Approximate Privacy: Foundations and Quantification (Extended Abstract)

Joan Feigenbaum*
Dept. of Computer Science
Yale Univ.
New Haven, CT, USA
joan.feigenbaum@yale.edu

Aaron D. Jaggard†
DIMACS
Rutgers Univ.
Piscataway, NJ, USA
adj@dimacs.rutgers.edu

Michael Schapira‡
Depts. of Computer Science
Yale Univ. and UC Berkeley
New Haven, CT and
Berkeley, CA, USA
michael.schapira@yale.edu

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. Despite many careful formulations and extensive study, there are still open questions about the feasibility of maintaining meaningful privacy 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 many well studied contexts: the 2^{nd}*-price Vickrey auction* [20], the *millionaires problem* of Yao [22], the *provisioning of a public good*, and also *set disjointness* and *set intersection*. We present both positive and negative results and many interesting directions for future research.

Categories and Subject Descriptors

K.4.1 [**Computers and Society**]: Public Policy Issues—*privacy*

General Terms

Economics, Security

Keywords

Approximate privacy, bisection auction

*Supported in part by NSF grants 0331548 and 0716223 and IARPA grant FA8750-07-0031.

†Supported in part by NSF grants 0751674 and 0753492.

‡Supported by NSF grant 0331548.

EC'10, June 7–11, 2010, Cambridge, Massachusetts, USA.

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 deals extensively with the question of 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 as much to contribute 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.

1.1 Our Approach

Consider an auction of a Bluetooth headset with two 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 $x_i \in \{0, \ldots, 7\}$ that is the maximum he is willing to pay for the headset. The item is sold in a 2^{nd}-price Vickrey auction, *i.e.*, the higher bidder gets the item (with ties broken in favor of bidder 1), and the price he pays is the lower bid. The demand for privacy arises naturally in such scenarios [18]: In a straightforward protocol, the auction-

eer 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.

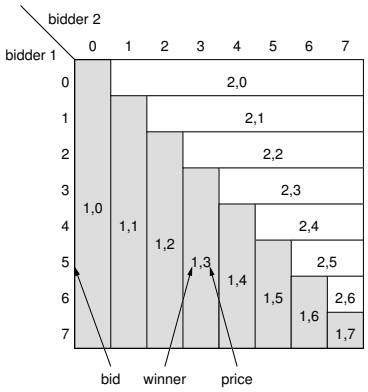

Figure 1: The minimal knowledge requirements for 2^{nd}-price auctions

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 illustrates the information the auctioneer *must* learn in order to determine the outcome of the 2^{nd}-price auction described above. 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 characterize a perfectly 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 [5] 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 $\{0, \ldots, 7\}$ 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.

Thus, intuitively, PAR captures the effect of a protocol on the privacy afforded to protocol participants. The best that participants can hope for is that an observer of the entire protocol transcript (*e.g.,* an auctioneer) learns only that their inputs are in the pre-image of the output. In this best case, the PAR is 1. If the best case is impossible or extremely costly to achieve, participants might still benefit from a protocol in which an observer learns that their inputs are in a subset of the pre-image; the larger (on average) this subset is relative to the entire pre-image, the smaller the (average-case) PAR. For example, the auctioneer must learn the losing bidder's private value, which is a lower bound for the winning bidder's private value. However, the auctioneer could use a protocol that also allows him to learn a non-trivial upper bound on the winning bidder's value. If this range of values isn't too small compared to the set of all values that would produce the same winner and price, then the bidders may view the effects of the protocol on their privacy as reasonable (perhaps in exchange for reduced communication required by the protocol). To formalize and generalize this intuitive notion of PAR, we make use of the machinery of 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. We discuss other notions of approximate privacy (probability-mass-based, semantic, entropy-based, and more) in Sec. 8.

1.2 Our Findings

Yao's millionaires problem and 2^{nd}-price auctions. We present both upper and lower bounds on the privacy-approximation ratio for both the *millionaires problem* of Yao [22] and *2^{nd}-price Vickrey auctions* with 2 bidders [20]. Our analysis of these two problems takes place within Yao's

Problem	Protocol	Objective PAR	Subjective PAR	Ratio of Subj. PARs
2ND-PRICE AUCTION$_k$	All	≥ 1	≥ 1	≥ 1
	English Auction	1	1	1
	BISECTION AUCTION$_{g(k)}$	$\frac{g(k)+3}{2} - \frac{2^{g(k)}}{2^{k+1}} + \frac{1}{2^{k+1}} - \frac{1}{2^{g(k)+1}}$	$\frac{g(k)+5}{4} - \frac{1}{2^{g(k)+2}} + \frac{g(k)}{2^{k+2}}$	—
	BISECTION AUCTION	$\frac{k}{2}+1$	$\frac{k+5}{4} + \frac{k-1}{2^{k+2}}$	~ 1
	Sealed-Bid Auction	$\frac{2^{k+1}}{3} + \frac{1}{3\cdot 2^k}$	$\frac{2^k}{3} + 1 - \frac{1}{3\cdot 2^k}$	~ 1
PUBLIC GOOD$_k$	All	$\geq 2^k - \frac{1}{2} + 2^{-(k+1)}$	—	—
TRUTHFUL PUBLIC GOOD$_{k,c}$	All	$\geq 1 + \frac{c^3}{2^{2k+1}}(1 - \frac{1}{c^2})$	—	—
TRUTHFUL PUBLIC GOOD$_k$	All	$\geq 2^{k-1} - \frac{1}{2} + \frac{1}{2^k}$	—	—
THE MILLIONAIRES PROBLEM$_k$	All	$\geq 2^k - \frac{1}{2} + 2^{-(k+1)}$	—	—
	BISECTION PROTOCOL	$\frac{3}{2}2^k - \frac{1}{2}$	$\frac{k}{2}+1$	1
DISJOINTNESS$_k$	All	$\geq \left(\frac{3}{2}\right)^k$	—	—
	Trivial	$\sim 2^k$	$\sim 2^k$	$\sim 2^k$
	1 First	$\sim 2^k$	$\sim \left(\frac{3}{2}\right)^k$	$\sim \frac{2}{k}\left(\frac{3}{2}\right)^k$
	Alternating	$\sim 2^k$	$\sim \frac{3+2\sqrt{2}}{2}\left(\frac{1+\sqrt{2}}{2}\right)^k$	$\sim \sqrt{2}$
INTERSECTION$_k$	All	$\geq \left(\frac{7}{4}\right)^k$	—	—
	Trivial/1 First	$\left(\frac{7}{4}\right)^k$	$\left(\frac{3}{2}\right)^k$	$\left(\frac{3}{2}\right)^k$
	Alternating	$\left(\frac{7}{4}\right)^k$	$\frac{6}{5}\left(\frac{5}{4}\right)^k$	$\frac{3}{2}$

Table 1: Main PAR results; all are average-case with respect to the uniform distribution. Asymptotic results are for $k \to \infty$.

2-party communication model [21], in which the private information of each party is a k-bit string, representing a value in $\{0, \ldots, 2^k - 1\}$. In the millionaires problem, the two parties (the millionaires) wish to keep their private information hidden *from each other*. We refer to this goal as the preservation of *subjective* privacy. In electronic-commerce environments, each party (bidder) often communicates with the auctioneer via a secure channel, and so the aim in the 2^{nd}-price auction is to prevent *a third party* (the auctioneer), who is unfamiliar with any of the parties' private inputs, from learning "too much" about the bidders. We refer to this goal as the preservation of *objective* privacy.

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 $\{0, \ldots, 2^k - 1\}$, 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 on inputs. 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.

Importantly, our results for the 2^{nd}-price Vickrey auction are obtained by proving a more general result about a large family of protocols for single-item auctions, known as *bounded-bisection auctions*, that contains both the celebrated ascending-price English auction and bisection auctions [12, 13]. In Cor. 4.4, we find the tradeoff between PAR and communication for this family of protocols.

Provisioning a public good. We show that our results for the millionaires problem extend to the classic economic problem of *provisioning a public good*, in which the goal is to determine whether the sum of participants' private values (representing the benefits that these participants would derive from a "public good" such as a bridge or a park) are at least as great as the cost of the good. We observe that, in terms of privacy-approximation ratios, the public-good and millionaires problems are equivalent. We also present upper and lower bounds on the privacy-approximation ratio for the truthful version of provisioning a public good.

Set problems. Finally, we apply our PAR framework to the *intersection* problem and to its decision variant, the *disjointness* problem. From both the privacy perspective and the communication-complexity perspective, these are extremely natural problems to study. The intersection problem has served as a motivating example in the study of privacy-preserving computation for decades, while the disjointness problem plays a central role in the theory and application of communication complexity. We consider three natural protocols that apply to both problems; the objective and subjective PARs are exponential in all cases, but we show that the protocol that is intuitively the best is quantifiably (and significantly) more fair than the others in the sense described below.

Summary of main PAR values. Table 1 contains our results for average-case PAR values for the various problems and protocols that we consider here. (Because of space limitations, we will not restate these results below as individual theorems.) The rows labeled with "All" describe bounds for all protocols for that problem (as reflected by the inequalities). The ratios of the subjective PARs are obtained by dividing the larger subjective PAR by the smaller one; so this ratio is always at least 1. Asymptotic results are for

$k \to \infty$; entries of "—" for bounds on subjective PARs indicate that we do not have results beyond those implied by the PARs for specific protocols.

1.3 Related Work on Privacy-Preserving Computation

Our points of departure in defining approximate privacy are the work of Chor and Kushilevitz [6] on characterization of perfectly privately computable functions and that of Kushilevitz [15] on the communication complexity of perfectly private computation. Starting from the same place, Bar-Yehuda *et al.* [2] 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 maintaining it. In Appendix A below, we show that the formulations in [2] are significantly different from the ones we present here. A full characterization of the relationship between their formulations and ours is an interesting direction for future work (see Sec. 8). As in [2,6,15], we focus on the two-party, deterministic communication model; as outlined in Sec. 8, this is just a first step in a more general research agenda.

There are many formulations of privacy-preserving computation, both perfect 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.

Secure, Multiparty Function Evaluation. The most extensively developed approach to privacy in distributed computation is that of *secure, multiparty function evaluation* (SMFE). Indeed, to achieve (perfect) 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 previously explained by Brandt and Sandholm [5], who applied Chor and Kushilevitz's tools to the study of perfect privacy in auctions, SMFE protocols fall into two main categories (*information-theoretically private protocols* and *multiparty protocols that use cryptography to achieve privacy*) that both have inherent disadvantages from the point of view of mechanism design.

Note that we are using the phrase "perfect privacy" differently from the way it is used in some of the SMFE literature. There, the terms "unconditional," "perfect," and "information-theoretic" are sometimes used interchangeably to describe protocols that preserve input privacy but do not use cryptography (or, more generally, do not rely on unproven complexity-theoretic assumptions). We use "perfectly private" (in contrast with "approximately private") to describe a protocol with privacy-approximation ratio 1 – nothing about the input is revealed except the fact that it is in the pre-image of the output; proving that an SMFE protocol has this property might require one to make an assumption about a cryptographic primitive.

In certain scenarios, the demand for perfect privacy preservation cannot be relaxed; for example, a protocol that revealed two out of nine digits of a social-security number might have a low privacy-approximation ratio but nonetheless be unacceptable in practice. In these cases, if the function cannot be computed in a perfectly privacy-preserving manner without the use of cryptography, there is no choice but to resort to a cryptographic protocol. There is an extensive body of work on cryptographic protocols, and we are not offering our notion of PAR as an extension of that work or an all-purpose substitute for it.

In ad exchanges and other natural e-commerce applications, however, perfect privacy may not be a realistic goal. More generally, we argue that, in e-commerce and other real-world scenarios, privacy preservation should be regarded as one of *several design goals*, along with low computational or communication complexity, protocol simplicity, incentive compatibility, and more. (See, *e.g.*, [17].) Therefore, it is necessary to be able to *quantify* privacy preservation in order to understand the tradeoffs among the different design goals and to obtain "reasonable" (but not necessarily perfect) privacy guarantees. As with any measure of privacy, a value of PAR that indicates good privacy (*i.e.*, PAR close to 1) does not provide guarantees about what an observer might be able to deduce using auxiliary information.

Private Approximations and Approximate Privacy. Here, 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 [3, 8, 14], but it is unrelated to the questions we ask here. Similarly, definitions and techniques from *differential privacy* [7] (see also [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 Sec. 3; we present our main results in Secs. 4–7. Discussion and future directions can be found in Sec. 8.

Many proofs have been omitted because of space limitations; all of them can be found in our two technical reports [9, 10]. As noted above, we do not state here most of the results from Table 1 as individual theorems, but we do state some additional results below that are not included in Table 1.

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; Kushilevitz and Nisan [16] give 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 $f : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}^t$.

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 $f : \{0,1\}^k \times \{0,1\}^k \to \{0,1\}^t$ 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, AND MONOCHROMATICITY) *A* region *in a matrix A is any subset of entries in A (not necessarily a submatrix of A). A* partition *of A is a collection of disjoint regions in A whose union equals A. 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 2. (RECTANGLES, TILINGS, AND REFINEMENTS) *A* rectangle *in a matrix A is a submatrix of A. A* tiling *of a matrix A is a partition of A into rectangles. 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: for every pair of private inputs (x_1, x_2), every communication protocol P for a function f 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*".

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 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 (who 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 3. 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. 3 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. Perfect privacy with respect to 2 is defined similarly.

OBSERVATION 2.1. *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 4. (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 5. (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 6. (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. (For each value in A, the maximal monochromatic region that corresponds to this value is unique; so the ideal monochromatic partition of A is unique.)*

OBSERVATION 2.2. (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 shown by our results, in many environments, perfect privacy can be either impossible or very costly (in

terms of communication complexity) to obtain. To measure a protocol's effect on privacy, relative to the ideal—but perhaps impossible to implement—computation of the outcome of a problem, 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 asses 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 7 (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 α-objective-privacy-preserving in the worst case.

DEFINITION 8. (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 9. (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 10. (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 11. (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 α-privacy-preserving with respect to i in the worst case.

DEFINITION 12. (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 13. (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 14. (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*

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

We say that P is α-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.

At first, it may seem more natural to define average-case PAR as the ratio of the probability masses of the ideal region and the induced rectangle (with respect to D) rather than as the ratio of their cardinalities. We show in Sec. 8.1 below that this intuition is wrong, because it fails to distinguish two very different distributions on inputs with respect to a natural set of protocols.

4. 2ND-PRICE AUCTIONS

4.1 Problem

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^{nd}-price auction [20]: Allocate the item to the highest bidder, and charge him the second-highest bid.

Problem: 2ND-PRICE AUCTION$_k$
Input: $x_1, x_2 \in \{0, \ldots, 2^k - 1\}$ (each as 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 [5] have shown that a perfectly privacy-preserving communication protocol exists for 2ND-PRICE AUCTION$_k$. 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 (in each step first asking bidder 1 and then, if necessary, asking bidder 2). 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 [5] that the English auction is essentially *the only* perfectly privacy-preserving protocol for 2ND-PRICE AUCTION$_k$. 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 Objective Privacy PARs

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 as follows: Use binary search to find a value v 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. Bisection auctions are incentive-compatible in ex-post Nash [12, 13].

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 2ND-PRICE AUCTION$_k$ obtains a subexponential objective PAR:

THEOREM 4.1. (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 $2^{\frac{k}{2}}$.*

By contrast, as shown in Table 1, reasonable privacy guarantees are achievable in the average case.

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 4.2. (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 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 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 BISECTION AUCTION$_{g(k)}$ as follows: Given an instance of 2ND-PRICE AUCTION$_k$, and a integer-valued function $g(k)$ such that $0 \le g(k) \le k$, run the 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 BISECTION AUCTION$_{g(k)}$ ranges from the ascending-price English auction to the BISECTION AUCTION. If we allow a fixed, positive number of bisections ($g(k) = b > 0$), computations show that for $b = 1, 2, 3$ 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 b, 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 average-case objective PAR of the BISECTION AUCTION$_{g(k)}$ with respect to the uniform distribution allows us to do these things; this is given by the following theorem and is also, like the corresponding result for subjective PAR, shown in Table 1.

THEOREM 4.3. *For the* 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 $0 \le c \le k$, and we let $i = k - c$. Figure 2 illustrates the induced tiling for $k = 4$, $c = 2$, and $i = 2$; as in Fig. 1 (which illustrates the ideal tiling for $k = 3$), the input pair $(0, 0)$ is in the upper-left corner. 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).

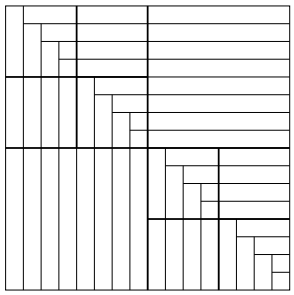

Figure 2: Illustration for the proof of Thm. 4.3

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

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

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

$$\mathsf{PAR} = \frac{1}{2^{2k}} \sum_S |S| \frac{|R^I(S)|}{|S|} = \frac{1}{2^{2k}} \sum_S |R^I(S)|, \qquad (1)$$

where the sums are now over protocol-induced rectangles S (we will simplify notation and write S instead of S^P), and $R^I(S)$ denotes the ideal region that contains the protocol-induced rectangle S. 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 S, let $j_S = 2^k - |R^I(S)|$. Let $a_{c,i}$ be the total number of tiles that appear in the tiling of the k-bit value space induced by the BISECTION AUCTION$_{g(k)}$ with $g(k) = c$, and let $b_{c,i} = \sum_S j_S$ (with this sum being over the protocol-induced tiles in this same partition). Then we may rewrite (1) as

$$\mathsf{PAR}_{c,i} = \frac{1}{2^{2k}} \sum_S (2^k - j_S) = \frac{a_{c,i} 2^k - b_{c,i}}{2^{2k}}. \qquad (2)$$

(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. 2) 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 $a_{c,i} = 2^c \left(2^i(c+2) - 1\right)$. The sum of j_S over protocol-induced rectangles S in the upper-left quadrant is $b_{c,i}$. For a rectangle S in the lower-right quadrant, j_S equals 2^{c+i} plus $j_{S'}$, where S' is the corresponding rectangle in the upper-left quadrant; there are $a_{c,i}$ such S, so the sum of j_S over protocol-induced rectangles S in the upper-left quadrant is $b_{c,i} + a_{c,i} 2^{c+i}$. Finally, the sum of j_S over S in the lower-left quadrant equals $\sum_{h=0}^{2^{c+i}-1} h$ and the sum over S 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 get $b_{c,i} = 2^{c+i-1}\left((1+2^c)\left(-1+2^i\right) + 2^{c+i}c\right)$. Rewriting (2), we obtain

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

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

For the protocols corresponding to values of $g(k)$ ranging from 0 to k (ranging from the ascending-price English auction to the 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 4.4. *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 $\Theta(1 + g(k))$.*

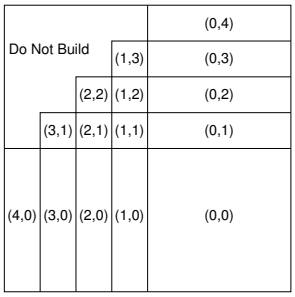

Figure 3: Ideal partition of the value space for TRUTHFUL PUBLIC GOOD$_{k,c}$ with $k = 3$ and $c = 4$.

5. PROVISIONING PUBLIC GOODS

We consider the public-good problem. There are two taxpayers, each with a private value in $\{0, \ldots, 2^k - 1\}$ that represents his benefit from the construction of a public project (public good), *e.g.*, a bridge.[1] The goal of the social planner is to build the public project only if the sum of the taxpayers' values is at least the cost of the bridge, where, as in [1], the cost is set to be $2^k - 1$.

Problem: PUBLIC GOOD$_k$
Input: $x_1, x_2 \in \{0, \ldots, 2^k - 1\}$ (each represented by a k-bit string)
Output: "Build" if $x_1 + x_2 \geq 2^k - 1$, "Do Not Build" otherwise.

We also consider a truthful version of this problem. Now, in addition to determining whether to build the bridge, the government incentivizes truthful disclosure of the private values by requiring taxpayer i to pay $c - \sum_{j \neq i} x_j$ if $\sum_{j \neq i} x_j < c$ but $\sum_i x_i \geq c$. (See, *e.g.*, [19] for a discussion of this approach.) The government should thus learn whether or not to build the bridge and how much, if anything, each taxpayer should pay. The formal description of the function is as follows.

Problem: TRUTHFUL PUBLIC GOOD$_{k,c}$
Input: $c, x_1, x_2 \in \{0, \ldots, 2^k - 1\}$ (each represented by a k-bit string)
Output: "Do Not Build" if $x_1 + x_2 < c$; "Build" and (t_1, t_2) if $x_1 + x_2 \geq c$, where $t_i = c - x_{3-i}$ if $x_{3-i} < c$ and $x_1 + x_2 \geq c$, and $t_i = 0$ otherwise.

The corresponding ideal partition of the value space is shown in Fig. 3, where regions for which the output is "Build" are just labeled with the appropriate value of (t_1, t_2). It is easy to show (via Observation 2.2) that no perfectly privacy-preserving communication protocol exists for either PUBLIC GOOD$_k$ or TRUTHFUL PUBLIC GOOD$_{k,c}$. Therefore, we are interested in the PARs for these problems. Table 1 shows our bounds on the average-case objective PARs for these problems.

6. 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.

[1]This is a discretization of the classic public-good problem, in which the private values are taken from an interval of reals, as in [1, 4].

Problem: THE MILLIONAIRES PROBLEM$_k$
Input: $x_1, x_2 \in \{0, \ldots, 2^k - 1\}$ (encoded as k-bit strings)
Output: the identity of the party with the higher value, *i.e.*, $\arg\max_{i \in \{0,1\}} x_i$ (breaking ties lexicographically).

There cannot be a perfectly privacy-preserving communication protocol for THE MILLIONAIRES PROBLEM$_k$ [15]. Hence, we are interested in the PARs for this well studied problem. The following theorem shows that not only is perfect subjective privacy unattainable for THE MILLIONAIRES PROBLEM$_k$, but a stronger result holds:

THEOREM 6.1. (A WORST-CASE LOWER BOUND ON SUBJECTIVE PAR) *No protocol for* THE MILLIONAIRES PROBLEM$_k$ *has a worst-case subjective PAR less than* $2^{\frac{k}{2}}$.

By contrast to this worst-case result, 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 this protocol applied to THE MILLIONAIRES PROBLEM$_k$ (see Table 1). (As described in [9], the i-induced tilings of the value space for THE MILLIONAIRES PROBLEM$_k$ induced by the BISECTION PROTOCOL may be transformed to obtain the tiling induced by applying the BISECTION AUCTION to 2ND-PRICE AUCTION$_k$.)

Now consider a third party who observes the interaction of the two millionaires. How much can this observer learn about the private information of the two millionaires? Unlike the case of subjective privacy, good PARs are unattainable even in the average case; as shown in Table 1, the average-case objective PAR grows exponentially in k, and the PAR of the BISECTION PROTOCOL grows at the same rate but with a larger constant factor. There are numerous different tilings of the value space that achieve the optimum bound and that can be realized by communication protocols.

Note that the question of whether $x_1 \geq x_2$ is equivalent, in terms of tiling the input space, to the question of whether $x_1 + \hat{x_2} \geq 2^k - 1$ for $\hat{x_2} = 2^k - 1 - x_2$. As discussed in [9], this allows us to use the same approach to this problem as to PUBLIC GOOD$_k$.

7. DISJOINTNESS AND INTERSECTION

7.1 Problem and protocols

We define the DISJOINTNESS$_k$ problem as follows:
Problem: DISJOINTNESS$_k$
Input: Sets $S_1, S_2 \subseteq \{1, \ldots, k\}$
Output: 1 if $S_1 \cap S_2 = \emptyset$, 0 if $S_1 \cap S_2 \neq \emptyset$

Figure 4 shows the ideal monochromatic partition of the 3-bit value space; inputs for which $S_1 \cap S_2 = \emptyset$ are white, and inputs for which $S_1 \cap S_2 \neq \emptyset$ are black.

We define the INTERSECTION$_k$ problem as follows:
Problem: INTERSECTION$_k$
Input: Sets $S_1, S_2 \subseteq \{1, \ldots, k\}$
Output: The set $S_1 \cap S_2$

We will typically encode $S \subseteq \{1, \ldots, k\}$ as a k-bit string so that 1011 encodes $\{1, 2, 4\} \subset \{1, 2, 3, 4\}$, and we will abuse notation and identify $x \in \{0, 1\}^k$ with the subset of $\{1, \ldots, k\}$ that it encodes.

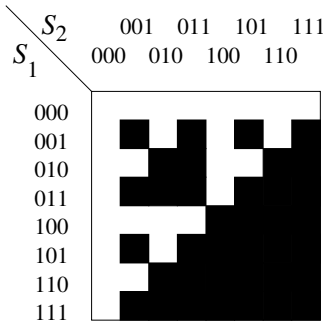

Figure 4: Ideal monochromatic partition for DISJOINTNESS$_k$ **with** $k = 3$.

We now turn to the three protocols that we consider here. With slight variations in their termination, each can be used for both DISJOINTNESS$_k$ and INTERSECTION$_k$.

Trivial protocol. In the trivial protocol, player 1 (without loss of generality) sends his input to player 2, who computes the output and sends this back to player 1. This requires the transmission of $k + 1$ bits for DISJOINTNESS$_k$ and $2k$ bits for INTERSECTION$_k$.

1-first protocol. In the 1-first protocol, player 1 announces a bit, and player 2 replies with his corresponding bit if its value might affect the output; this continues until the output is determined. In detail, player 1 announces the most significant (first) bit of x_1. After player 1 announces his j^{th} bit, if this bit is 0 and $j < k$, then player 1 announces his $(j+1)^{\text{st}}$ bit. If this bit is 0 and $j = k$, then the protocol terminates (with, if computing DISJOINTNESS$_k$, output 1). If this bit is 1, then player 2 announces the value of his j^{th} bit. If player 2's j^{th} bit is also 1, then for DISJOINTNESS$_k$ the protocol terminates with output 0, and for INTERSECTION$_k$ the protocol continues (with $k + 1 - j$ in the output set); if player 2's bit is 0 and $j < k$, then player 1 announces his $(j+1)^{\text{st}}$ bit, while, if $j = k$, the protocol terminates.

Alternating protocol. In the alternating protocol, the role of being the first player to announce the value of a particular bit alternates between the players whenever the first player to announce the value of his j^{th} bit announces "0" (in which case the other player does not announce the value of his corresponding bit). This continues until the output is determined. In detail, player 1 starts by announcing the most significant (first) bit of x_1. After player i announces the value of his j^{th} bit, if this bit is 0 and $j < k$, then the other player announces his $j + 1^{\text{st}}$ bit; if i's j^{th} bit is 0 and $j = k$, the protocol terminates (with output 1 if computing DISJOINTNESS$_k$).

If i's j^{th} bit is 1 and the other player had previously announced his j^{th} bit (which would necessarily be 1, else player i would not be announcing his j^{th} bit), then the protocol terminates with output 0 if computing DISJOINTNESS$_k$, or it continues with the other player announcing his $(j+1)^{\text{st}}$ bit (and with $k+1-j$ being part of the output set). If i's j^{th} bit is 1 and the other player had not previously announced his j^{th} bit, then the other player announces his j^{th} bit; if that bit is 0, then player i proceeds as above. If that bit is 1 and DISJOINTNESS$_k$ is being computed, the protocol terminates with output 0; if the bit is 1 and INTERSECTION$_k$ is being computed, then player i proceeds as above (and $k+1-j$ will

be in the output set). Figure 5 shows the partition of the 1-, 2-, and 3-bit input spaces induced by the alternating protocol for DISJOINTNESS$_k$; each induced rectangle is labeled with the corresponding transcript. (Note that some rectangles appear as non-contiguous regions in the figure.) We note that both the ideal monochromatic partitions and the monochromatic tilings induced by these protocols exhibit recursive structure that is useful in proving PAR results for these problems. (See [10].)

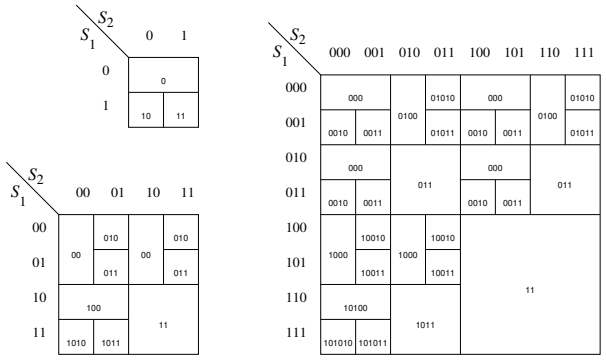

Figure 5: Partition of the value space for $k = 1$ (top left), 2 (bottom left), and 3 (right) induced by the alternating protocol for DISJOINTNESS$_k$.

7.2 PARs

Table 1 shows our PAR results for DISJOINTNESS$_k$ and INTERSECTION$_k$. (The trivial and 1-first protocols induce the same tiling for INTERSECTION$_k$ and thus have the same PARs; so those results are grouped together in Table 1.) Both problems have exponentially bad average-case objective PARs. For INTERSECTION$_k$, all three protocols provide the best possible average-case objective PAR; for DISJOINTNESS$_k$, we conjecture that the bound for the problem can be tightened to equal the average-case objective PAR of these three protocols. Although we do not know the average-case subjective PARs of these problems, for each problem all three protocols also have exponentially bad subjective privacy. We conjecture that this is true for all protocols for each of these problems.

CONJECTURE 7.1. *The average-case subjective PAR of both* DISJOINTNESS$_k$ *and* INTERSECTION$_k$ *with respect to the uniform distribution is exponential in k.*

Intuitively, the alternating protocol seems better than the other two protocols from a privacy perspective, because players take turns disclosing information about their inputs that might not be matched by disclosure about the other player's input. However, the average-case objective and subjective PARs do not reflect this. For both problems, though, we do see a clear difference between the protocols when we consider the ratio between the average-case PARs with respect to each participant. In the trivial and 1-first protocols, the average-case PAR with respect to player 2 is exponentially worse than that with respect to player 1. However, for the alternating protocol, these two PARs differ asymptotically by only a constant factor ($\sqrt{2}$ and $\frac{3}{2}$). Even though the subjective privacy is exponentially bad, the alternating protocol is at least fairer than the others in the sense that it has a similar effect on the players' (subjective) privacy.

8. DISCUSSION & FUTURE DIRECTIONS

8.1 Other Notions of Approximate Privacy

By our definitions, the worst-case/average-case PARs of a protocol are determined by the worst-case/expected value of the expression $\frac{|R^I(\mathbf{x})|}{|R^P(\mathbf{x})|}$, where $R^P(\mathbf{x})$ is the monochromatic rectangle induced by P for input $\mathbf{x}$, and $R^I(\mathbf{x})$ is the monochromatic region containing $A(f)_\mathbf{x}$ in the ideal monochromatic partition of $A(f)$. That is, informally, we are interested in the ratio of the *size* of the ideal monochromatic region for a specific pair of inputs to the *size* of the monochromatic rectangle induced by the protocol for that pair. More generally, we can define worst-case/average-case PARs with respect to a function g by considering the ratio $\frac{g(R^I(\mathbf{x}),\mathbf{x})}{g(R^P(\mathbf{x}),\mathbf{x})}$. Our definitions of PARs set $g(R,\mathbf{x})$ to be the cardinality of R. This captures the intuitive notion of the indistinguishability of inputs that is natural to consider in the context of privacy preservation. Other definitions of PARs may be appropriate in analyzing other notions of privacy. We suggest a few here and discuss them in some more detail in [10]. Further investigation of these and other definitions provides many interesting avenues for future work.

Probability mass. Given a probability distribution D over the parties' inputs, a seemingly natural choice of g is the probability mass. That is, for any region R, $g(R) = Pr_D(R)$, the probability (according to D) that the input corresponds to an entry in R. However, a simple example illustrates that this intuitive choice of g is problematic: Consider a problem for which $\{0,\ldots,n\} \times \{i\}$ is a maximal monochromatic region for $0 \le i \le n-1$. Let P be the communication protocol consisting of a single round in which party 1 reveals whether or not his value is 0; this induces the monochromatic tiling with tiles $\{(0,i)\}$ and $\{(1,i),\ldots,(n,i)\}$ for each i. Now, let D_1 and D_2 be the probability distributions over the inputs $\mathbf{x} = (x_1,x_2)$ such that, for $0 \le i \le n-1$ and $1 \le j \le n$, $Pr_{D_1}[(x_1,x_2) = (0,i)] = \frac{\epsilon}{n}$, $Pr_{D_1}[(x_1,x_2) = (j,i)] = \frac{1-\epsilon}{n^2}$, $Pr_{D_2}[(x_1,x_2) = (0,i)] = \frac{1-\epsilon}{n}$, and $Pr_{D_2}[(x_1,x_2) = (j,i)] = \frac{\epsilon}{n^2}$ for some small $\epsilon > 0$. Intuitively, any reasonable definition of PAR should imply that, for D_1, P provides "bad" privacy guarantees (because w.h.p. it reveals the value of x_1), and, for D_2, P provides "good" privacy (because w.h.p. it reveals little about x_1). In sharp contrast, choosing g to be the probability mass results in the same average-case PAR in both cases.

Other additive functions. In our definition of PAR and in the probability-mass approach, each input $\mathbf{x}$ in a rectangle contributes to $g(R,\mathbf{x})$ in a way that is independent of the other inputs in R. Below, we discuss some natural approaches that violate this condition, but we start by noting that other functions that satisfy this condition may be of interest. For example, taking $g(R,\mathbf{x}) = 1 + \sum_{\mathbf{y} \in R \setminus \mathbf{x}} d(\mathbf{x},\mathbf{y})$, where d is some distance defined on the input space, gives our original definition of PAR when $d(x,y) = 1 - \delta_{\mathbf{x},\mathbf{y}}$ and might capture other interesting definitions (in which indistinguishable inputs that are farther away from $\mathbf{x}$ contribute more to the privacy for $\mathbf{x}$). (The addition of 1 ensures that the ratio $g(R^I,\mathbf{x})/g(R^P,\mathbf{x})$ is defined, but that can be accomplished in other ways if needed.) Importantly, here and below, the notion of distance that is used might not be a

Euclidean metric on the n-player input space $[0, 2^k - 1]^n$. It could instead (and likely would) focus on the problem-specific interpretation of the input space. Of course, there are may possible variations on this (*e.g.*, also accounting for the probability mass).

Maximum distance. We might take the view that a protocol does not reveal much about an input $\mathbf{x}$ if there is another input that is "very different" from $\mathbf{x}$ that the protocol cannot distinguish from $\mathbf{x}$ (even if the total number of things that are indistinguishable from $\mathbf{x}$ under the protocol is relatively small). For some distance d on the input space, we might than take g to be something like $1 + \max_{\mathbf{y} \in R \setminus \{\mathbf{x}\}} d(\mathbf{y}, \mathbf{x})$.

Plausible deniability. One drawback to the maximum-distance approach is that it does not account for the probability associated with inputs that are far from $\mathbf{x}$ (according to a distance d) and that are indistinguishable from $\mathbf{x}$ under the protocol. While there might be an input $\mathbf{y}$ that is far away from $\mathbf{x}$ and indistinguishable from $\mathbf{x}$, the probability of $\mathbf{y}$ might be so small that the observer feels comfortable assuming that $\mathbf{y}$ does not occur. A more realistic approach might be one of "plausible deniability." This makes use of a plausibility threshold—intuitively, the minimum probability that the "far away" inputs(s) (which is/are indistinguishable from $\mathbf{x}$) must be assigned in order to "distract" the observer from the true input $\mathbf{x}$. This threshold might correspond to, *e.g.*, "reasonable doubt" or other levels of certainty. We then consider how far we can move away from $\mathbf{x}$ while still having "enough" mass (*i.e.*, more than the plausibility threshold) associated with the elements indistinguishable from $\mathbf{x}$ that are still farther away. We could then take g to be something like $1+\max\{d_0 | Pr_D(\{\mathbf{y} \in R | d(\mathbf{y}, \mathbf{x}) \geq d_0\})/Pr_D(R) \geq t\}$; other variations might focus on mass that is concentrated in a particular direction from $\mathbf{x}$. (In quantifying privacy, we would expect to only consider those R with positive probability, in which case dividing by $Pr_D(R)$ would not be problematic.) Here we use $Pr_D(R)$ to normalize the weight that is far away from x before comparing it to the threshold t; intuitively, an observer would know that the value is in the same region as x, and so this seems to make the most sense.

Relative rectangle size. One observation is that a bidder likely has a very different view of an auctioneer's being able to tell (when some particular protocol is used) whether his bid lies between 995 and 1005 than he does of the auctioneer's being able to tell whether his bid lies between 5 and 15. In each case, however, the bids in the relevant range are indistinguishable under the protocol from 11 possible bids. In particular, the privacy gained from an input's being distinguishable from a fixed number of other inputs may (or may not) depend on the context of the problem and the intended interpretation of the values in the input space. This might lead to a choice of g such as $diam_d(R)/|\mathbf{x}|$, where $diam_d$ is the diameter of R with respect to some distance d and $|\mathbf{x}|$ is some (problem-specific) measure of the size of $\mathbf{x}$ (*e.g.*, bid value in an auction). Numerous variations on this are natural and may be worth investigating.

Information-theoretic approaches. Information-theoretic approaches using conditional entropy are also natural to consider when studying privacy, and these have been used in various settings. Most relevantly, Bar-Yehuda *et al.* [2] defined multiple measures based on the conditional mutual information about one player's value (viewed as a random variable) revealed by the protocol trace and knowledge of the other player's value. It would also be natural to study objective-PAR versions using the entropy of the random variable corresponding to the (multi-player) input conditioned only on the protocol output (and not the input of any player). Such approaches might facilitate the comparison of privacy between different problems.

8.2 Open Questions

There are many interesting directions for future research:

- Explore other definitions of PARs, along the lines discussed in the previous subsection.
- We have shown that, for both 2ND-PRICE AUCTION$_k$ and THE MILLIONAIRES PROBLEM$_k$, reasonable average-case PARs with respect to the uniform distribution are achievable. Prove or disprove the conjecture that our upper bounds for these problems extend to *all* possible distributions over inputs.
- Prove lower bounds on the average-case PARs of problems for which we do not have them.
- Apply our PAR framework to other functions, and extend the framework to n-party communication.
- Starting from the same place that we did, namely [6, 15], Bar-Yehuda *et al.* [2] provided three quantifications of privacy. We show in Appendix A below that the formulation in [2] is not equivalent to ours, but there is more to do along these lines. The definition in [2] that seems most relevant to the study of privacy-approximation ratios is the notion of *h-privacy*. Determine when and how it is possible to express PARs in terms of h-privacy and *vice versa*.

9. ACKNOWLEDGEMENTS

We are grateful to the anonymous reviewers and to audiences at University Residential Centre of Bertinoro, Boston University, DIMACS, Northwestern, University of Massachusetts, and Rutgers for helpful questions and feedback.

10. REFERENCES

[1] Moshe Babaioff, Liad Blumrosen, Moni Naor, and Michael Schapira. Informational overhead of incentive compatibility. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 88–97, 2008.

[2] Reuven Bar-Yehuda, Benny Chor, Eyal Kushilevitz, and Alon Orlitsky. Privacy, additional information, and communication. *IEEE Transactions on Information Theory*, 39(6):55–65, 1993.

[3] 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.

[4] Liad Blumrosen and Noam Nisan. Auctions with severely bounded communications. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*, pages 406–415, 2002.

[5] Felix Brandt and Tuomas Sandholm. On the existence of unconditionally privacy-preserving auction protocols. *ACM Trans. Inf. Syst. Secur.*, 11(2):1–21, 2008.

[6] Benny Chor and Eyal Kushilevitz. A zero-one law for boolean privacy. *SIAM J. Discrete Math*, 4(1):36–47, 1991.

[7] Cynthia Dwork. Differential privacy. In *Proceedings of the International Colloquium on Automata, Languages and Programming*, pages 1–12, 2006.

[8] 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.

[9] Joan Feigenbaum, Aaron D. Jaggard, and Michael Schapira. Approximate privacy: Foundations and quantification. DIMACS technical report 2009-14 and arXiv:0910.5714, 2009.

[10] Joan Feigenbaum, Aaron D. Jaggard, and Michael Schapira. Approximate privacy: PARs for set problems. DIMACS technical report 2010-01 and arXiv:1001.3388, 2010.

[11] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally utility-maximizing privacy mechanisms. In *Proceedings of the ACM Symposium on Theory of Computing*, pages 351–360, 2009.

[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] S. Muthukrishnan. Challenges in Designing Ad Exchanges. Talk at the *NSF Workshop on Research Issues at the Interface of Computer Science and Economics*, Cornell University, September 2009.

[18] 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.

[19] Noam Nisan. Introduction to mechanism design (for computer scientists). In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, **Algorithmic Game Theory**, chapter 9. Cambridge University Press, 2007.

[20] William Vickrey. Counterspeculation, auctions and competitive sealed tenders. *J. of Finance*, 16(1):8-37, 1961.

[21] 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.

[22] Andrew C. Yao. Protocols for secure computation. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*, pages 160–164, 1982.

APPENDIX

A. CONTRAST WITH BAR-YEHUDA *ET AL.* [2]

Although the work presented here and that of Bar-Yehuda *et al.* [2] have common roots, there are significant differences in what the two frameworks capture. Specifically:

First, the results in [2] deal with what can be learned by a party who knows one of the inputs. By contrast, our notion of objective PAR captures the effect of a protocol on privacy with respect to an external observer who does not know either of the players' private values.

Second, and more importantly, the framework of [2] does not address the size of monochromatic regions. As illustrated by the following example, the ability to do so is necessary to capture the effects of protocols on interesting aspects of privacy that are captured by our definitions of PAR.

Consider the function $f : \{0, \ldots, 2^n - 1\} \times \{0, \ldots, 2^n - 1\} \to \{0, \ldots, 2^{n-2}\}$ defined by $f(x, y) = floor(\frac{x}{2})$ if $x < 2^{n-1}$ and $f(x, y) = 2^{n-2}$ otherwise. Consider the following two protocols for f: In P, player 1 announces his value x if $x < 2^{n-1}$ and otherwise sends 2^{n-1} (which indicates that $f(x, y) = 2^{n-2}$); in Q, player 1 announces $floor(\frac{x}{2})$ if $x < 2^n - 1$ and x if $x = 2^n - 1$. Observe that both P and Q induce $2^{n-1} + 1$ rectangles.

Intuitively, P and Q have different effects on privacy. For half of the input pairs, P reduces by a factor of 2 the number of input pairs from which they are indistinguishable while not affecting the indistinguishability of the other half of the input pairs. Q does not affect the indistinguishability of the input pairs affected by P, but it does reduce the number of input pairs indistinguishable from any pair in which $x \geq 2^{n-1}$ by at least a factor of 2^{n-2}.

Our notion of PAR is able to capture the different effects on privacy of the protocols P and Q. (With respect to the uniform distribution, the average-case objective PAR of P is constant, and that of Q is exponential in n.) By contrast, the three quantifications of privacy from [2]—I_c, I_i, and I_{c-i}—do not distinguish between these two protocols; we now sketch the arguments for this claim.

For each protocol, any function h for which the protocol is weakly h-private must take at least $2^{n-1}+1$ different values. This bound is tight for both P and Q. Thus, I_c is unable to distinguish between the effects of P and Q on f.

The number of rectangles induced by P that intersect each row and column equals the number induced by Q. Considering the geometric interpretation of I_P and I_Q, as well as the discussion in Sec. VII.A of [2], we see that I_i and I_{c-i} (applied to protocols) cannot distinguish between the effects of P and Q on f.