1}}.$$

For the sealed-bid auction, the average-case PAR with respect to 2 is

$$\frac{1}{2^{2k}} \left[ \sum_{j=1}^{2^k} j^2 + \sum_{j=1}^{2^k-1} j \right] = \frac{2^k}{3} + 1 - \frac{1}{3 \cdot 2^k}.$$

$\square$

### 4.3.2 Bounded-bisection auctions

We now present a middle ground between the perfectly-private yet highly inefficient (in terms of communication) ascending English auction and the communication-efficient  $\text{BISECTION AUCTION}$  whose average-case objective PAR is linear in  $k$  (and is thus unbounded as  $k$  goes to infinity): We bound the number of bisections, using an ascending English auction to determine the outcome if it is not resolved by the limited number of bisections.

We define the  $\text{BISECTION AUCTION}_{g(k)}$  as follows: Given an instance of  $\text{2ND-PRICE AUCTION}_k$ , and an integer-valued function  $g(k)$  such that  $0 \leq g(k) \leq k$ , run the  $\text{BISECTION AUCTION}$  as above but do at most  $g(k)$  bisection operations. (Note that we will never do more than  $k$  bisections.) If the outcome is undetermined after  $g(k)$  bisection operations, so that both players' values lie in an interval  $I$  of size  $2^{k-g(k)}$ , apply the ascending-price English auction to this interval to determine the identity of the winning bidder and the value of the losing bidder.

As  $g(k)$  ranges from 0 to  $k$ , the  $\text{BISECTION AUCTION}_{g(k)}$  ranges from the ascending-price English auction to the  $\text{BISECTION AUCTION}$ . If we allow a fixed, positive number of bisections ( $g(k) = c > 0$ ), computations show that for  $c = 1, 2, 3$  we obtain examples of protocols that do not provide perfect privacy but that do have bounded average-case objective PARs with respect to the uniform distribution. We wish to see if this holds for all positive  $c$ , determine the average-case objective PAR for general  $g(k)$ , and connect the amount of communication needed with the approximation of privacy in this family of protocols. The following theorem allows us to do these things.

**Theorem 4.12** *For the  $\text{BISECTION AUCTION}_{g(k)}$ , the average-case objective PAR with respect to the uniform distribution equals*

$$\frac{g(k) + 3}{2} - \frac{2^{g(k)}}{2^{k+1}} + \frac{1}{2^{k+1}} - \frac{1}{2^{g(k)+1}}.$$

**Proof:** Fix  $k$ , the number of bits used for bidding, and let  $c = g(k)$  be the number of bisections; we have  $0 \leq c \leq k$ , and we let  $i = k - c$ . Figure 5 illustrates this tiling for  $k = 4$ ,  $c = 2$ , and  $i = 2$ ; note that the upper-left and lower-right quadrants have identical structure and that the lower-left and upper-right quadrants have no structure other than that of the ideal partition and the quadrant boundaries (which are induced by the first bisection operation performed).




Figure 5: Illustration for the proof of Thm. 4.12

Our general approach is the following. The average-case objective PAR with respect to the uniform distribution equals

$$\text{PAR} = \frac{1}{2^{2k}} \sum_{(x_1, x_2)} \frac{|R^I(x_1, x_2)|}{|R^P(x_1, x_2)|},$$

where the sum is over all pairs  $(x_1, x_2)$  in the value space; recall that  $R^I(x_1, x_2)$  is the region in the ideal partition that contains  $(x_1, x_2)$ , and  $R^P(x_1, x_2)$  is the rectangle in the tiling induced by the protocol that contains  $(x_1, x_2)$ . We may combine all of the terms corresponding to points in the same protocol-induced rectangle to obtain

$$\text{PAR} = \frac{1}{2^{2k}} \sum_R |R| \frac{|R^I(R)|}{|R|} = \frac{1}{2^{2k}} \sum_R |R^I(R)|, \quad (1)$$

where the sums are now over protocol-induced rectangles  $R$  (we will simplify notation and write  $R$  instead of  $R^P$ ), and  $R^I(R)$  denotes the ideal region that contains the protocol-induced rectangle  $R$ . Each ideal region in which bidder 1 wins is a rectangle of width 1 and height at most  $2^k$ ; each ideal region in which bidder 2 wins is a rectangle of height 1 and width strictly less than  $2^k$ . For a protocol-induced rectangle  $R$ , let  $j_R = 2^k - |R^I(R)|$ . Let  $a_{c,i}$  be the total number of tiles that appear in the tiling of the  $k$ -bit value space induced by the  $\text{BISECTION AUCTION}_{g(k)}$  with  $g(k) = c$ , and let  $b_{c,i} = \sum_R j_R$  (with this sum being over the protocol-induced tiles in this same partition). Then we may rewrite (1) as

$$\text{PAR}_{c,i} = \frac{1}{2^{2k}} \sum_R 2^k - j_R = \frac{a_{c,i} 2^k - b_{c,i}}{2^{2k}}. \quad (2)$$

(Note that (1) holds for general protocols; we now add the subscripts “ $c, i$ ” to indicate the particular PAR we are computing.) We now determine  $a_{c,i}$  and  $b_{c,i}$ .

Considering the tiling induced by  $c+1$  bisections of a  $c+i+1$ -bit space (which has  $a_{c+1,i}$  total tiles), the upper-left and lower-right quadrants each contain  $a_{c,i}$  tiles while the lower-left and upper-right quadrants (as depicted in Fig. 5) each contribute  $2^{c+i}$  tiles, so  $a_{c+1,i} = 2a_{c,i} + 2^{c+i+1}$ . When there are no bisections, the  $i$ -bit value space has  $a_{0,i} = 2^{i+1} - 1$  tiles, from which we obtain  $a_{c,i} = 2^c (2^i (c+2) - 1)$ . The sum of  $j_R$  over protocol-induced rectangles  $R$  in the upper-left quadrant is  $b_{c,i}$ . For a rectangle  $R$  in the lower-right quadrant,  $j_R$  equals  $2^{c+i}$  plus  $j_{R'}$ , where  $R'$  is the corresponding rectangle in the upper-left quadrant; there are  $a_{c,i}$  such  $R$ , so the sum of  $j_R$  over protocol-induced rectangles  $R$  in the upper-left quadrant is  $b_{c,i} + a_{c,i} 2^{c+i}$ . Finally, the sum of  $j_R$  over  $R$  in the lower-left quadrant equals  $\sum_{h=0}^{2^{c+i}-1} h$  and the sum over  $R$  in the top-right quadrant equals  $\sum_{h=1}^{2^{c+i}} h$ . Thus,  $b_{c+1,i} = 2b_{c,i} + a_{c,i} 2^{c+i} + 2^{2(c+i)}$ ; with  $b_{0,i} = \sum_{h=0}^{2^i-1} h + \sum_{h=1}^{2^i-1} h$ , we obtain  $b_{c,i} = 2^{c+i-1} ((1+2^c)(-1+2^i) + 2^{c+i}c)$ . Rewriting (2), we obtain

$$\text{PAR}_{c,i} = \frac{c+3}{2} - \frac{2^c}{2^{c+i+1}} + \frac{1}{2^{c+i+1}} - \frac{1}{2^{c+1}}.$$

Recalling that  $k = c + i$ , the proof is complete.  $\square$

For the protocols corresponding to values of  $g(k)$  ranging from 0 to  $k$  (ranging from the ascending-price English auction to the  $\text{BISECTION AUCTION}$ ), we may thus relate the amount of communication saved (relative to the English auction) to the effect of this on the PAR.

**Corollary 4.13** *Let  $g$  be a function that maps non-negative integers to non-negative integers. Then the average-case objective PAR with respect to the uniform distribution for the*

BISECTION AUCTION $_{g(k)}$  is bounded if  $g$  is bounded and is unbounded if  $g$  is unbounded. We then have that the BISECTION AUCTION $_{g(k)}$  may require the exchange of  $\Theta(k + 2^{k-g(k)})$  bits, and it has an average-case objective PAR of  $\Theta(1 + g(k))$ .

Finally, Table 2 summarizes the average-case PAR results (with respect to the uniform distribution) for 2ND-PRICE AUCTION $_k$ .

	Avg.-Case Obj. PAR	Avg.-Case Subj. PAR
English Auction	1	1
BISECTION AUCTION $_{g(k)}$	$\frac{g(k)+3}{2} - \frac{2^{g(k)}}{2^{k+1}} + \frac{1}{2^{k+1}} - \frac{1}{2^{g(k)+1}}$	
BISECTION AUCTION	$\frac{k}{2} + 1$	$\frac{k+5}{4} + \frac{k-1}{2^{k+2}}$
Sealed-Bid Auction	$\frac{2^{k+1}}{3} + \frac{1}{3 \cdot 2^k}$	$\frac{2^k}{3} + 1 - \frac{1}{3 \cdot 2^k}$

Table 2: Average-case PARs (with respect to the uniform distribution) for 2ND-PRICE AUCTION $_k$

## 5 Discussion and Future Directions

There are many interesting directions for future research:

- We have shown that, for both THE MILLIONAIRES PROBLEM $_k$  and 2ND-PRICE AUCTION $_k$ , reasonable average-case PARs with respect to the uniform distribution are achievable. We conjecture that our upper bounds for these problems extend to *all* possible distributions over inputs.
- An interesting open question is whether one can prove lower bounds on the average-case PARs for THE MILLIONAIRES PROBLEM $_k$  and 2ND-PRICE AUCTION $_k$ .
- It would be interesting to apply the PAR framework presented in this paper to other functions.
- The extension of our PAR framework to the  $n$ -party communication model is a challenging direction for future research.
- Starting from the same place that we did, namely [6, 15], Bar-Yehuda *et al.* [1] provided three definitions of approximate privacy. The one that seems most relevant to the study of privacy-approximation ratios is their notion of *h-privacy*. It would be interesting to know exactly when it is possible to express PARs in terms of *h-privacy* and *vice versa*.

## References

- [1] Reuven Bar-Yehuda, Benny Chor, Eyal Kushilevitz, and Alon Orlitsky. Privacy, additional information, and communication. *IEEE Transactions on Information Theory*, 39:55–65, 1993.
- [2] Amos Beimel, Paz Carmi, Kobbi Nissim, and Enav Weinreb. Private approximation of search problems. In *Proceedings of the ACM Symposium on Theory of Computing*, pages 119–128, 2006.
- [3] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic, fault-tolerant computation. In *Proceedings of the ACM Symposium on Theory of Computing*, 1988.
- [4] Felix Brandt and Tuomas Sandholm. On the existence of unconditionally privacy-preserving auction protocols. *ACM Trans. Inf. Syst. Secur.*, 11(2):1–21, 2008.
- [5] David Chaum, Claude Crépeau, and Ivan Damgaard. Multiparty, unconditionally secure protocols. In *Proceedings of the ACM Symposium on Theory of Computing*, pages 11–19, 1988.
- [6] Benny Chor and Eyal Kushilevitz. A zero-one law for boolean privacy. *SIAM J. Discrete Math.*, 4:36–47, 1991.
- [7] Yevgeniy Dodis, Shai Halevi, and Tal Rabin. A cryptographic solution to a game-theoretic problem. In *Advances in Cryptology—Crypto 2000*, Springer Verlag LNCS.
- [8] Cynthia Dwork. Differential privacy. In *Proceedings of the International Colloquium on Automata, Languages and Programming*, pages 1–12, 2006.
- [9] Joan Feigenbaum, Yuval Ishai, Tal Malkin, Kobbi Nissim, Martin J. Strauss, and Rebecca N. Wright. Secure multiparty computation of approximations. *ACM Transactions on Algorithms*, 2(3):435–472, 2006.
- [10] Joan Feigenbaum and Scott Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *Proceedings of the ACM Mobicom Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, pages 1–13, 2002.
- [11] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally utility-maximizing privacy mechanisms. In *Proceedings of the ACM Symposium on Theory of Computing*, 2009. To appear.
- [12] Elena Grigorieva, P. Jean-Jacques Herings, and Rudolf Müller. The communication complexity of private value single-item auctions. *Oper. Res. Lett.*, 34(5):491–498, 2006.

- [13] Elena Grigorieva, P. Jean-Jacques Herings, Rudolf Müller, and Dries Vermeulen. The private value single item bisection auction. *Economic Theory*, 30(1):107–118, January 2007.
- [14] Shai Halevi, Robert Krauthgamer, Eyal Kushilevitz, and Kobbi Nissim. Private approximation of NP-hard functions. In *Proceedings of the ACM Symposium on Theory of Computing*, pages 550–559, 2001.
- [15] Eyal Kushilevitz. Privacy and communication complexity. *SIAM J. Discrete Math.*, 5(2):273–284, 1992.
- [16] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [17] Moni Naor, Benny Pinkas, and Reuben Sumner. Privacy preserving auctions and mechanism design. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 129–139, 1999.
- [18] William Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.
- [19] Andrew C. Yao. Some complexity questions related to distributive computing (preliminary report). In *Proceedings of the ACM Symposium on Theory of Computing*, pages 209–213, 1979.
- [20] Andrew C. Yao. Protocols for secure computation. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*, pages 160–164, 1982.
- [21] Andrew C. Yao. How to generate and exchange secrets. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*, pages 162–167, 1986.