

The Power of Shared Randomness in Uncertain Communication

Badih Ghazi * Madhu Sudan[†]

May 2, 2017

Abstract

In a recent work (Ghazi et al., SODA 2016), the authors with Komargodski and Kothari initiated the study of *communication with contextual uncertainty*, a setup aiming to understand how efficient communication is possible when the communicating parties imperfectly share a huge context. In this setting, Alice is given a function f and an input string x , and Bob is given a function g and an input string y . The pair (x, y) comes from a known distribution μ and f and g are guaranteed to be close under this distribution. Alice and Bob wish to compute $g(x, y)$ with high probability. The lack of agreement between Alice and Bob on the function that is being computed captures the uncertainty in the context. The previous work showed that any problem with one-way communication complexity k in the standard model (i.e., without uncertainty¹) has *public-coin* communication at most $O(k(1 + I))$ bits in the uncertain case, where I is the mutual information between x and y . Moreover, a lower bound of $\Omega(\sqrt{I})$ bits on the public-coin uncertain communication was also shown.

However, an important question that was left open is related to the power that public randomness brings to uncertain communication. Can Alice and Bob achieve efficient communication amid uncertainty without using public randomness? And how powerful are public-coin protocols in overcoming uncertainty? Motivated by these two questions:

- We prove the first separation between private-coin uncertain communication and public-coin uncertain communication. Namely, we exhibit a function class for which the communication in the standard model and the public-coin uncertain communication are $O(1)$ while the private-coin uncertain communication is a growing function of n (the length of the inputs). This lower bound (proved with respect to the uniform distribution) is in sharp contrast with the case of public-coin uncertain communication which was shown by the previous work to be within a constant factor from the certain communication. This lower bound also implies the first separation between public-coin uncertain communication and deterministic uncertain communication. Interestingly, we also show that if Alice and Bob *imperfectly share* a sequence of random bits (a setup weaker than public randomness), then achieving a constant blow-up in communication is still possible.
- We improve the lower-bound of the previous work on public-coin uncertain communication. Namely, we exhibit a function class and a distribution (with mutual information $I \approx n$) for which the one-way certain communication is k bits but the one-way public-coin uncertain communication is at least $\Omega(\sqrt{k} \cdot \sqrt{I})$ bits.

Our proofs introduce new problems in the standard communication complexity model and prove lower bounds for these problems. Both the problems and the lower bound techniques may be of general interest.

*MIT, CSAIL, Cambridge MA 02139. badih@mit.edu.

[†]Harvard John A. Paulson School of Engineering and Applied Sciences. madhu@cs.harvard.edu.

¹in other words, under the promise that $f = g$.

Contents

1	Introduction	1
1.1	Our Contributions	2
1.2	Implications	5
2	Overview of Proofs	6
2.1	Overview of Proof of Theorem 1.1	6
2.2	Overview of Proof of Theorem 1.3	7
3	Preliminaries	9
4	Construction for Private-Coin Uncertain Protocols	10
4.1	Proof of Lemma 4.3	12
4.2	Proof of Part (iii) of Theorem 1.1	13
4.2.1	Description of Protocol Π'	13
4.2.2	The $\text{stretch}_{r,a}$ Procedure	13
4.2.3	Analysis of Protocol Π'	14
5	Construction for Public-Coin Uncertain Protocols	17
5.1	Proof Preliminaries for Subsections 5.2 and 5.3	18
5.2	Simulation Protocol	20
5.3	Proof of Lemma 5.2	24
6	Proof of Theorem 1.2	25
7	Open Questions	27
A	Two-Dimensional Berry-Esseen Theorem for Independent Random Variables	29

1 Introduction

In many forms of communication (e.g., human, computer-to-computer), the communicating parties share some context (e.g, knowledge of a language, operating system, communication protocol, encoding/decoding mechanisms.). This context is usually a) *huge* and b) *imperfectly shared* among the parties. Nevertheless, in human communication, very efficient communication is usually possible. Can we come up with a mathematical analogue of this phenomenon of efficient communication based on a huge but imperfectly shared context? Motivated by this general question, the study of “communication amid uncertainty” has been the subject of a series of recent work starting with Goldreich, Juba and Sudan [JS08, GJS12] followed by [JKKS11, JS11, JW13, HS14, CGMS15]. While early works were very abstract and general, later works (starting with Juba, Kalai, Khanna and Sudan [JKKS11]) tried to explore the ramifications of uncertainty in Yao’s standard communication complexity model [Yao79]. In particular, the more recent works relax the different pieces of context that were assumed to be *perfectly shared* in Yao’s model, such as shared randomness [CGMS15], and in a recent work of the authors with Komargodski and Kothari the *function being computed* [GKKS16].

Specifically, [GKKS16] study the following functional notion of uncertainty in communication. Their setup builds on – and generalizes – Yao’s classical model of (distributional) communication complexity, where Alice has an input x and Bob has an input y , with (x, y) being sampled from a distribution μ . Their goal is to communicate minimally so as to compute some function $g(x, y)$ (with high probability over the choice of (x, y)). The understated emphasis of the model is that for many functions g , the communication required is much less than the lengths of x or y , the entropy of x or y or even the conditional entropy of x given y .

The question studied by [GKKS16] is: How much of this gain in communication is preserved when the communicating parties do not exactly agree on the function being computed? (We further discuss the importance of this question in Section 1.2.) This variation of the problem is modelled as follows: Alice is given a Boolean function f and an input string x , and Bob is given a Boolean function g and an input string y where (x, y) is sampled from a known distribution μ as before, and (f, g) is chosen (adversarially) from a known class \mathcal{F} of pairs of functions that are close in terms of the Hamming distance Δ_μ (weighted according to μ). Alice and Bob wish to compute $g(x, y)$. Alice’s knowledge of the function f (which is close but not necessarily equal to g) captures the uncertainty in the knowledge of the context.

We define the *public-coin uncertain communication complexity* $\text{PubCCU}_\epsilon^\mu(\mathcal{F})$ as the minimum length of a two-way public-coin protocol whose output is correct with probability at least $1 - \epsilon$ over its internal randomness and that of (x, y) . We similarly define the *private-coin uncertain communication complexity* $\text{PrivCCU}_\epsilon^\mu(\mathcal{F})$ by restricting to private-coin protocols. Clearly, $\text{PubCCU}_\epsilon^\mu(\mathcal{F}) \leq \text{PrivCCU}_\epsilon^\mu(\mathcal{F})$. The quantities $\text{owPubCCU}_\epsilon^\mu(\mathcal{F})$ and $\text{owPrivCCU}_\epsilon^\mu(\mathcal{F})$ are similarly defined by restricting to one-way protocols.²

The previous work ([GKKS16]) gave an upper bound on $\text{owPubCCU}_\epsilon^\mu(\mathcal{F})$ whenever \mathcal{F} consists of functions g whose one-way distributional complexity is small. More precisely, denote by $\text{owCC}_\epsilon^\mu(g)$ the one-way communication complexity of g in the standard distributional model.³

²Note that the uncertain model is clearly a generalization of Yao’s model which corresponds to the particular case where $\mathcal{F} = \{(f, f)\}$ for some fixed function f . On the other hand, the uncertain model can also be viewed as a particular case of Yao’s model via an exponential blow-up in the input size. For more on this view (which turns out to be ineffective in our setup), we refer the reader to Note 1.2 at the end of this section.

³By the “easy direction” of Yao’s min-max principle, we can without loss of generality consider deterministic

Namely, $\text{owCC}_\epsilon^\mu(g)$ is the minimum length of a one-way deterministic protocol computing g with probability at least $1 - \epsilon$ over the randomness of (x, y) . Then, [GKKS16] showed that if \mathcal{F} consists of pairs (f, g) of functions that are at distance δ , and if $\text{owCC}_\epsilon^\mu(f), \text{owCC}_\epsilon^\mu(g) \leq k$, then for every positive θ , $\text{owPubCCU}_{\epsilon+2\delta+\theta}^\mu(\mathcal{F}) \leq O_\theta(k \cdot (1 + I(x; y)))$, where $I(x; y)$ denotes the mutual information between x and y .⁴ Note that if μ is a product distribution and if we let the parameter θ be a small constant, then the blow-up in communication is only a constant factor. However, the protocol of [GKKS16] crucially uses public randomness, and one of the main motivations behind this work is to understand how large the blow-up would be in the case where Alice and Bob have access to weaker types of randomness (or no randomness at all).

We point out that understanding the type of randomness that is needed in order to cope with uncertainty is a core question in the setup of communication with contextual uncertainty: *If Alice and Bob do not (perfectly) agree on the function being computed, why can we assume that they (perfectly) agree on the shared randomness?*

1.1 Our Contributions

We prove several results about the power of shared randomness in uncertain communication.

Private and Imperfectly Shared Randomness Our first result (Theorem 1.1) shows that private-coin protocols are much weaker than public-coin protocols in the setup of communication with contextual uncertainty. Far from obtaining a constant factor blow-up in communication, private-coin protocols incur an increase that is a growing function of n when dealing with uncertainty.

Let $\mathcal{U} \triangleq \mathcal{U}_{2^n}$ be the uniform distribution on $\{0, 1\}^{2^n}$. For positive integers t and n , we define $\log^{(t)}(n)$ by setting $\log^{(1)}(n) = \log n$, and $\log^{(i)}(n) = \max(\log \log^{(i-1)}(n), 1)$ for all $i \in \{2, \dots, t\}$.

Theorem 1.1 (Lower-bound on private-coin uncertain protocols). *For every sufficiently small $\delta > 0$, there exist a positive integer $\ell \triangleq \ell(\delta)$ and a function class $\mathcal{F} \triangleq \mathcal{F}_\delta$ such that*

(i) *For each $(f, g) \in \mathcal{F}$, we have that $\Delta_{\mathcal{U}}(f, g) \leq \delta$.*

(ii) *For each $(f, g) \in \mathcal{F}$, we have that $\text{owCC}_0^{\mathcal{U}}(f), \text{owCC}_0^{\mathcal{U}}(g) \leq \ell$.*

(iii) *For every $\eta > 0$ and $\epsilon \in (4\delta, 0.5]$, we have that $\text{PrivCCU}_{\epsilon/2-2\delta-\eta}^{\mathcal{U}}(\mathcal{F}) = \Omega(\eta^2 \cdot \log^{(t)}(n))$ for some positive integer $t = \Theta((\epsilon/\delta)^2)$.*

In Theorem 1.1, the inputs x and y are binary strings of length n and \mathcal{F} is a family of pairs of functions, which each function mapping $\{0, 1\}^n \times \{0, 1\}^n$ to $\{0, 1\}$. Also, the parameter η can possibly depend on n . We point out that Theorem 1.1 also implies the first separation between *deterministic* uncertain protocols and public-coin uncertain protocols⁵.

(instead of public-coin) protocols when defining $\text{owCC}_\epsilon^\mu(g)$. We point out that this is not true in the uncertain case.

⁴One interpretation of the dependence of the communication in the uncertain setup on the mutual information $I(x; y)$ is that the players are better able to use the correlation of their inputs in the standard case than in the uncertain case.

⁵This uses the fact that private-coin communication complexity is no larger than deterministic communication complexity, both in the certain and uncertain setups.

Note. We point out that the relative power of private-coin and public-coin protocols in the uncertain model is both conceptually and technically different from the standard model. Specifically, the randomness is potentially used in the standard model in order to fool an adversary selecting the input pair (x, y) , whereas in the uncertain model, it is used to fool an adversary selecting the pair (f, g) of functions that are promised to be close. This promise makes the task of proving lower bounds against private-coin protocols in the uncertain model (e.g., Theorem 1.1) significantly more challenging than in the standard model.⁶ Moreover, a well-known theorem due to Newman [New91] shows that in the standard model, any public-coin protocol can be simulated by a private-coin protocol while increasing the communication by an additive $O(\log n)$ bits. By contrast, there is no known analogue of Newman’s theorem in the uncertain case!

Note. The construction that we use to prove Theorem 1.1 cannot give a separation larger than $\Theta(\log \log n)$. Thus, showing a separation of $\omega(\log \log n)$ between private-coin and public-coin protocols in the uncertain case would require a new construction. For more details, see Note 4.2.3.

In light of Theorem 1.1, it is necessary for Alice and Bob to share some form of randomness in order to only incur a constant blow-up in communication for product distributions. Fortunately, it turns out that it is not necessary for Alice and Bob to *perfectly* share a sequence of random coins. If Alice is given a uniform-random string r of bits and Bob is given a string r' obtained by independently flipping each coordinate of r with probability 0.49, then efficient communication is still possible!

More formally, for $\rho \in [0, 1]$, define $\text{owlsrCCU}_{\epsilon, \rho}^{\mu}(\mathcal{F})$ in the same way that we defined $\text{owPubCCU}_{\epsilon}^{\mu}(\mathcal{F})$ except that instead of Alice and Bob having access to public randomness, Alice will have access to a sequence r of independent uniformly-random bits, and Bob will have access to a sequence r' of bits obtained by independently flipping each coordinate of r with probability $(1 - \rho)/2$. Note that this setup of imperfectly shared randomness interpolates between the public randomness and private randomness setups, i.e., $\text{owlsrCCU}_{\epsilon, 1}^{\mu}(\mathcal{F}) = \text{owPubCCU}_{\epsilon}^{\mu}(\mathcal{F})$ and $\text{owlsrCCU}_{\epsilon, 0}^{\mu}(\mathcal{F}) = \text{owPrivCCU}_{\epsilon}^{\mu}(\mathcal{F})$.

Theorem 1.2 (Uncertain protocol using imperfectly shared randomness). *Let $\rho \in (0, 1]$ and μ be a product distribution. Let \mathcal{F} consist of pairs (f, g) of functions with $\Delta_{\mu}(f, g) \leq \delta$, and $\text{owCC}_{\epsilon}^{\mu}(f), \text{owCC}_{\epsilon}^{\mu}(g) \leq k$. Then, for every positive θ , $\text{owlsrCCU}_{\epsilon+2\delta+\theta, \rho}^{\mu}(\mathcal{F}) \leq O_{\theta}(k/\rho^2)$.*

The *imperfectly shared randomness* model in Theorem 1.2 was recently independently introduced (in the setup of communication complexity) by Bavarian, Gavinsky and Ito [BGI14] and by Canonne, Guruswami, Meka and Sudan [CGMS15] (and it was further studied in [GKS16]). Moreover, our proof of Theorem 1.2 is based on combining the uncertain protocol of [GKKS16] and the locality-sensitive-hashing based protocol of [CGMS15].

We point out that Theorem 1.2 also holds for more general i.i.d. sources of correlated randomness than the one described above. More precisely, for i.i.d. (not necessarily binary) sources of (imperfectly) shared randomness with *maximal correlation*⁷ ρ , the work of Witsenhausen [Wit75] along with the protocols of [CGMS15] and [GKKS16] imply an uncertain protocol with $O_{\theta}(k/\rho^2)$ bits of communication.

⁶In particular, the diagonalization-based arguments that imply a separation between the public-coin and the private-coin communication complexities of the Equality function in the standard model completely fail when we impose such a promise.

⁷The *maximal correlation* of a pair (X, Y) of random variables (with support $\mathcal{X} \times \mathcal{Y}$) is defined as $\rho(X, Y) \triangleq \sup \mathbb{E}[F(X)G(Y)]$ where the supremum is over all functions $F : \mathcal{X} \rightarrow \mathbb{R}$ and $G : \mathcal{Y} \rightarrow \mathbb{R}$ with $\mathbb{E}[F(X)] = \mathbb{E}[G(Y)] = 0$ and $\text{Var}[F(X)] = \text{Var}[G(Y)] = 1$. It is not hard to show that the binary source of imperfectly shared randomness defined in the paragraph preceding Theorem 1.2 has maximal correlation ρ .

Public Randomness We now turn to our next result where we consider the dependence of the upper bound of [GKKS16] on the mutual information $I \triangleq I(X; Y)$ in the case of public-coin protocols. The previous work [GKKS16] proved a lower bound of $\Omega(\sqrt{I})$ on this dependence, but their lower-bound does not grow with k . We improve this lower bound to $\Omega(\sqrt{k} \cdot \sqrt{I})$.

Theorem 1.3 (Improved lower-bound on public-coin uncertain protocols). *For every sufficiently small $\delta > 0$ and every positive integers k, n such that $k = o(\exp(\sqrt{n}))$, there exist an input distribution μ on input pairs $(X, Y) \in \{0, 1\}^{k \cdot n} \times \{0, 1\}^{k \cdot n}$ with mutual information $I \approx k \cdot n$ and a function class $\mathcal{F} \triangleq \mathcal{F}_{\delta, k, n}$ such that⁸*

(i) For each $(f, g) \in \mathcal{F}$, we have that $\Delta_\mu(f, g) \leq \delta$.

(ii) For each $(f, g) \in \mathcal{F}$, we have that $\text{owCC}_0^\mu(f), \text{owCC}_0^\mu(g) \leq k$.

(iii) $\text{owPubCCU}_\epsilon^\mu(\mathcal{F}) = \Omega(\sqrt{k} \cdot \sqrt{I})$ for some absolute constant $\epsilon > 0$ independent of δ .

As will be explained in detail in Section 2.2, the proof of Theorem 1.3 is based on an extension of the lower bound construction of [GKKS16], which is then analyzed using different techniques.

Note. *The construction that we use to prove Theorem 1.3 cannot give a lower bound larger than $\tilde{\Theta}(\sqrt{k} \cdot \sqrt{I})$. Thus, improving on the lower bound in Theorem 1.3 by more than logarithmic factors in k and I would require a new construction.*

New Communication Problems Our lower bounds in Theorems 1.1 and 1.3 are derived by defining new problems in *standard* communication complexity (i.e., without uncertainty) and proving lower bounds for these problems. We describe these problems and our results on these next.

The construction that we use to prove Theorem 1.1 requires us to understand the following “subset-majority with side information” setup. Alice is given a subset $S \subseteq [n]$ and a string $x \in \{\pm 1\}^n$, and Bob is given a subset $T \subseteq [n]$ and a string $y \in \{\pm 1\}^n$. The subsets S and T are adversarially chosen but are promised to satisfy $S \subseteq T$, $|T| = \ell$ and $|T \setminus S| \leq \delta \cdot \ell$ for some fixed parameters ℓ and δ . The strings x and y are chosen independently and uniformly at random. Alice and Bob wish to compute the function $\text{SubsetMaj}((S, x), (T, y)) \triangleq \text{Sign}(\sum_{i \in T} x_i y_i)$. In words, $\text{SubsetMaj}((S, x), (T, y))$ is equal to 0 if x and y differ on a majority of the coordinates in subset T , and 1 otherwise. Note that S does not directly appear in the definition of the function SubsetMaj but it can serve as useful side-information for Alice.⁹ What is the private-coin communication complexity of computing SubsetMaj on every (S, T) -pair satisfying the above promise and with high probability over the random choice of (x, y) and over the private randomness? We prove the following (informally stated) lower bound.

Lemma 1.4. *Any private-coin protocol computing SubsetMaj on every (S, T) -pair satisfying the promise and with high probability over the random choice of (x, y) and over the private randomness should communicate at least $\log^{(t)}(n)$ bits for some positive integer t that depends on δ and the error probability.*

⁸We note that $I \approx k \cdot n$ means that $I/(k \cdot n) \rightarrow 0$ as $n \rightarrow \infty$.

⁹Note that we could have alternatively defined SubsetMaj in terms of S , and let T serve as the potentially useful side-information. Our lower bound would also apply to this setup.

The proof of Theorem 1.3 is based on a construction that leads to the question described next regarding the communication complexity of a particular block-composed function. Namely, consider the following “majority composed with subset-parity with side information” setup. Alice is given a sequence of subsets $S \triangleq (S^{(i)} \subseteq [n])_{i \in [k]}$ and a sequence of strings $x \triangleq (x^{(i)} \in \{0, 1\}^n)_{i \in [k]}$, and Bob is given a sequence of subsets $T \triangleq (T^{(i)} \subseteq [n])_{i \in [k]}$ and a sequence of strings $y \triangleq (y^{(i)} \in \{0, 1\}^n)_{i \in [k]}$. We consider the following distribution μ on $((S, x), (T, y))$. Independently for each $i \in [k]$, we sample $((S^{(i)}, x^{(i)}), (T^{(i)}, y^{(i)}))$ as follows: $S^{(i)}$ is a uniform-random subset and $T^{(i)}$ is an ϵ -noisy¹⁰ copy of $S^{(i)}$, and independently $x^{(i)}$ is a uniform-random string and $y^{(i)}$ is an ϵ -noisy copy of $x^{(i)}$. Here, ϵ is a positive parameter that can depend on n and k . Alice and Bob wish to compute the function $\text{Maj} \circ \text{SubsetParity}((S, x), (T, y)) \triangleq \text{Sign}(\sum_{i=1}^k (-1)^{\langle T^{(i)}, x^{(i)} \oplus y^{(i)} \rangle})$ where $T^{(i)}$ denotes both the subset and its 0/1 indicator vector, the inner product is over \mathbb{F}_2 , and $x^{(i)} \oplus y^{(i)}$ is the coordinate-wise XOR of $x^{(i)}$ and $y^{(i)}$. What is the communication complexity of computing $\text{Maj} \circ \text{SubsetParity}$ with high probability over the distribution μ ? We prove the following lower bound.

Lemma 1.5. *Any 1-way protocol computing $\text{Maj} \circ \text{SubsetParity}$ with high probability over the distribution μ should communicate $\Omega(k \cdot \epsilon \cdot n)$ bits.*

In Section 2.1, we outline the proof of Theorem 1.1 and explain how it leads to the setup of Lemma 1.4 and how we prove Lemma 1.4. In Section 2.2, we outline the proof of Theorem 1.3 and explain how it leads to the setup of Lemma 1.5 and how we prove Lemma 1.5.

Before doing so, we discuss some conceptual implications of our results.

1.2 Implications

Functional uncertainty models much of the day-to-day interactions among humans, where a person is somewhat aware of the objectives of the other person she is interacting with, but do not know them precisely. Neither person typically knows exactly what aspects of their own knowledge may be relevant to the interaction, yet they do manage to have a short conversation. This is certainly a striking phenomenon, mostly unexplained in mathematical terms. This line of works aims to explore such phenomena. It is important to understand what mechanisms may come into play here, and what features play a role. Is the ability to make random choices important? Is shared information crucial? Is there a particular measure of distance between functions that makes efficient communication feasible? In order to understand such questions, one first needs to have a ground-level understanding of communication with functional uncertainty. This work tackles several basic questions that remain unexplored.

An ideal model for communication would only assume a constant amount of perfectly shared context between the sender and receiver, such as the knowledge of an encoding/decoding algorithm, one universal Turing machine, etc. Solutions to most interesting communication problems seem to assume a shared information which grows with the length of the inputs. Recent work showed that in many of these scenarios some assumptions about the shared context can be relaxed to an imperfect sharing, but these results are often brittle and break when two or more contextual elements are simultaneously assumed to be imperfectly shared. Our work raises the question of whether *imperfectly shared randomness* would be sufficient to overcome *functional uncertainty*. We show that this is indeed the case for product distributions, but the loss for non-product distributions

¹⁰This means that the indicator vector of $T^{(i)}$ is obtained by independently flipping each coordinate of the indicator vector of $S^{(i)}$ with probability ϵ .

might be much larger (for this and other open questions, we refer the reader to our conclusion Section 7). Such results highlight the delicate nature of the role of shared context in communication. They beg for a more systematic study of communication which at the very least should be able to mimic the aims, objectives and phenomena encountered in human communication.

Note. *As mentioned in Footnote 2, the uncertain model is clearly a generalization of Yao’s model. Strictly speaking, the uncertain model can also be viewed as a particular case of Yao’s model by regarding the function(s) that is being computed as part of the inputs of Alice and Bob, which results in an exponential blow-up in the input-size. This latter view turns out to be fruitless for our purposes. Indeed, from this perspective, all the different well-studied communication functions (such as Equality, Set Disjointness, Pointer Jumping, etc.) are regarded as special cases of one “universal function”! More importantly, this view completely blurs the distinction between the goal of the communication (i.e., the function to compute) and the inputs of the parties. On a technical level, it does not simplify the task of proving the lower bounds in Theorems 1.1 and 1.3 in any way since it does not capture the promise that the two functions (given to Alice and Bob) are close in Hamming distance. Thus, in the rest of this paper, we stick to the former view and use the expressions “uncertain model” and “standard model” to refer to the setups with and without uncertainty, respectively.*

2 Overview of Proofs

2.1 Overview of Proof of Theorem 1.1

Reduction to Lemma 1.4. In order to prove Theorem 1.1, we need to devise a function class for which circumventing the uncertainty is much easier using public randomness than using private randomness. One general setup in which Bob can leverage public randomness to resolve some uncertainty regarding Alice’s knowledge is the following “small-set intersection” problem. Assume that Alice is given a subset $S \subseteq [n]$, and Bob is given a subset $T \subseteq [n]$ such that T contains S and $|T| = \ell$, where we think of ℓ as being a large constant. Here, Bob knows that Alice has a subset of his own T but he is uncertain which subset Alice has. Using public randomness, a standard 1-way hashing protocol communicating $\tilde{O}(\ell)$ bits allows Bob to determine S with high probability. On the other hand, using only private randomness, the communication complexity of this task is $\Theta(\log \log n)$ bits.

With the above general setup in mind, we consider functions f_S indexed by small subsets S of coordinates on which they depend. Since we want the functions f_S and f_T to be close in Hamming distance, we enforce $|T \setminus S|$ to be small for every pair (f_S, f_T) of functions in our class, and we let each function f_S be “noise-stable”. Since we want our function f_S to genuinely depend on all coordinates in S , the majority function $f_S(x, y) = \text{Sign}(\sum_{i \in S} x_i y_i)$ for $x, y \in \{\pm 1\}^n$ arises as a natural choice. We also let x and y be independent uniform-random strings. In this case, it can be seen that if $|T \setminus S|$ is a small constant fraction of $|T|$, then the quadratic polynomials $\sum_{i \in S} x_i y_i$ and $\sum_{i \in T} x_i y_i$ behave like standard Gaussians with correlation close to 1, and the quadratic threshold functions $f_S(x, y)$ and $f_T(x, y)$ are thus close in Hamming distance.

Note that in the certain case, i.e., when both Alice and Bob agree on S , they can easily compute $f_S(x, y)$ by having Alice send to Bob the ℓ bits $(x_i)_{i \in S}$. Moreover, if Alice and Bob are given access to public randomness in the uncertain case, Bob can figure out S via the hashing protocol mentioned

above using $\tilde{O}(\ell)$ bits of communication, which would reduce the problem to the certain case¹¹. The bulk of the proof will be to lower-bound the private-coin uncertain communication. Note that by the choice of our function class and distribution, this is equivalent to proving Lemma 1.4.

Proof of Lemma 1.4. To prove Lemma 1.4, the high-level intuition is that a protocol solving the uncertain problem should be essentially revealing to Bob the subset S that Alice holds. Formalizing this intuition turns out to be challenging, especially that a private-coin protocol solving the uncertain problem is only required to output a *single bit* which is supposed to equal the Boolean function $f_T(x, y)$ with high probability over (x, y) and over the private randomness. In fact, this high-level intuition can be shown not to hold in certain regimes¹². Moreover, the standard proofs that lower bound the communication of small-set intersection do not extend to lower-bound the communication complexity of f_T .

To lower-bound the private-coin communication of solving the uncertain task by a growing function of n , we consider the following *shift communication game*. Bob is given a sorted tuple $\sigma = (\sigma_1, \dots, \sigma_t)$ of integers with $1 \leq \sigma_1 < \dots < \sigma_t \leq n$, and Alice is either given the prefix $(\sigma_1, \dots, \sigma_{t-1})$ of length $t - 1$ of σ or the suffix $(\sigma_2, \dots, \sigma_t)$ of length $t - 1$ of σ . Bob needs to determine the input of Alice. We show that a celebrated lower bound of Linial [Lin92] on the *chromatic number* of certain related graphs implies a lower bound of $\log^{(t+1)}(n)$ on the private-coin communication of the shift communication game. We then show that any private-coin protocol solving the uncertain task can be turned into a private-coin protocol solving the shift-communication game with a constant (i.e., independent of n) blow-up in the communication (see Protocol 1).

2.2 Overview of Proof of Theorem 1.3

Reduction to Lemma 1.5. The proof of Theorem 1.3 builds on the lower-bound construction of [GKKS16] which we recall next. Let μ be the distribution over pairs $(x, y) \in \{0, 1\}^{2n}$ where x is uniform-random and y is an ϵ -noisy copy of x with $\epsilon = \sqrt{\delta/n}$. Then, the mutual information between x and y satisfies $I \approx n$. For each $S \subseteq [n]$, consider the function $f_S(x, y) \triangleq \langle S, x \oplus y \rangle$ where the inner product is over \mathbb{F}_2 , $x \oplus y$ denotes the coordinate-wise XOR of x and y , and S is used to denote both the subset and its 0/1 indicator vector. Moreover, consider the class \mathcal{F} of all pairs of functions (f_S, f_T) where $|S \Delta T| \leq \sqrt{\delta n}$. It can be seen that for such S and T , the distance between f_S and f_T under μ is at most δ . If Alice and Bob both know S , then Alice can send the single bit $\langle S, x \rangle$ to Bob who can then output the correct answer $\langle S, x \oplus y \rangle = \langle S, x \rangle \oplus \langle S, y \rangle$. This means that the certain communication is 1 bit. Using the well-known discrepancy method, [GKKS16] showed a lower bound of $\Omega(\sqrt{n})$ bits on the communication of the associated uncertain problem. Since in this case $I \approx n$, this in fact lower-bounds the uncertain communication by $\Omega(\sqrt{I})$ bits. For this construction, this lower bound turns out to be tight up to a logarithmic factor.

To improve the lower-bound from \sqrt{I} to $\sqrt{k} \cdot \sqrt{I}$, we consider the following “block-composed” framework. Let $\{f_{S^{(i)}}(x^{(i)}, y^{(i)}) : i \in [k]\}$ be k independent copies of the above base problem of [GKKS16] and consider computing the composed function $g(f_{S^{(1)}}(x^{(1)}, y^{(1)}), \dots, f_{S^{(k)}}(x^{(k)}, y^{(k)}))$ for some outer function $g : \{0, 1\}^k \rightarrow \{0, 1\}$. For any choice of g , the certain communication of the composed function would be at most k bits. When choosing the outer function g to use in our

¹¹Alternatively, Alice and Bob can run the protocol of [GKKS16] which would communicate $O(\ell)$ bits.

¹²For example, for constant error probabilities, the 1-way randomized communication complexity of small-set intersection is known to be $\Theta(\ell \cdot \log(\ell))$ bits (see, e.g., [BCK⁺14]) whereas the public-coin protocol of [GKKS16] can compute f_T with $O(\ell)$ bits of communication.

lower bound, we thus have two objectives to satisfy. First, g has to be sufficiently *hard* in the sense that its average-case decision tree complexity with respect to the uniform distribution on $\{0, 1\}^k$ should be $\Omega(k)$; otherwise, it will not be the case that the uncertain communication of computing g on k copies of the base problem is at least k times the uncertain communication of the base problem. Second, g has to be *noise stable* in order to be able to upper bound the distance between $g(f_{S^{(1)}}(\cdot), \dots, f_{S^{(k)}}(\cdot))$ and $g(f_{T^{(1)}}(\cdot), \dots, f_{T^{(k)}}(\cdot))$.

Note that setting g to be a dictator function would satisfy the noise-stability property, but it clearly would not satisfy the hardness property, as the composed function would be equal to the base function and would thus have uncertain communication $\tilde{O}(\sqrt{n})$ bits. Another potential choice of g is to set it to the parity function on k bits. This function would satisfy the hardness property, but it would strongly violate the noise stability property that is crucial to us. This leads us to setting g to the majority function on k bits, which is well-known to be noise stable, and has average-case decision-tree complexity $\Omega(k)$ with respect to the uniform distribution on $\{0, 1\}^k$. In fact, the noise stability of the majority function readily implies an upper bound of $O(\sqrt{\delta})$ on the distance between any pair of composed functions that are specified by tuples of subsets $(S^{(1)}, \dots, S^{(k)})$ and $(T^{(1)}, \dots, T^{(k)})$ with $|S^{(i)} \Delta T^{(i)}| \leq \sqrt{\delta n}$ for each $i \in [k]$. The crux of the proof will be to lower-bound the uncertain communication of the majority-composed function by $\Omega(k\sqrt{n})$, which amounts to proving Lemma 1.5. Since in this block-composed framework the mutual information satisfies $I \approx kn$, this would imply the lower bound of $\Omega(\sqrt{k}\sqrt{I})$ on the uncertain communication in Part (iii) of Theorem 1.3.

Proof of Lemma 1.5. We first point out that the average-case *quantum* decision tree complexity of Maj_k with respect to the uniform distribution is $\tilde{O}(\sqrt{k})$ [ADW01]. This implies that any communication complexity lower-bound method that extends to the quantum model cannot prove a lower bound larger than $\tilde{O}(\sqrt{k} \cdot \sqrt{n})$ on our uncertain communication¹³. In particular, we cannot solely rely on the discrepancy bound (as done in [GKKS16]), since this bound is known to lower-bound quantum communication. Similarly, the techniques of [She08, SZ07, LZ10] rely on the generalized discrepancy bound (originally due to [Kla01]) which also lower-bounds quantum communication. Moreover, the recent results of [MWY15] only apply to product distributions (i.e., where Alice’s input is independent of Bob’s input) in contrast to our case where the inputs of Alice and Bob are very highly-correlated. Finally, the recent works of [GLM⁺15, GPW15] do not imply lower bounds on the average-case complexity with respect to the distribution that arises in our setup.

To circumvent the above obstacles, we use a new approach that is tailored to our setup and that is outlined next. Let Π be a 1-way protocol solving the uncertain task with high probability. We consider the information that Π reveals about the *inputs to the outer function*, i.e., about the length- k binary string $(f_{S^{(1)}}(x^{(1)}, y^{(1)}), \dots, f_{S^{(k)}}(x^{(k)}, y^{(k)}))$. We call this quantity the *intermediate information cost* of Π , and we argue that it is at least $\Omega(k)$ bits. To do so, we recall the Hamming distance function HD_k defined by $\text{HD}_k(u, v) = 1$ if the Hamming distance between u and v is at least $k/2$ and $\text{HD}_k(u, v) = 0$ otherwise. We upper bound the information complexity of computing HD_k over the uniform distribution on $\{0, 1\}^{2k}$ by the intermediate information cost of Π . We do so by giving an information-cost preserving procedure (Protocol 2) where Alice and Bob are given independent uniformly distributed u and v (respectively) and use their private and public coins in order to simulate the input distribution (X, Y) of our uncertain problem. The known 1-way lower

¹³Thus, since $I \approx k \cdot n$ in our block-composed framework, such methods cannot be used to improve the lower-bound of $\Omega(\sqrt{I})$ of [GKKS16] by more than logarithmic factors.

bound of [Woo07] on HD_k under the uniform distribution then implies that Π reveals $\Omega(k)$ bits of information to Bob about the tuple $(f_{S^{(1)}}(x^{(1)}, y^{(1)}), \dots, f_{S^{(k)}}(x^{(k)}, y^{(k)}))$. This allows Bob to guess this tuple with probability 0.51^k . We then apply the strong direct product theorem for discrepancy of [LSS08] which, along with the discrepancy-based lower bound on the communication of the base uncertain problem of [GKKS16], implies that Π should be communicating at least $\Omega(k\sqrt{n})$ bits.

Organization of the rest of the paper In Section 3, we give some preliminaries that will be useful to us. In Section 4, we prove Theorem 1.1. In Section 5, we prove Theorem 1.3. The proof of Theorem 1.2 is given in Section 6. In Section 7, we conclude with some interesting open questions. A useful lemma that is used in Section 4 appears in Appendix A.

3 Preliminaries

For a real number x , we define $\text{Sign}(x)$ to be 1 if $x \geq 0$ and 0 if $x < 0$. For a set S , we write $X \in_R S$ to indicate that X is a random variable that is uniformly distributed on S . For a positive integer n , we let $[n] \triangleq \{1, \dots, n\}$. For a real number x , we denote $\exp(x)$ a quantity of the form $2^{\Theta(x)}$. For any two subsets $S, T \subseteq [n]$, we let $S \setminus T$ be the set of all elements of S that are not in T . We let $S \Delta T$ be the symmetric difference of S and T , i.e., the union of $S \setminus T$ and $T \setminus S$. For functions $f, g : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ and any distribution μ on $\mathcal{X} \times \mathcal{Y}$, we define the distance $\Delta_\mu(f, g) \triangleq \Pr_{(x, y) \sim \mu}[f(x, y) \neq g(x, y)]$ as the Hamming distance between the values of f and g , weighted with respect to μ . If μ is the uniform distribution on $\mathcal{X} \times \mathcal{Y}$, we drop the subscript μ and denote Δ_μ by Δ . We next recall the standard communication complexity model of Yao [Yao79]. In Definitions 3.1, 3.2 and 3.3 below, we let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ be a Boolean function.

Definition 3.1 (Deterministic Communication Complexity). *The two-way (resp. one-way) deterministic communication complexity of f , denoted by $\text{CC}(f)$ (resp. $\text{owCC}(f)$), is defined as the minimum over all two-way (resp. one-way) deterministic protocols Π that compute f correctly on every input pair, of the communication cost of Π .*

Definition 3.2 (Private-Coin Communication Complexity). *The two-way (resp. one-way) private-coin communication complexity of f with error ϵ , denoted by $\text{PrivCC}_\epsilon(f)$ (resp. $\text{owPrivCC}_\epsilon(f)$), is defined as the minimum over all two-way (resp. one-way) private-coin protocols Π that compute f with probability at least $1 - \epsilon$ on every input pair, of the communication cost of Π .*

The quantities in Definitions 3.1 and 3.2 can be similarly defined for partial functions f .

Definition 3.3 (Distributional Communication Complexity). *Let μ be a distribution over $\mathcal{X} \times \mathcal{Y}$. The two-way (resp. one-way) distributional communication complexity of f over μ with error ϵ , denoted by $\text{CC}_\epsilon^\mu(f)$ (resp. $\text{owCC}_\epsilon^\mu(f)$), is the minimum over all two-way (resp. one-way) protocols Π that compute f with probability $1 - \epsilon$ over μ , of the communication cost of Π .*

We next recall the model of communication with contextual uncertainty. For more details on this model, we refer the reader to [GKKS16]. In this setup, Alice knows a function f and is given an input x , and Bob knows a function g and is given an input y . Let $\mathcal{F} \subseteq \{f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}\}^2$ be a family of pairs of Boolean functions with domain $\mathcal{X} \times \mathcal{Y}$, and μ be a distribution on $\mathcal{X} \times \mathcal{Y}$. We say that a public-coin (resp. private-coin) protocol Π ϵ -computes \mathcal{F} over μ if for every $(f, g) \in \mathcal{F}$, we have that Π outputs the value $g(x, y)$ with probability at least $1 - \epsilon$ over the randomness of $(x, y) \sim \mu$ and over the public (resp. private) randomness of Π .

Definition 3.4 (Contextually Uncertain Communication Complexity). *Let μ be a distribution on $\mathcal{X} \times \mathcal{Y}$ and $\mathcal{F} \subseteq \{f: \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}\}^2$. The two-way (resp. one-way) public-coin communication complexity of \mathcal{F} under contextual uncertainty, denoted $\text{PubCCU}_\epsilon^\mu(\mathcal{F})$ (resp. $\text{owPubCCU}_\epsilon^\mu(\mathcal{F})$), is the minimum over all two-way (resp. one-way) public-coin protocols Π that ϵ -compute \mathcal{F} over μ , of the maximum communication complexity of Π over $(f, g) \in \mathcal{F}$, (x, y) from the support of μ and settings of the public coins.*

Similarly, the two-way (resp. one-way) private-coin communication complexity of \mathcal{F} under contextual uncertainty $\text{PrivCCU}_\epsilon^\mu(\mathcal{F})$ (resp. $\text{owPrivCCU}_\epsilon^\mu(\mathcal{F})$) is defined by restricting to two-way (resp. one-way) private-coin protocols.

4 Construction for Private-Coin Uncertain Protocols

We now describe the construction that is used to prove Theorem 1.1. Each function in our universe is specified by a subset $S \subseteq [n]$ and is of the form $f_S: \{\pm 1\}^n \times \{\pm 1\}^n \rightarrow \{0, 1\}$ with $f_S(X, Y) \triangleq \text{Sign}(\sum_{i \in S} X_i Y_i)$ for all $X, Y \in \{\pm 1\}^n$. The function class is then defined by

$$\mathcal{F}_\delta \triangleq \{(f_S, f_T) : S \subseteq T, |T| = \ell \text{ and } |T \setminus S| \leq \delta' \cdot \ell\},$$

where $\delta' = \alpha \cdot \delta^2$ for some sufficiently small positive absolute constant α , and $\ell = \ell(\delta)$ is a sufficiently large function of δ . The input pair (X, Y) is drawn from the uniform distribution on $\{\pm 1\}^{2n}$. We start with the proof of Part (i) of Theorem 1.1. It essentially follows from the fact that the polynomials $\sum_{i \in S} X_i Y_i$ and $\sum_{i \in T} X_i Y_i$ behave like zero-mean Gaussians with unit-variance and correlation $\sqrt{1 - \delta'}$.

We will need the following well-known fact which follows from Sheppard's formula [She99].

Fact 4.1. *If (X, Y) is a pair of zero-mean Gaussians with correlation $\mathbb{E}[XY] = \rho$, then*

$$\Pr[\text{Sign}(X) \neq \text{Sign}(Y)] = \frac{\arccos(\rho)}{\pi}.$$

We now prove Part (i) of Theorem 1.1.

Proof of Part (i) of Theorem 1.1. Let $S \subseteq T \subseteq [n]$ be such that $|T| = \ell$ and $|T \setminus S| \leq \delta' \cdot \ell$. For fixed ℓ , the distance $\Delta(f_S, f_T)$ decreases when $|T \setminus S|$ decreases. So it suffices to upper bound $\Delta(f_S, f_T)$ when $|T \setminus S| = \delta' \cdot \ell$. Assume that the coordinates in T are $1 \leq t_1 < t_2 < \dots < t_\ell$. Then, we define the random vectors $X', Y' \in \{0, 1\}^\ell$ as $X'_i = X_{t_i} Y_{t_i}$ for all $i \in [\ell]$, and $Y'_i = X_{t_i} Y_{t_i}$ if $t_i \in S$ and $Y'_i = 0$ if $t_i \notin S$. We will apply the two-dimensional Berry-Esseen Theorem A.1 to (X', Y') . To do so, first note that the random pairs $(X'_1, Y'_1), (X'_2, Y'_2), \dots, (X'_\ell, Y'_\ell)$ are independent. Moreover, for every $i \in [\ell]$ such that $t_i \in S$, the covariance matrix of (X_i, Y_i) is given by $\Sigma_i = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$. On the other hand, for $i \in [\ell]$ such that $t_i \notin S$, the covariance matrix of (X_i, Y_i) is given by $\Sigma_i = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$. Hence, the average (across the ℓ coordinates) covariance matrix is given by $\Sigma = \ell^{-1} \cdot \sum_{i \in [\ell]} \Sigma_i = \begin{bmatrix} 1 & 1 - \delta' \\ 1 - \delta' & 1 - \delta' \end{bmatrix}$.

The smallest and largest eigenvalues of Σ are respectively given by

$$\begin{aligned}\lambda &\triangleq \frac{1 - \delta'}{2} - \frac{\sqrt{5 \cdot (1 - \delta')^2 - 2 \cdot (1 - \delta') + 1}}{2} + \frac{1}{2}, \\ \Lambda &\triangleq \frac{1 - \delta'}{2} + \frac{\sqrt{5 \cdot (1 - \delta')^2 - 2 \cdot (1 - \delta') + 1}}{2} + \frac{1}{2}.\end{aligned}\tag{1a}$$

In Equation (1a), it can be checked that for $\delta' \in (0, 1)$, $\lambda > 0$. By the two-dimensional Berry-Esseen Theorem A.1, we get that

$$\begin{aligned}\Delta(f_S, f_T) &= \Pr[\text{Sign}(\sum_{i \in S} X_i Y_i) \neq \text{Sign}(\sum_{i \in T} X_i Y_i)] \\ &= \Pr[\text{Sign}(\sum_{i \in [\ell]} X'_i) \neq \text{Sign}(\sum_{i \in [\ell]} Y'_i)] \\ &= \Pr[\text{Sign}(X'') \neq \text{Sign}(Y'')] \pm O\left(\frac{1}{\lambda^{3/2} \cdot \sqrt{\ell}}\right),\end{aligned}\tag{2}$$

where (X'', Y'') is a pair of zero-mean Gaussians with covariance matrix Σ . We can scale Y'' to make it have unit-variance; this does not change its mean or the probability in Equation (2). The covariance matrix becomes: $\Sigma' = \begin{bmatrix} 1 & \sqrt{1 - \delta'} \\ \sqrt{1 - \delta'} & 1 \end{bmatrix}$. By Sheppard's formula (i.e., Fact 4.1), we deduce that

$$\begin{aligned}\Delta(f_S, f_T) &= \frac{\arccos(\sqrt{1 - \delta'})}{\pi} \pm O\left(\frac{1}{\lambda^{3/2} \cdot \sqrt{\ell}}\right) \\ &= O(\sqrt{\delta'}) \\ &\leq \delta,\end{aligned}$$

where we have used the facts that $\arccos(1 - x) = O(\sqrt{x})$ for small positive values of x , that $\delta' = \alpha \cdot \delta^2$ for a sufficiently small positive absolute constant α , and that ℓ is a sufficiently large function of δ . \square

We now give the proof of Part (ii) of Theorem 1.1, which is quite immediate.

Proof of Part (ii) of Theorem 1.1. When both Alice and Bob know the subset $T \subseteq [n]$ which satisfies $|T| \leq \ell$, Alice can send the sequence $(X_j : j \in T)$ of at most ℓ bits to Bob who can then output $f_T(X, Y)$. \square

To prove Part (iii) of Theorem 1.1, the next definition – which is based on the graphs studied by Linial [Lin92] – will be crucial to us.

Definition 4.2 (Shift Communication Game $\mathcal{G}_{m,t}$). *Let m and t be positive integers with $t \leq m$. In the communication problem $\mathcal{G}_{m,t}$, Bob is given a sorted tuple $\sigma = (\sigma_1, \dots, \sigma_t)$ of distinct integers with $1 \leq \sigma_1 < \dots < \sigma_t \leq m$. In the YES case, Alice is given the prefix $(\sigma_1, \dots, \sigma_{t-1})$ of length $t - 1$ of σ . In the NO case, Alice is given the suffix $(\sigma_2, \dots, \sigma_t)$ of length $t - 1$ of σ . Alice and Bob need to determine which of the YES and NO cases occurs.*

Lemma 4.3 lower-bounds the private-coin communication complexity of $\mathcal{G}_{m,t}$. Its proof uses Linial's lower bound on the chromatic number of related graphs.

Lemma 4.3. *There is an absolute constant c such that for every sufficiently small $\epsilon > 0$, we have that $\text{PrivCC}_\epsilon(\mathcal{G}_{m,t}) \geq c \cdot \log^{(t+2)}(m)$.*

We prove Lemma 4.3 in Section 4.1. The proof of Part (iii) of Theorem 1.1 – which is the main part in the proof of Theorem 1.1 – is given in Section 4.2.

4.1 Proof of Lemma 4.3

The following family of graphs was first studied by Linial (in the setup of distributed graph algorithms) [Lin92].

Definition 4.4 (Shift Graph $G_{m,t}$). *Let m and t be positive integers with $t \leq m$. In the graph $G_{m,t} = (V_{m,t}, E_{m,t})$, the vertices are all sorted tuples $\sigma = (\sigma_1, \dots, \sigma_t)$ of distinct integers with $1 \leq \sigma_1 < \dots < \sigma_t \leq m$. Two such tuples σ and π are connected by an edge in $E_{m,t}$ if and only if either $(\sigma_1, \dots, \sigma_{t-1}) = (\pi_2, \dots, \pi_t)$ or $(\sigma_2, \dots, \sigma_t) = (\pi_1, \dots, \pi_{t-1})$.*

Recall that the chromatic number $\chi(G)$ of an undirected graph G is the minimum number of colors needed to color its vertices such that no two adjacent vertices share the same color. The following theorem is due to Linial.

Theorem 4.5 ([Lin92], Proof of Theorem 2.1). *Let m and t be positive integers such that $t \leq m$ and t is odd. Then, $\chi(G_{m,t}) \geq \log^{(t-1)}(m)$.*

The next lemma uses Theorem 4.5 to lower-bound the deterministic two-way communication complexity of the shift communication game $\mathcal{G}_{m,t}$.

Lemma 4.6. *Let m and t be positive integers such that $t \leq m$ and t is odd. Then, it is the case that $\text{CC}(\mathcal{G}_{m,t}) \geq \log^{(t+1)}(m)$.*

Proof of Lemma 4.6. Assume for the sake of contradiction that there exists a deterministic two-way protocol that computes $\mathcal{G}_{m,t}$ and that has communication cost smaller than $\log^{(t+1)}(m)$. Then, using the fact that the one-way communication complexity of any function is at most exponential in its two-way communication complexity, we get that there is a one-way protocol Π that computes $\mathcal{G}_{m,t}$ and that has communication cost smaller than $\log^{(t)}(m)$. Let M be the single message sent from Alice to Bob under Π . Then, the length of M satisfies $|M| < \log^{(t)}(m)$. Note that Alice's input is an element of the vertex-set $V_{m,t}$ of the shift graph $G_{m,t}$ (Definition 4.4). Since M is a deterministic function of Alice's input, it induces a coloring of $V_{m,t}$ into less than $2^{\log^{(t)}(m)} = \log^{(t-1)}(m)$ colors. The fact that no two adjacent vertices in $G_{m,t}$ share the same color follows from the correctness of Π in computing $\mathcal{G}_{m,t}$. This contradicts Theorem 4.5. \square

The following known fact gives a generic lower bound on the bounded-error private-coin communication complexity in terms of the deterministic communication complexity.

Fact 4.7 ([KN97], Theorem 3.14). *For every communication function f and every non-negative ϵ that is bounded below $1/2$, we have that $\text{PrivCC}_\epsilon(f) = \Omega(\log(\text{CC}(f)))$.*

Lemma 4.3 now follows by combining Lemma 4.6 and Fact 4.7.

4.2 Proof of Part (iii) of Theorem 1.1

In this section, we prove Part (iii) of Theorem 1.1. We assume for the sake of contradiction that there exists a 1-way private-coin protocol Π computing \mathcal{F}_δ w.r.t. the uniform distribution on $\{0, 1\}^{2n}$ with error at most $\epsilon/2 - 2\delta - \eta$ and with communication cost $|\Pi| = o(\eta^2 \cdot \log^{(t)}(n))$ for some positive integer $t = \Theta((\epsilon/\delta)^2)$ that will be exactly specified later on. We will use Π to give a 1-way private-coin protocol Π' solving the shift communication game $\mathcal{G}_{m,t}$ with high constant probability and with communication cost $|\Pi'| \leq O(\eta^{-2} \cdot |\Pi|) = o(\log^{(t)}(n))$, which would contradict Lemma 4.3.

4.2.1 Description of Protocol Π'

The operation of protocol Π' , which uses Π as a black-box, is described in Protocol 1. Note that the parameters in Protocol 1 are defined in terms of ϵ (which is a number in $[\delta, 0.5]$ that is given in the statement of Theorem 1.1) and δ' (which, as mentioned above, is set to $\alpha \cdot \delta^2$ for a sufficiently small positive absolute constant α). Also as above, ℓ is set to a sufficiently large function of δ . Note that $\epsilon' = O(\epsilon^2)$, and hence $t = O((\epsilon/\delta)^2)$. In Protocol 1, Bob is given as input a sorted tuple σ , and Alice is given as input either the prefix ϕ of σ or its suffix ψ . In steps 1 and 2, Alice and Bob “stretch” their tuples, which amounts to each of them repeating each bit of the corresponding 0/1 indicator vector a certain number of times and appending a certain number of zeros (see Definition 4.8 and Figure 1 in Section 4.2.2 below for a thorough definition of $\text{stretch}_{r,a}$). The aim of steps 1 and 2 is for Alice and Bob to produce a pair of subsets (S, T) of $[n]$ such that $(f_S, f_T) \in \mathcal{F}_\delta$. The goal of Protocol 1 is for Bob to determine if Alice was given the prefix or the suffix of his tuple σ . To do so, Alice and Bob sample (using their private coins) k random inputs $(X^{(i)}, Y^{(i)})_{i \in [k]}$ (steps 3 and 4) and then simulate the given private-coin protocol Π to compute the function $f_T(X, Y)$ on the k random input-pairs that were sampled (steps 6 and 7). Moreover, Alice sends to Bob the influential bits of $X^{(i)}$ for each $i \in [k]$ (step 8). The main idea will be to be for Bob to compute the empirical error corresponding to each of the prefix (step 11) and suffix (step 12), and then output the hypothesis with the smallest empirical error (step 13).

4.2.2 The $\text{stretch}_{r,a}$ Procedure

In this section, we thoroughly define and illustrate the stretching procedure used in Protocol 1 and mentioned in Section 4.

Definition 4.8 (The $\text{stretch}_{r,a}$ procedure). *Let d, t, r and a be positive integers with $t \leq d$. For any sorted tuple $\sigma = (\sigma_1, \dots, \sigma_t)$ of distinct integers with $1 \leq \sigma_1 < \dots < \sigma_t \leq d$, we first let $z \in \{0, 1\}^d$ be the 0/1 indicator vector of the subset of $[d]$ corresponding to σ . Let $w \in \{0, 1\}^{d \cdot r + a}$ be the string obtained from z by repeating each of its coordinates r times, and then appending a ones. Namely, for each $i \in [d]$ and $j \in [r]$, we set $w_{(i-1) \cdot r + j} = z_i$ and for each $j \in [a]$, we set $w_{d \cdot r + j} = 1$. Then, we let $1 \leq \Sigma_1 < \Sigma_2 < \dots < \Sigma_{t \cdot r + a} \leq d \cdot r + a$ be the indices of the coordinates of w that are equal to 1. The output of $\text{stretch}_{r,a}(\sigma)$ is then the pair (Σ, W) where $\Sigma \triangleq (\Sigma_1, \Sigma_2, \dots, \Sigma_{t \cdot r + a})$ and $W \subseteq [d \cdot r + a]$ is the support of w .*

The operation of $\text{stretch}_{r,a}$ is illustrated in Figure 1 in the particular case where $r = 2, a = 3$ and $d = 9$. The “appending parameter” a allows us to control how small is the *normalized* Hamming distance between the binary strings corresponding to Σ and Φ (respectively Ψ). The purpose of the

Protocol 1 Reduction Protocol Π'

Parameters. $\epsilon' = 1 - \cos(\epsilon\pi)$, $t = \lceil \epsilon'/\delta' \rceil$, $r = \delta' \cdot \ell$, $a = \ell \cdot (1 - t \cdot \delta')$, $s = \ell \cdot (1 - \delta')$, $k = \Theta(1/\eta^2)$.
Inputs. Bob is given a sorted tuple $\sigma = (\sigma_1, \dots, \sigma_t)$ of integers with $1 \leq \sigma_1 < \dots < \sigma_t \leq (n-a)/r$. Alice is given a sorted tuple $\lambda \in \{\phi, \psi\}$ where $\phi \triangleq (\sigma_1, \dots, \sigma_{t-1})$ and $\psi \triangleq (\sigma_2, \dots, \sigma_t)$.

- 1: Alice sets $(\Lambda, S) \leftarrow \text{stretch}_{r,a}(\lambda)$.
 - 2: Bob sets $(\Sigma, T) \leftarrow \text{stretch}_{r,a}(\sigma)$.
 - 3: Alice uses her private randomness to sample k i.i.d. strings $X^{(1)}, \dots, X^{(k)} \in_R \{0, 1\}^n$.
 - 4: Bob uses his private randomness to sample k i.i.d. strings $Y^{(1)}, \dots, Y^{(k)} \in_R \{0, 1\}^n$.
 - 5: **for** $i = 1, \dots, k$ **do**
 - 6: Alice and Bob simulate the protocol Π on inputs $((S, X^{(i)}), (T, Y^{(i)}))$.
 - 7: Bob computes the resulting output bit B_i .
 - 8: Alice sends to Bob the sequence of bits $(X_{\Lambda_1}^{(i)}, X_{\Lambda_2}^{(i)}, \dots, X_{\Lambda_s}^{(i)})$.
 - 9: **end for**
 - 10: Bob sets $\Phi \leftarrow \text{stretch}_{r,a}(\phi)$ and $\Psi \leftarrow \text{stretch}_{r,a}(\psi)$.
 - 11: Bob computes the “prefix error” $\text{err}^{(p)} \triangleq \sum_{i \in [k]} \mathbb{1} \left[\text{Sign} \left(\sum_{j \in [s]} X_{\Lambda_j}^{(i)} Y_{\Phi_j}^{(i)} \right) \neq B_i \right]$.
 - 12: Bob computes the “suffix error” $\text{err}^{(s)} \triangleq \sum_{i \in [k]} \mathbb{1} \left[\text{Sign} \left(\sum_{j \in [s]} X_{\Lambda_j}^{(i)} Y_{\Psi_j}^{(i)} \right) \neq B_i \right]$.
 - 13: Bob returns YES if $\text{err}^{(p)} \leq \text{err}^{(s)}$ and NO otherwise.
-

“repetition parameter” r is the following. Consider the tuples $\phi = (2, 4, 5, 7)$ and $\psi = (4, 5, 7, 9)$ in Figure 1. The number of differing coordinates between the tuples ϕ and ψ is 4. After stretching, the number of differing coordinates between the resulting tuples Φ and Ψ is amplified to $4 \cdot r = 8$. These two notions of distance (the number of coordinates on which the tuples differ and the normalized Hamming distance between the corresponding binary strings) are important to us. The fact that these two distances are important to us is the reason why we let the $\text{stretch}_{r,a}$ procedure have two equivalent outputs (a subset W and a tuple Σ).

$$\begin{array}{l}
 \text{Bob's subset: } 0 \mathbf{1} 0 1 1 0 1 0 \mathbf{1} \xrightarrow{\text{stretch}_{r,a}} 0 0 \mathbf{1 1} 0 0 1 1 1 1 1 0 0 1 1 0 0 \mathbf{1 1 1 1 1} \\
 \sigma = (2, 4, 5, 7, 9) \qquad \qquad \qquad \Sigma = (3, 4, 7, 8, 9, 10, 13, 14, 17, 18, 19, 20, 21) \\
 \\
 \text{Prefix of } \sigma: 0 \mathbf{1} 0 1 1 0 1 0 \mathbf{0} \xrightarrow{\text{stretch}_{r,a}} 0 0 \mathbf{1 1} 0 0 1 1 1 1 1 0 0 1 1 0 0 \mathbf{0 0 1 1 1} \\
 \phi = (2, 4, 5, 7) \qquad \qquad \qquad \Phi = (3, 4, 7, 8, 9, 10, 13, 14, 19, 20, 21) \\
 \\
 \text{Suffix of } \sigma: 0 \mathbf{0} 0 1 1 0 1 0 \mathbf{1} \xrightarrow{\text{stretch}_{r,a}} 0 0 \mathbf{0 0} 0 0 1 1 1 1 1 0 0 1 1 0 0 \mathbf{1 1 1 1 1} \\
 \psi = (4, 5, 7, 9) \qquad \qquad \qquad \Psi = (7, 8, 9, 10, 13, 14, 17, 18, 19, 20, 21)
 \end{array}$$

Figure 1: Operation of the $\text{stretch}_{r,a}$ procedure for $r = 2$ and $a = 3$.

4.2.3 Analysis of Protocol Π'

We now turn to the formal analysis of Protocol 1. First, note that the communication cost of protocol Π' satisfies $|\Pi'| \leq O(\eta^{-2} \cdot |\Pi|)$. Let Λ, Σ, Φ and Ψ be the ordered sequences defined in the

operation of Protocol 1. Define the functions g and h as

$$g(X, Y) \triangleq \text{Sign}\left(\sum_{j \in [s]} X_{\Lambda_j} Y_{\Phi_j}\right), \quad (3a)$$

$$h(X, Y) \triangleq \text{Sign}\left(\sum_{j \in [s]} X_{\Lambda_j} Y_{\Psi_j}\right). \quad (3b)$$

Note that steps 11 and 12 of the protocol compute the empirical errors of functions g and h respectively. Since $\Lambda \in \{\Phi, \Psi\}$, let's assume WLOG that $\Lambda = \Phi$ and show that the protocol Π' returns YES with high probability. The case where $\Lambda = \Psi$ is symmetric. When $\Lambda = \Phi$, we have that $g = f_S$ where S is the subset of $[n]$ that Alice gets in step 1. The operation of the $\text{stretch}_{r,a}$ procedure (Definition 4.8) guarantees that $S \subseteq [n]$, $|T| = \ell$ and $|T \setminus S| = \delta' \cdot \ell$. Part (i) of Theorem 1.1 then implies that $\Delta(g, f_T) \leq \delta$. In order to show that in step 13 Bob returns YES with high probability, the main idea will be to lower bound the distance between the functions g and h . This is done in the next lemma.

Lemma 4.9. *The functions g and h defined in Equations (3a) and (3b) satisfy $\Delta(g, h) \geq \epsilon - \delta$.*

To prove Lemma 4.9, we use the next lemma which spells out the distribution of the sequence $(X_{\Lambda_j} Y_{\Phi_j}, X_{\Lambda_j} Y_{\Psi_j})_{j \in [s]}$ of random variables.

Lemma 4.10. *The random variables $(X_{\Lambda_1} Y_{\Phi_1}, X_{\Lambda_1} Y_{\Psi_1}), (X_{\Lambda_2} Y_{\Phi_2}, X_{\Lambda_2} Y_{\Psi_2}), \dots, (X_{\Lambda_s} Y_{\Phi_s}, X_{\Lambda_s} Y_{\Psi_s})$ are independent, and they are distributed as follows:*

1. For $1 \leq j \leq s - a$, $(X_{\Lambda_j} Y_{\Phi_j}, X_{\Lambda_j} Y_{\Psi_j})$ is uniformly distributed on $\{\pm 1\}^2$.
2. For $s - a < j \leq s$, $(X_{\Lambda_j} Y_{\Phi_j}, X_{\Lambda_j} Y_{\Psi_j})$ is uniformly distributed on $\{(-1, -1), (+1, +1)\}$.

Proof of Lemma 4.10. By the operation of the $\text{stretch}_{r,a}$ procedure (Definition 4.8) in steps 1 and 2 of Protocol 1, for every $s - a < j \leq s$, it holds that $(X_{\Lambda_j} Y_{\Phi_j}, X_{\Lambda_j} Y_{\Psi_j}) = (X_j Y_j, X_j Y_j)$ and that X_j and Y_j do not contribute to any other $(X_{\Lambda_{j'}} Y_{\Phi_{j'}}, X_{\Lambda_{j'}} Y_{\Psi_{j'}})$ pair. This implies part 1 and that for each $s - a < j \leq s$, the pair $(X_{\Lambda_j} Y_{\Phi_j}, X_{\Lambda_j} Y_{\Psi_j})$ is independent of all other pairs in the sequence.

We next prove by induction on $1 \leq j \leq s - a$ that conditioned on $(X_{\Lambda_{<j}} Y_{\Phi_{<j}}, X_{\Lambda_{<j}} Y_{\Psi_{<j}})$ taking any particular value, $(X_{\Lambda_j} Y_{\Phi_j}, X_{\Lambda_j} Y_{\Psi_j})$ is uniformly distributed on $\{\pm 1\}^2$. To see this (assuming WLOG that $\Lambda = \Phi$), note that $(X_{\Lambda_j} Y_{\Phi_j}, X_{\Lambda_j} Y_{\Psi_j}) = (X_{\Lambda_j} Y_{\Lambda_j}, X_{\Lambda_j} Y_{\Psi_j})$ where X_{Λ_j} and Y_{Ψ_j} do not appear in $(X_{\Lambda_{<j}} Y_{\Phi_{<j}}, X_{\Lambda_{<j}} Y_{\Psi_{<j}})$. Hence, conditioned on $(X_{\Lambda_{<j}} Y_{\Phi_{<j}}, X_{\Lambda_{<j}} Y_{\Psi_{<j}})$ taking any particular value, $X_{\Lambda_j} Y_{\Phi_j}$ is a uniformly random bit (because of X_{Λ_j}). Moreover, conditioned on $(X_{\Lambda_{<j}} Y_{\Phi_{<j}}, X_{\Lambda_{<j}} Y_{\Psi_{<j}})$ and $X_{\Lambda_j} Y_{\Phi_j}$ taking any particular values, $X_{\Lambda_j} Y_{\Psi_j}$ is a uniformly random bit (because of Y_{Ψ_j}). This completes the proof of the lemma. \square

We now use Lemma 4.10 along with a two-dimensional Central Limit Theorem in order to prove Lemma 4.9.

Proof of Lemma 4.9. By Lemma 4.10, independently for each $1 \leq j \leq s - a$, $(X_{\Lambda_j} Y_{\Phi_j}, X_{\Lambda_j} Y_{\Psi_j})$ is a ρ_j -correlated random pair with $\rho_j = 0$, and independently for each $s - a < j \leq s$, $(X_{\Lambda_j} Y_{\Phi_j}, X_{\Lambda_j} Y_{\Psi_j})$ is a ρ_j -correlated random pair with $\rho_j = 1$. Hence, the average correlation across coordinates is

$$\rho \triangleq k^{-1} \cdot \sum_{j \in [k]} \rho_j = \frac{a}{s} = \frac{\ell \cdot (1 - t \cdot \delta')}{\ell \cdot (1 - \delta')} = 1 - \frac{(t - 1) \cdot \delta'}{1 - \delta'} \leq 1 - (\epsilon' - \delta'),$$

where the last inequality used the setting of $t = \lceil \epsilon' / \delta' \rceil$ in Protocol 1. Denoting by Σ_j the covariance matrix of $(X_{\Lambda_j} Y_{\Phi_j}, X_{\Lambda_j} Y_{\Psi_j})$ for every $j \in [s]$, the average (across coordinates) covariance matrix is then given by $\Sigma \triangleq k^{-1} \cdot \sum_{j \in [k]} \Sigma_j = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$. Note that the smallest eigenvalue of Σ is $\lambda = 1 - |\rho|$.

Let (X', Y') be a pair of zero-mean ρ -correlated Gaussians. By the two-dimensional Berry-Esseen Theorem A.1 and Sheppard's formula (i.e., Fact 4.1), we get that

$$\begin{aligned}
\Delta(g, h) &= \Pr[g(X, Y) \neq h(X, Y)] \\
&= \Pr[g(X, Y) = 0, h(X, Y) = 1] + \Pr[g(X, Y) = 1, h(X, Y) = 0] \\
&= \Pr[X' < 0, Y \geq 0] \pm O\left(\frac{1}{(1 - |\rho|)^{3/2} \cdot \sqrt{s}}\right) + \Pr[X' \geq 0, Y < 0] \pm O\left(\frac{1}{(1 - |\rho|)^{3/2} \cdot \sqrt{s}}\right) \\
&= \Pr[\text{Sign}(X') \neq \text{Sign}(Y')] \pm O\left(\frac{1}{(1 - |\rho|)^{3/2} \cdot \sqrt{\ell}}\right) \\
&= \frac{\arccos(\rho)}{\pi} \pm O\left(\frac{1}{(1 - |\rho|)^{3/2} \cdot \sqrt{\ell}}\right) \\
&\geq \frac{\arccos(1 - (\epsilon' - \delta'))}{\pi} - O\left(\frac{1}{(1 - |\rho|)^{3/2} \cdot \sqrt{\ell}}\right) \\
&\geq \frac{\arccos(1 - \epsilon')}{\pi} - O(\sqrt{\delta'}) - O\left(\frac{1}{(1 - |\rho|)^{3/2} \cdot \sqrt{\ell}}\right) \\
&\geq \epsilon - \delta,
\end{aligned}$$

where the last inequality follows by setting $\epsilon' = 1 - \cos(\epsilon\pi)$ for the given $\epsilon \in [\delta, 0.5]$, setting ℓ to be a sufficiently large function of δ , and setting $\delta' = \alpha \cdot \delta^2$ for a sufficiently small positive absolute constant α . \square

We are now ready to complete the proof of Part (iii) of Theorem 1.1. In Protocol 1, the tuples $((X^{(i)}, Y^{(i)}), B_i)_{i \in [k]}$ are i.i.d. samples corresponding to a function q which, by the error guarantee of protocol Π , is $(\epsilon/2 - 2\delta - \eta)$ -close to f_T . Since $\Delta(g, f_T) \leq \delta$, we get that $\Delta(g, q) \leq (\epsilon/2 - \delta - \eta)$. By Lemma 4.9, we also get that $\Delta(h, q) \geq (\epsilon/2 + \eta)$. By Hoeffding's bound (i.e., Fact 4.11), for $k = \Theta(1/\eta^2)$, the empirical error $\epsilon^{(p)}$ of g on the samples $((X^{(i)}, Y^{(i)}), B_i)_{i \in [k]}$ is less than the empirical error $\epsilon^{(s)}$ of h on these samples, with high constant probability. Hence, Bob returns YES with high constant probability. A symmetric argument shows that when $\Lambda = \Psi$, Bob returns NO with high constant probability, which completes the proof.

Fact 4.11 (Hoeffding's bound). *Consider a coin that shows up head with probability p . Let $H(k)$ be the number of heads obtained in k independent tosses of this coin. Then, for every $\epsilon > 0$, $\Pr[|H(k) - pk| > \epsilon k] \leq 2e^{-2\epsilon^2 k}$.*

Note. *As mentioned in Note 1.1, the above construction cannot give a separation larger than $\Theta(\log \log n)$. This is because using private randomness, Bob can learn the set S using $O(\log \log n)$ bits of communication (see, e.g., [BCK⁺14]). Additionally, Alice can send the coordinates of X indexed by the elements of S to Bob who can then compute $f_S(X, Y)$.*

5 Construction for Public-Coin Uncertain Protocols

In this section, we describe the construction that is used to prove Theorem 1.3. We set

$$\delta' \triangleq c \cdot \delta^2 \quad \text{and} \quad \epsilon \triangleq \sqrt{\delta'/n}, \quad (\dagger)$$

where δ is the parameter from the statement of Theorem 1.3, and $c > 0$ is a small-enough absolute constant. To define our input distribution, we first define a slightly more general distribution μ_η . The support of μ_η is $\{0, 1\}^{kn} \times \{0, 1\}^{kn}$ and we will view the coordinates of a sample $(x, y) \sim \mu_\eta$ as $x = (x^{(i)})_{i \in [k]}$ and $y = (y^{(i)})_{i \in [k]}$ with $x^{(i)}, y^{(i)} \in \{0, 1\}^n$ for all $i \in [k]$. A sample $(x, y) \sim \mu_\eta$ is generated by letting $x \in_R \{0, 1\}^{kn}$ and for all $i \in [k]$ and $j \in [n]$, independently setting $y_j^{(i)}$ to be an η -noisy copy of $x_j^{(i)}$. In other words, we set $y_j^{(i)} = x_j^{(i)}$ w.p. $1 - \eta$ and $y_j^{(i)} = 1 - x_j^{(i)}$ w.p. η . Our input distribution is then $\mu_{2\epsilon - 2\epsilon^2}$.

We now define our function class \mathcal{F}_ϵ . Each function in our universe is specified by a sequence of subsets $S \triangleq (S^{(i)} \subseteq [n])_{i \in [k]}$ and it is of the form $f_S : \{0, 1\}^{2 \cdot kn} \rightarrow \{0, 1\}$ with¹⁴

$$f_S(x, y) \triangleq \text{Sign}\left(\sum_{i \in [k]} (-1)^{\langle S^{(i)}, x^{(i)} \oplus y^{(i)} \rangle}\right) \quad (4)$$

for all $x, y \in \{0, 1\}^{k \cdot n}$, where in Eq. (4) the inner product is over \mathbb{F}_2 , the sum is over \mathbb{R} and $x^{(i)} \oplus y^{(i)}$ denotes the coordinate-wise XOR of the two length- n binary strings $x^{(i)}$ and $y^{(i)}$. The function class is then defined by $\mathcal{F}_\epsilon \triangleq \{(f_S, f_T) : |S^{(i)} \Delta T^{(i)}| \leq \epsilon \cdot n \text{ for all } i \in [k]\}$.

We now give the proof of Part (i) of Theorem 1.3. It follows from known bounds on the noise stability of the majority function.

Proof of Part (i) of Theorem 1.3. Let $(f_S, f_T) \in \mathcal{F}_\epsilon$, and denote $a_i \triangleq (-1)^{\langle S^{(i)}, x^{(i)} \oplus y^{(i)} \rangle}$ and $b_i \triangleq (-1)^{\langle T^{(i)}, x^{(i)} \oplus y^{(i)} \rangle}$ for every $i \in [k]$. Also, let $a \triangleq (a_i)_{i \in [k]}$ and $b \triangleq (b_i)_{i \in [k]}$. Note that (a_i, b_i) is a pair of ρ_i -correlated random strings with $\rho_i \geq (1 - 2\delta')$. Since $\Delta_{\mu_{2\epsilon - 2\epsilon^2}}(f_S, f_T)$ increases when ρ_i decreases, we assume WLOG that $\rho_i = 1 - 2\delta' \triangleq \rho$ for all $i \in [k]$.

Recall that the *noise stability* of a function $h : \{0, 1\}^k \rightarrow \{\pm 1\}$ is defined as $\text{Stab}_\rho(h) = \mathbb{E}[h(x)h(y)]$ where (x, y) is a random pair of ρ -correlated strings. Let the function $\widetilde{\text{Maj}}_k$ be defined by $\widetilde{\text{Maj}}_k(x) = (-1)^{\text{Maj}_k(x)}$ for all $x \in \{0, 1\}^k$. Recall (see [O'D14]) that the noise stability of $\widetilde{\text{Maj}}_k$ satisfies

$$\text{Stab}_\rho(\widetilde{\text{Maj}}_k) \geq 1 - \frac{2}{\pi} \arccos(\rho).$$

Hence, we get that

$$\begin{aligned} \Delta_{\mu_{2\epsilon - 2\epsilon^2}}(f_S, f_T) &= \Pr[f_S(x, y) \neq f_T(x, y)] \\ &\leq \frac{1 - \text{Stab}_\rho(\widetilde{\text{Maj}}_k)}{2} \\ &\leq \frac{\arccos(\rho)}{\pi} \\ &= O(\sqrt{\delta'}) \\ &\leq \delta, \end{aligned}$$

¹⁴We will use the symbol $S^{(i)}$ to denote both the subset of $[n]$ and its corresponding 0/1 indicator vector.

where the last equality uses the facts that $\rho = 1 - 2\delta'$ and that $\arccos(1 - x) = O(\sqrt{x})$ for small positive values of x , and the last inequality follows from the setting of δ' in (\dagger) . \square

We now give the (straightforward) proof of Part (ii) of Theorem 1.3.

Proof of Part (ii) of Theorem 1.3. Note that $\langle S^{(i)}, x^{(i)} \oplus y^{(i)} \rangle = \langle S^{(i)}, x^{(i)} \rangle \oplus \langle S^{(i)}, y^{(i)} \rangle$. Thus, when both Alice and Bob know S , Alice can send the k bits $(\langle S^{(i)}, x^{(i)} \rangle)_{i \in [k]}$ to Bob who can then output the value $f_S(x, y)$. \square

In order to prove Part (iii) of Theorem 1.3, we first define (as in [GKKS16]) a communication problem in the standard distributional model that reduces to solving the contextually-uncertain problem specified by the function class \mathcal{F}_ϵ and the distribution $\mu_{2\epsilon-2\epsilon^2}$. For distributions ϕ and ψ , we denote by $\phi \otimes \psi$ the joint distribution of a sample from ϕ and an independent sample from ψ . The new problem is defined as follows.

Inputs: Alice's input is a pair (S, x) where $S \triangleq (S^{(i)} \subseteq [n])_{i \in [k]}$ and $x \in \{0, 1\}^{k \cdot n}$. Bob's input is a pair (T, y) where $T \triangleq (T^{(i)} \subseteq [n])_{i \in [k]}$ and $y \in \{0, 1\}^{k \cdot n}$.

Distribution: Let \mathcal{D}_q be the distribution on the pair (S, T) of sequences of k subsets of $[n]$, which is defined by independently setting, for each $i \in [k]$, $S^{(i)}$ to be a uniformly-random subset of $[n]$, and $T^{(i)}$ to be a q -noisy copy of $S^{(i)}$. The distribution on the inputs of Alice and Bob is then given by $\nu_\epsilon \triangleq \mathcal{D}_\epsilon \otimes \mu_{2\epsilon-2\epsilon^2}$ with $\epsilon = \sqrt{\delta'/n}$.

Function: The goal is to compute the function $F : \{0, 1\}^{2kn} \times \{0, 1\}^{2kn} \rightarrow \{0, 1\}$ defined by $F((S, x), (T, y)) = f_T(x, y) = \text{Sign}(\sum_{i \in [k]} (-1)^{\langle T^{(i)}, x^{(i)} \oplus y^{(i)} \rangle})$.

The next proposition follows from a simple application of the Chernoff bound.

Proposition 5.1. *For any $\theta > 0$, $\text{owPubCCU}_\theta^{\mu_{2\epsilon-2\epsilon^2}}(\mathcal{F}_\epsilon) \geq \text{owCC}_{\theta+\theta'}^{\nu_\epsilon}(F)$ with $\theta' = 2^{-\Theta(\epsilon \cdot n)}$.*

We will prove the following lower bound on $\text{owCC}_\theta^{\nu_\epsilon}(F)$, which along with Proposition 5.1 and the settings of ϵ and δ' in (\dagger) , implies Part (iii) of Theorem 1.3:

Lemma 5.2. *For every sufficiently small positive constant θ , $\text{owCC}_\theta^{\nu_\epsilon}(F) = \Omega(k \cdot \epsilon \cdot n)$.*

We now prove Lemma 5.2 (which is the bulk of the proof of Theorem 1.3). Subsection 5.1 summarizes some known results that we use in Subsections 5.2 and 5.3. In Subsection 5.2, we prove a ‘‘Simulation Lemma’’ that will be useful to us. In Subsection 5.3, we prove Lemma 5.2.

5.1 Proof Preliminaries for Subsections 5.2 and 5.3

In this subsection, we state some tools and known results that we use in the proofs in Subsections 5.2 and 5.3. We use the following strong direct product theorem for discrepancy of [LSS08]¹⁵.

¹⁵We point out that the statement of Corollary 23 of [LSS08] has a small inaccuracy: the additive $2^{-k \cdot (1 - H_b(\tau))}$ term in Lemma 5.3 is missing. This term is clearly needed as one can always guess $f^{(k)}$ with probability 2^{-k} . The statement that we use (Lemma 5.3) can be obtained by combining Theorem 22 of [LSS08] and the proof of Proposition 1.4 of [VW08].

Lemma 5.3 (Corollary 23 of [LSS08]). *Let $f : X \times Y \rightarrow \{0, 1\}$ be a Boolean function and P a probability distribution over $X \times Y$. If $\text{CC}_{1/2-w/2}^P(f) \geq C$ is proved using the discrepancy method, then the success probability under distribution $P^{\otimes k}$ of any $kC/3$ bit protocol computing the vector of solutions $f^{(k)}$ is at most $(8w)^{\tau \cdot k} + 2^{-k \cdot (1-H_b(\tau))}$ where τ is any positive constant less than 0.5.*

Let ξ_ϵ be the distribution that is obtained by projecting ν_ϵ on one of the k blocks and marginalizing over the remaining $k-1$ blocks. Namely, $\nu_\epsilon = \xi_\epsilon^{\otimes k}$. Define the “base function” $G((\tilde{S}, \tilde{x}), (\tilde{T}, \tilde{y}))$ by $G((\tilde{S}, \tilde{x}), (\tilde{T}, \tilde{y})) = \langle \tilde{T}, \tilde{x} \oplus \tilde{y} \rangle$ for every $\tilde{S}, \tilde{T} \subseteq [n]$ and $\tilde{x}, \tilde{y} \in \{0, 1\}^n$. We will use the following lower-bound on the distributional complexity of G over ξ_ϵ that was proved in [GKKS16] using the discrepancy method.

Lemma 5.4 ([GKKS16]). *For any $w = 2^{-o(\sqrt{\delta'n})}$, we have that $\text{CC}_{1/2-w/2}^{\xi_\epsilon}(G) \geq \Omega(\epsilon \cdot n)$, and it is proved using the discrepancy method.*

Combining Lemma 5.3 and Lemma 5.4 implies the next corollary.

Corollary 5.5. *For every positive constant γ , any deterministic protocol computing $G^{(k)}$ correctly with probability at least $(0.5 + \gamma)^k$ with respect to the distribution $\xi_\epsilon^{\otimes k} = \nu_\epsilon$ should be communicating $\Omega(k \cdot \epsilon \cdot n)$ bits.*

We define the Hamming distance function $\text{HD}_k : \{0, 1\}^k \times \{0, 1\}^k \rightarrow \{0, 1\}$ as follows. For all $x, y \in \{0, 1\}^k$, $\text{HD}_k(x, y) = 1$ if the Hamming distance between x and y is at least $\lfloor k/2 \rfloor$ and $\text{HD}_k(x, y) = 0$ otherwise. Let \mathcal{U}_{2k} denote the uniform distribution on $\{0, 1\}^{2k}$.

Lemma 5.6 ([Woo07]). *For every sufficiently small $\epsilon > 0$, it holds that $\text{owCC}_\epsilon^{\mathcal{U}_{2k}}(\text{HD}_k) = \Omega(k)$.*

The next lemma of [JRS03] compresses a 1-way private-coin protocol with external information cost¹⁶ I into a 1-way deterministic protocol with communication cost $O(I)$.

Lemma 5.7 (Result 1 of [JRS03]). *Suppose that Π is a 1-way private-coin randomized protocol for $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$. Let the average error of Π under a probability distribution μ on the inputs $\mathcal{X} \times \mathcal{Y}$ be θ . Let X, Y denote the random variables corresponding to Alice’s and Bob’s inputs respectively. Let M denote the single message sent by Alice to Bob. Suppose $I(X, Y; M) \leq a$. Let $\zeta > 0$. Then, there is another deterministic 1-way protocol Π' with the following properties:*

1. *The communication cost of Π' is at most $\frac{2(a+1)}{\zeta^2} + \frac{2}{\zeta}$ bits.*
2. *The distributional error of Π' under μ is at most $\theta + 2\zeta$.*

We will also use the next lemma.

Lemma 5.8. *Let (Q, W, B) be correlated random variables with $Q \in \mathcal{Q}$, $W \in \mathcal{W}$ and $B \in \{0, 1\}^k$. Let $\alpha \in (0, 1]$ be any constant. If $I(B; W|Q) \geq \alpha \cdot k$, then there exists a positive constant β that only depends on α , and a deterministic function $E : \mathcal{Q} \times \mathcal{W} \rightarrow \{0, 1\}^k$ such that $E(Q, W) = B$ with probability at least $(0.5 + \beta)^k$ over the random choice of (Q, W, B) .*

¹⁶The external information cost of a protocol is the amount of information that it reveals about the inputs to an external observer. For a 1-way private-coin protocol, it is given by $I(X, Y; M)$ where M is the single message sent from Alice to Bob.

Proof of Lemma 5.8. Consider the deterministic function $E : \mathcal{Q} \times \mathcal{W} \rightarrow \{0, 1\}^k$ defined as follows. For each $(q, w) \in \mathcal{Q} \times \mathcal{W}$, $E(q, w)$ is set to an arbitrary element of the set

$$\arg \max_{\hat{B} \in \{0, 1\}^k} \Pr[B = \hat{B} | Q = q, W = w].$$

We now argue that $E(Q, W) = B$ with probability at least $(0.5 + \beta)^k$ over the randomness of (Q, W, B) , where β is a positive constant that only depends on α . Since $I(B; W | Q) \geq \alpha \cdot k$, we have that

$$\begin{aligned} H(B | Q, W) &= H(B | Q) - I(B; W | Q) \\ &\leq H(B | Q) - \alpha \cdot k \\ &\leq H(B) - \alpha \cdot k \\ &\leq (1 - \alpha) \cdot k, \end{aligned}$$

where the third inequality above uses the fact that conditioning does not increase entropy, and the fourth inequality follows from the fact that $B \in \{0, 1\}^k$. By an averaging argument, with probability at least $\alpha/10$ over $(q, w) \sim (Q, W)$, it should be the case that

$$H(B | Q = q, W = w) \leq (1 - \alpha/10) \cdot k. \quad (5)$$

Let $\mathcal{G} \subseteq \mathcal{Q} \times \mathcal{W}$ denote the set of all pairs (q, w) that satisfy Equation (5). We now fix $(q, w) \in \mathcal{G}$, and consider the min-entropy

$$H_{\min}(B | Q = q, W = w) \triangleq \min_{\hat{B} \in \{0, 1\}^k} \log_2 \left(\frac{1}{\Pr[B = \hat{B} | Q = q, W = w]} \right).$$

Using the fact that min-entropy lower-bounds Shannon entropy and Equation (5), we deduce that there exists $\hat{B} \triangleq \hat{B}(q, w) \in \{0, 1\}^k$ such that $\Pr[B = \hat{B} | Q = q, W = w] \geq 2^{-(1-\alpha/10) \cdot k}$. Hence, for any fixed $(q, w) \in \mathcal{G}$, conditioned on $(Q = q, W = w)$, the value $E(q, w)$ is equal to B with probability at least $2^{-(1-\alpha/10) \cdot k}$. Since the probability that $(Q, W) \in \mathcal{G}$ is at least $\alpha/10$, we conclude that $E(Q, W)$ is equal to B with probability at least

$$(\alpha/10) \cdot 2^{-(1-\alpha/10) \cdot k} \geq (0.5 + \beta)^k,$$

for some constant β that only depends on α . □

5.2 Simulation Protocol

Recall that the distribution ν_ϵ over the inputs $((S, X), (T, Y))$ to F was defined as $\nu_\epsilon \triangleq \mathcal{D}_\epsilon \otimes \mu_{2\epsilon-2\epsilon^2}$. In the following simulation lemma, the error probability will be measured w.r.t. distribution ν_ϵ whereas the information cost will be measured w.r.t. another distribution κ_ϵ over $((S, X), (T, Y))$ inputs, which is defined as $\kappa_\epsilon \triangleq \mathcal{D}_\epsilon \otimes \mu_\epsilon$.

Lemma 5.9 (Simulation Lemma). *Let Π be any deterministic 1-way protocol computing F with error at most θ on the distribution ν_ϵ over $((S, X), (T, Y))$ inputs, and let $M \triangleq M(X, S)$ be the corresponding single message that is sent from Alice to Bob under Π . Then, we have that*

$$I_{((S, X), (T, Y)) \sim \kappa_\epsilon} \left(\left(\langle T^{(i)}, X^{(i)} \rangle \right)_{i \in [k]}; M(X, S) \mid Y, T \right) \geq \beta \cdot k \quad (6)$$

for some constant $\beta > 0$ that only depends on θ .

We point out that in Lemma 5.9, the error probability is measured w.r.t. the distribution ν_ϵ while the information cost is measured w.r.t. the distribution κ_ϵ .

Definition 5.10. A sequence $\hat{T} \triangleq (\hat{T}^{(i)} \subseteq [n])_{i \in [k]}$ of subsets is said to be typical if $|\hat{T}^{(i)}| \in [n/3, 2n/3]$ for all $i \in [k]$.

Recall that the total variation distance between two distributions ϕ and ψ defined on the same finite support Ω is given by $\Delta_{TV}(\phi, \psi) = \max_{A \subseteq \Omega} |\phi(A) - \psi(A)| = 0.5 \cdot \sum_{x \in \Omega} |\phi(x) - \psi(x)|$. We will use the next lemma.

Lemma 5.11 (Closeness Lemma). For a given sequence $\hat{T} \triangleq (\hat{T}^{(i)} \subseteq [n])_{i \in [k]}$ of subsets, we define the distribution $\mu_{\hat{T}, \epsilon}$ as follows. To sample $(X, Y) \sim \mu_{\hat{T}, \epsilon}$, we independently sample $U, V \in_R \{0, 1\}^k$, $Z \in_R \{0, 1\}^{k \cdot n}$, X to be an ϵ -noisy copy of Z conditioned on $(\langle \hat{T}^{(i)}, X^{(i)} \rangle)_{i \in [k]} = U$, and Y to be an ϵ -noisy copy of Z conditioned on $(\langle \hat{T}^{(i)}, Y^{(i)} \rangle)_{i \in [k]} = V$.

Then, for every fixed typical \hat{T} , we have that

$$\Delta_{TV}(\mu_{2\epsilon - 2\epsilon^2}, \mu_{\hat{T}, \epsilon}) \leq k \cdot \exp(-\epsilon \cdot n).$$

Proof of Lemma 5.11. We denote $\langle \hat{T}, X \rangle \triangleq (\langle \hat{T}^{(i)}, X^{(i)} \rangle)_{i \in [k]}$ and similarly $\langle \hat{T}, Y \rangle \triangleq (\langle \hat{T}^{(i)}, Y^{(i)} \rangle)_{i \in [k]}$. Note that both $\langle \hat{T}, X \rangle$ and $\langle \hat{T}, Y \rangle$ are elements of $\{0, 1\}^k$. For every fixed $\hat{X}, \hat{Y} \in \{0, 1\}^{k \cdot n}$, define $\hat{U} \triangleq \langle \hat{T}, \hat{X} \rangle$ and $\hat{V} \triangleq \langle \hat{T}, \hat{Y} \rangle$. We have that

$$\mu_{2\epsilon - 2\epsilon^2}(X = \hat{X}, Y = \hat{Y}) = \frac{1}{2^{k \cdot n}} \cdot (2\epsilon - 2\epsilon^2)^{\Delta(\hat{X}, \hat{Y})} \cdot (1 - 2\epsilon + 2\epsilon^2)^{k \cdot n - \Delta(\hat{X}, \hat{Y})}$$

On the other hand, we have that

$$\begin{aligned} & \mu_{\hat{T}, \epsilon}(X = \hat{X}, Y = \hat{Y}) \\ &= \mu_{\hat{T}, \epsilon}(U = \hat{U}, V = \hat{V}, X = \hat{X}, Y = \hat{Y}) \\ &= \sum_{\hat{Z} \in \{0, 1\}^{k \cdot n}} \mu_{\hat{T}, \epsilon}(Z = \hat{Z}, U = \hat{U}, V = \hat{V}, X = \hat{X}, Y = \hat{Y}) \\ &= \sum_{\hat{Z} \in \{0, 1\}^{k \cdot n}} \mu_{\hat{T}, \epsilon}(Z = \hat{Z}, U = \hat{U}, V = \hat{V}) \cdot \mu_{\hat{T}, \epsilon}(X = \hat{X}, Y = \hat{Y} | Z = \hat{Z}, U = \hat{U}, V = \hat{V}) \\ &= \sum_{\hat{Z} \in \{0, 1\}^{k \cdot n}} \frac{1}{4^k} \cdot \frac{1}{2^{k \cdot n}} \cdot \mu_{\hat{T}, \epsilon}(X = \hat{X} | Z = \hat{Z}, U = \hat{U}) \cdot \mu_{\hat{T}, \epsilon}(Y = \hat{Y} | Z = \hat{Z}, V = \hat{V}) \end{aligned} \quad (7)$$

Denote by $N_\epsilon(Z)$ the distribution of a random variable that is an ϵ -noisy copy version of Z . Then,

$$\mu_{\hat{T}, \epsilon}(X = \hat{X} | Z = \hat{Z}, U = \hat{U}) = \frac{\Pr_{X' \sim N_\epsilon(Z)}[X' = \hat{X}]}{\Pr_{X' \sim N_\epsilon(Z)}[\langle \hat{T}, X' \rangle = \hat{U}]}, \quad (8)$$

where

$$\begin{aligned}
\Pr_{X' \sim N_\epsilon(Z)} [\langle \hat{T}, X' \rangle = \hat{U}] &= \Pr_{X' \sim N_\epsilon(Z)} [\langle \hat{T}^{(i)}, X'^{(i)} \rangle_{i \in [k]} = \hat{U}] \\
&= \prod_{i=1}^k \Pr_{X'^{(i)} \sim N_\epsilon(Z^{(i)})} [\langle \hat{T}^{(i)}, X'^{(i)} \rangle = \hat{U}_i] \\
&= \prod_{i=1}^k \left(\frac{1}{2} \pm \exp(-\epsilon \cdot n) \right) \\
&= \frac{1}{2^k} \cdot \prod_{i=1}^k (1 \pm \exp(-\epsilon \cdot n)) \\
&= \frac{1}{2^k} \cdot (1 \pm k \cdot \exp(-\epsilon \cdot n)),
\end{aligned}$$

where the third equality above follows from the fact that \hat{T} is typical. Plugging back this last expression in Equation (8), we get

$$\begin{aligned}
\mu_{\hat{T}, \epsilon}(X = \hat{X} | Z = \hat{Z}, U = \hat{U}) &= 2^k \cdot \frac{\Pr_{X' \sim N_\epsilon(Z)} [X' = \hat{X}]}{1 \pm k \cdot \exp(-\epsilon \cdot n)} \\
&= 2^k \cdot \Pr_{X' \sim N_\epsilon(Z)} [X' = \hat{X}] \cdot (1 \pm k \cdot \exp(-\epsilon \cdot n)). \tag{9}
\end{aligned}$$

Similarly, we have that

$$\mu_{\hat{T}, \epsilon}(Y = \hat{Y} | Z = \hat{Z}, V = \hat{V}) = 2^k \cdot \Pr_{Y' \sim N_\epsilon(Z)} [Y' = \hat{Y}] \cdot (1 \pm k \cdot \exp(-\epsilon \cdot n)). \tag{10}$$

Combining Equations (7), (9) and (10) yields

$$\begin{aligned}
\mu_{\hat{T}, \epsilon}(X = \hat{X}, Y = \hat{Y}) &= \sum_{\hat{Z} \in \{0,1\}^{k \cdot n}} \frac{1}{2^{k \cdot n}} \cdot \Pr_{X' \sim N_\epsilon(Z)} [X' = \hat{X}] \cdot \Pr_{Y' \sim N_\epsilon(Z)} [Y' = \hat{Y}] \cdot (1 \pm k \cdot \exp(-\epsilon \cdot n))^2 \\
&= (1 \pm k \cdot \exp(-\epsilon \cdot n)) \cdot \sum_{\hat{Z} \in \{0,1\}^{k \cdot n}} \frac{1}{2^{k \cdot n}} \cdot \Pr_{X' \sim N_\epsilon(Z)} [X' = \hat{X}] \cdot \Pr_{Y' \sim N_\epsilon(Z)} [Y' = \hat{Y}] \\
&= (1 \pm k \cdot \exp(-\epsilon \cdot n)) \cdot \mu_{2\epsilon - 2\epsilon^2}(X = \hat{X}, Y = \hat{Y}), \tag{11}
\end{aligned}$$

where the last equality above follows from the fact that one way to sample a $(2\epsilon - 2\epsilon^2)$ -noisy pair (X, Y) is to first sample a uniform-random Z , and then independently sample each of X and Y to be an ϵ -noisy copy of Z . Using Equation (11) and the definition of the total-variation distance, we conclude that

$$\begin{aligned}
\Delta_{TV}(\mu_{2\epsilon - 2\epsilon^2}, \mu_{\hat{T}, \epsilon}) &= \frac{1}{2} \cdot \sum_{\hat{X}, \hat{Y} \in \{0,1\}^{k \cdot n}} |\mu_{2\epsilon - 2\epsilon^2}(X = \hat{X}, Y = \hat{Y}) - \mu_{\hat{T}, \epsilon}(X = \hat{X}, Y = \hat{Y})| \\
&\leq k \cdot \exp(-\epsilon \cdot n) \sum_{\hat{X}, \hat{Y} \in \{0,1\}^{k \cdot n}} \mu_{2\epsilon - 2\epsilon^2}(X = \hat{X}, Y = \hat{Y}) \\
&= k \cdot \exp(-\epsilon \cdot n). \quad \square
\end{aligned}$$

We are now ready to prove Lemma 5.9.

Proof of Lemma 5.9. Assume for the sake of contradiction that there exists a deterministic 1-way protocol Π that computes F with error at most θ on ν_ϵ and that violates Equation (6). Namely, if we define the *intermediate information cost* of Π as

$$\text{IIC}_{\kappa_\epsilon}(\Pi | T) \triangleq I_{((S,X),(T,Y)) \sim \kappa_\epsilon} \left((\langle T^{(i)}, X^{(i)} \rangle)_{i \in [k]}; M(X, S) | Y, T \right), \quad (12)$$

then we assume that $\text{IIC}_{\kappa_\epsilon}(\Pi | T) = o(k)$. We say that a particular value \hat{T} of T is *nice* if it simultaneously satisfies the following three properties:

1. \hat{T} is typical.
2. The conditional error of Π w.r.t. to ν_ϵ conditioned on $T = \hat{T}$ is at most $O(\theta)$.
3. The intermediate information cost conditioned on $T = \hat{T}$ satisfies

$$\text{IIC}_{\kappa_\epsilon}(\Pi | T = \hat{T}) = O\left(\text{IIC}_{\kappa_\epsilon}(\Pi | T)\right) = o(k).$$

We now argue that there exists a \hat{T} that is nice. To do so, we show that a random \hat{T} satisfies the above three properties with high probability. First, by Definition 5.10, a Chernoff bound and union bound, a random \hat{T} satisfies property 1 with probability at least $1 - o(1)$ as long as $k \cdot \exp(-n) = o(1)$. Moreover, by an averaging argument, a random \hat{T} satisfies property 2 with probability $1 - o(1)$. Finally, by an averaging argument and the definition of the conditional mutual information in Equation (12), we get that a random \hat{T} satisfies property 3 with probability $1 - o(1)$. By a union bound, we conclude that a random \hat{T} satisfies all three properties with high probability. Henceforth, we fix such a nice \hat{T} and use it to give a deterministic 1-way protocol computing the function HD_k w.h.p. over the uniform distribution on $\{0, 1\}^{2k}$ and with communication $o(k)$ bits. This would contradict the lower bound of [Woo07] (i.e., Lemma 5.6).

Consider the simulation protocol Π' described in Protocol 2. In this protocol, Alice is given as input a binary string U of length k and Bob is given as input a binary string V of length k . We will argue that

- (a) The output of Π' is equal to $\text{HD}_k(U, V)$ with probability $1 - O(\theta)$ over the randomness of $(U, V) \sim U_{2k}$ and over the private and shared randomness of Π' .
- (b) The information cost of Π' satisfies

$$I(U; M'(U) | V, R) = \text{IIC}_{\kappa_\epsilon}(\Pi | T = \hat{T}),$$

where M' is the single (randomized) message sent from Alice to Bob under Π' , and R is the public randomness of Π' .

We start by proving property (a). Let λ be the probability distribution of the sequence S of subsets that is sampled in Protocol 2. In other words, $S \sim \lambda$ is an ϵ -noisy copy of \hat{T} . Then, when (U, V) is drawn uniformly at random, the induced distribution on (S, X, Y) in Protocol 2 is $\lambda \otimes \mu_{\hat{T}, \epsilon}$. Property 2 above guaranteed that the error probability of protocol Π on pairs $((S, X), (\hat{T}, Y))$ such

that $(S, X, Y) \sim \lambda \otimes \mu_\epsilon$ is at most $O(\theta)$. Using Lemma 5.11, the fact that Protocol Π' simulates Π and the fact that

$$F((S, X), (\hat{T}, Y)) = \text{HD}_k \left((\langle \hat{T}^{(i)}, X^{(i)} \rangle)_{i \in [k]}, (\langle \hat{T}^{(i)}, Y^{(i)} \rangle)_{i \in [k]} \right),$$

we get that the error probability of Π' (over the randomness of $(U, V) \sim \mathcal{U}_{2k}$ and over the private and shared randomness) is at most $O(\theta) + O(k \cdot \exp(-\epsilon n))$, which is $O(\theta)$.

We next prove property (b). The information cost of Π' is given by

$$\begin{aligned} I(U; M'(U) \mid V, R) &= I((\langle \hat{T}^{(i)}, X^{(i)} \rangle)_{i \in [k]}; M(S, X) \mid V, Z) \\ &= I((\langle \hat{T}^{(i)}, X^{(i)} \rangle)_{i \in [k]}; M(S, X) \mid Z) \\ &= I((\langle \hat{T}^{(i)}, X^{(i)} \rangle)_{i \in [k]}; M(S, X) \mid Z, T = \hat{T}) \\ &= \text{IC}_{\kappa_\epsilon}(\Pi \mid T = \hat{T}) \end{aligned}$$

where the second equality above follows from the fact that (S, X) and V are conditionally independent given Z , and the last equality follows from the fact that $(X, Z) \sim \mu_\epsilon$ in Protocol 2.

To sum up, the 1-way protocol Π' computes HD_k with error probability at most $O(\theta)$ and has information cost $o(k)$ bits (by Properties 3 and (b) above). By averaging over the shared randomness, we can convert Π' into a 1-way private-coin protocol Π'' with the same error and information cost guarantees. Note that since U and V are independent, we have that $I(U, V; M''(U)) = I(U; M''(U) \mid V)$ where M'' is the single message sent from Alice to Bob under Π'' . Applying the generic compression result of [JRS03] (i.e., Lemma 5.7) with error parameter $\zeta = \theta$, we get that there exists a 1-way deterministic protocol Π''' that computes HD_k with error probability at most $O(\theta)$ over the uniform distribution and with communication cost $o(k)$ bits. This contradicts the lower bound of Woodruff [Woo07] (i.e., Lemma 5.6).

Protocol 2 Simulation Protocol Π'

Inputs. Alice is given $U \in \{0, 1\}^k$ and Bob is given $V \in \{0, 1\}^k$.

Parameters. A fixed sequence $\hat{T} \triangleq (\hat{T}^{(i)} \subseteq [n])_{i \in [k]}$ of subsets and noise parameters $\epsilon, q > 0$.

- 1: Alice and Bob use their shared randomness to sample $Z \in_R \{0, 1\}^{k \cdot n}$.
 - 2: Alice uses her private randomness to sample $X \in \{0, 1\}^{k \cdot n}$ to be an ϵ -noisy copy of Z conditioned on $(\langle \hat{T}^{(i)}, X^{(i)} \rangle)_{i \in [k]} = U$.
 - 3: Bob uses his private randomness to sample $Y \in \{0, 1\}^{k \cdot n}$ to be an ϵ -noisy copy of Z conditioned on $(\langle \hat{T}^{(i)}, Y^{(i)} \rangle)_{i \in [k]} = V$.
 - 4: Alice uses her private randomness to sample a sequence $S \triangleq (S^{(i)} \subseteq [n])_{i \in [k]}$ of subsets which is set to be a q -noisy copy of \hat{T} .
 - 5: Alice and Bob simulate the 1-way deterministic protocol Π on inputs $((S, X), (\hat{T}, Y))$ and return the resulting output.
-

□

5.3 Proof of Lemma 5.2

Assume for the sake of contradiction that there is a deterministic 1-way protocol Π computing F with error at most θ over the distribution ν_ϵ , and that has communication cost $o(k \cdot \epsilon \cdot n)$ bits. Let

$M \triangleq M(X, S)$ be the single message that is sent from Alice to Bob under Π . By the Simulation Lemma 5.9, we should have that

$$I_{((S,X),(T,Y)) \sim \kappa_\epsilon} \left((\langle T^{(i)}, X^{(i)} \rangle)_{i \in [k]}; M(X, S) \mid Y, T \right) \geq \beta \cdot k \quad (13)$$

for some constant β that only depends on θ .

By Lemma 5.8 and Equation (13), there exists a deterministic function $E(Y, T, M(X, S)) \in \{0, 1\}^k$ such that $E(Y, T, M(X, S)) = (\langle T^{(i)}, X^{(i)} \rangle)_{i \in [k]}$ with probability at least $(0.5 + \gamma)^k$ for some positive constant γ that only depends on θ . Hence, by applying the function $E(Y, T, M(X, S))$ to his inputs (Y, T) and to the message $M(X, S)$ that he receives from Alice, Bob can guess the sequence $(\langle T^{(i)}, X^{(i)} \rangle)_{i \in [k]}$ with probability $(0.5 + \gamma)^k$. By Corollary 5.5 – which combines the strong direct product theorem for discrepancy of [LSS08] (i.e., Lemma 5.3) and the base lower bound of [GKKS16] that was proved using the discrepancy method (i.e., Lemma 5.4) – we conclude that the protocol Π should have communication cost $\Omega(k \cdot \epsilon \cdot n)$ bits.

6 Proof of Theorem 1.2

We start by recalling the statement of Theorem 1.2.

Theorem 1.2. *Let $\rho \in (0, 1]$ and μ be a product distribution. Let \mathcal{F} consist of pairs (f, g) of functions with $\Delta_\mu(f, g) \leq \delta$, and $\text{owCC}_\epsilon^\mu(f), \text{owCC}_\epsilon^\mu(g) \leq k$. Then, for every positive θ , $\text{owlsrCCU}_{\epsilon+2\delta+\theta}^\mu(\mathcal{F}) \leq O_\theta(k/\rho^2)$.*

In order to prove Theorem 1.2, we start by defining a communication problem that will be useful to us (a similar definition is used in [CGMS15]).

Definition 6.1 (Gap Inner Product; $\text{GIP}_{c,s}^d$). *Let $-1 \leq s < c \leq 1$. Let Alice be given a vector $u \in \{\pm 1\}^d$ and Bob be given a vector $v \in \{\pm 1\}^d$. The goal is for Alice and Bob to distinguish the case where $\mathbb{E}_{i \in_R [d]}[u_i v_i] \geq c$ from the case where $\mathbb{E}_{i \in_R [d]}[u_i v_i] \leq s$.*

We let ISR_ρ denote the source of *imperfect shared randomness* where Alice gets a string r of independent uniform-random bits and Bob gets a string r' of bits obtained by independently flipping each coordinate of r with probability $(1 - \rho)/2$. For a function f , we denote by $\text{owlsrCC}_{\epsilon,\rho}(f)$ the minimum cost of a protocol that has access to ISR_ρ and that on input (x, y) in the domain of f , outputs $f(x, y)$ with probability at least $1 - \epsilon$, where the probability is over the randomness of ISR_ρ .

The following theorem –which upper bounds the communication complexity with ISR_ρ of $\text{GIP}_{c,s}^d$ – was proved by [CGMS15] using a locality-sensitive-hashing based protocol.

Theorem 6.2 ([CGMS15]). *Let $\rho \in (0, 1]$. Then, $\text{owlsrCC}_{\epsilon,\rho}(\text{GIP}_{c,s}^d) = O((c - s)^{-2} \rho^{-2} \log(1/\epsilon))$ via a protocol where Alice’s message depends only her input, her part of the randomness, the values of ρ and ϵ and the difference $c - s$.*

Theorem 6.2 can be used in order to estimate the weighted inner product of two vectors up to an arbitrary additive accuracy, and when the weighting is done according to an arbitrary distribution on coordinates that is known to both Alice and Bob.

Lemma 6.3. *Let $t, d \in \mathbb{N}$ and P be a distribution over $[d]$ that is known to both Alice and Bob. Let Alice be given a vector $u \in \{\pm 1\}^d$ and Bob be given t vectors $v^{(1)}, v^{(2)}, \dots, v^{(t)} \in \{\pm 1\}^d$. Let $\theta > 0$ and $\rho \in (0, 1]$ be given. Then, there exists a 1-way protocol with communication cost $O(\theta^{-2}\rho^{-2} \log(t/\theta))$ bits such that, with probability $1 - \theta$, for every $j \in [t]$, Bob computes $\mathbb{E}_{i \sim P}[u_i v_i^{(j)}]$ up to an additive accuracy of θ .*

Proof. First, note that we can reduce the case of general distributions P on $[d]$ to the case of the uniform distribution on $[d']$ for some integer $d' \in \mathbb{N}$, by having Alice and Bob repeat coordinate $i \in [d]$ of their vectors a number of times proportional to $P(i)$. More precisely, we can assume WLOG that $P(i)$ is a rational number for each $i \in [d]$ (because of the density of the rationals in the reals), and then have Alice and Bob repeat coordinate i a number of times equal to $\ell \cdot P(i)$ where ℓ is the least-common multiple of the denominators in $\{P(i) : i \in [d]\}$.

Henceforth, we assume that P is the uniform distribution on $[d']$. Alice sends a message of the protocol for GIP in Theorem 6.2 with parameters $c - s = \theta$ and $\epsilon = \theta^2/t$. Bob then divides the interval $[-1, +1]$ into $2/\theta$ sub-intervals, each of length θ . Then, he completes the protocol for GIP in Theorem 6.2 on Alice's message, for each subinterval and for each of his vectors $v^{(1)}, v^{(2)}, \dots, v^{(t)}$. For each fixed $j \in [t]$, by a union bound over the $2/\theta$ sub-intervals, we get that with probability $1 - \theta/t$, Bob can deduce the value of $\mathbb{E}_{i \sim P}[u_i v_i^{(j)}]$ up to additive accuracy θ . Another union bound over all t vectors of Bob implies that with probability $1 - \theta$, for each $j \in [t]$, he computes the value of $\mathbb{E}_{i \sim P}[u_i v_i^{(j)}]$ up to additive accuracy θ .

Moreover, by our setting of $c - s = \theta$ and $\epsilon = \theta^2/t$, we get that the communication cost of the protocol is $O(\theta^{-2}\rho^{-2} \log(t/\theta))$ bits. \square

We are now ready to prove Theorem 1.2.

Proof of Theorem 1.2. Let $\rho \in [0, 1]$ and μ be a product distribution. Consider a pair $(f, g) \in \mathcal{F}$. Since $\text{owCC}_\epsilon^\mu(g) \leq k$, there exist an integer $L \leq 2^k$ and (deterministic) functions $\pi: X \rightarrow [L]$ and $\{B_i: Y \rightarrow \{0, 1\}\}_{i \in [L]}$, such that Alice's message on input x is $\pi(x)$, and Bob's output on message i from Alice and on input y is $B_i(y)$. For every $x \in X$, we define the function $f_x: Y \rightarrow \{0, 1\}$ as $f_x(y) \triangleq f(x, y)$ for all $y \in Y$. We also denote $\tilde{f}_x(y) \triangleq (-1)^{f_x(y)}$, and similarly $\tilde{B}_i(y) \triangleq (-1)^{B_i(y)}$. The operation of the protocol is given in Protocol 3.

Protocol 3 The protocol that handles contextual uncertainty with ISR_ρ

Setting: Let μ be a product distribution over a message space $X \times Y$. Alice and Bob are given functions f and g , and inputs x and y , respectively, where $\Delta_\mu(f, g) \leq \delta$, $\text{owCC}_\epsilon^\mu(f), \text{owCC}_\epsilon^\mu(g) \leq k$ and $(x, y) \sim \mu$.

Protocol:

1. Alice and Bob run the protocol in Lemma 6.3 with $t = L$, $d = |Y|$, $P = \mu_Y$, accuracy $\theta/3$, $u = \tilde{f}_x$, $v^{(j)} = \tilde{B}_j$ for every $j \in [L]$.
 2. For every $j \in [L]$, let agr_j denote Bob's estimate of $\mathbb{E}_{y \sim \mu_Y}[\tilde{f}_x(y)\tilde{B}_j(y)]$.
 3. Bob determines $j_{\max} \triangleq \text{argmax}_{j \in [L]}\{\text{agr}_j\}$ and outputs $B_{j_{\max}}(y)$ and halts.
-

The same argument as in the analysis of the protocol of [GKKS16] (specialized for product distributions) then implies that the probability that $B_{j_{\max}}(y)$ is not equal to $g(x, y)$ is at most

$\epsilon + 2\delta + \theta$. By Lemma 6.3 and our setting of $t = L \leq 2^k$, we get that the communication cost of Protocol 3 is $O_\theta(k/\rho^2)$ bits, as desired. \square

7 Open Questions

As mentioned in Notes 1.1 and 1.1 in Section 1, significantly improving the bounds from Theorems 1.1 and 1.3 seems to require fundamentally new constructions, and is a very important question. Moreover, is there an analogue of the protocol in Theorem 1.2 for *non-product* distributions?

Another very important and intriguing open question is whether efficient communication under contextual uncertainty is possible in the multi-round setup. Namely, if k is the r -round certain communication, can we upper bound the r -round uncertain communication by some function of k , I and possibly r ? Even for $r = 2$ and when the uncertain protocol is allowed to use public randomness, no non-trivial protocols are known in this setting. On the other hand, no separations are known for this case (beyond those known for $r = 1$) even if the protocols are restricted to be deterministic.

Acknowledgements. The authors would like to thank Ilan Komargodski, Pravesh Kothari and Mohsen Ghaffari and the anonymous reviewers for very helpful discussions and pointers.

References

- [ADW01] Andris Ambainis and Ronald De Wolf. Average-case quantum query complexity. *Journal of Physics A: Mathematical and General*, 34(35):6741, 2001. 8
- [BCK⁺14] Joshua Brody, Amit Chakrabarti, Ranganath Kondapally, David P Woodruff, and Grigory Yaroslavtsev. Beyond set disjointness: the communication complexity of finding the intersection. In *Proceedings of the 2014 ACM symposium on Principles of distributed computing*, pages 106–113. ACM, 2014. 7, 16
- [BGI14] Mohammad Bavarian, Dmitry Gavinsky, and Tsuyoshi Ito. On the role of shared randomness in simultaneous communication. In *Automata, Languages, and Programming*, pages 150–162. Springer, 2014. 3
- [BR86] Rabindra N Bhattacharya and Ramaswamy Ranga Rao. *Normal approximation and asymptotic expansions*, volume 64. SIAM, 1986. 30
- [CGMS15] Clément Louis Canonne, Venkatesan Guruswami, Raghu Meka, and Madhu Sudan. Communication with imperfectly shared randomness. In *Innovations in Theoretical Computer Science, ITCS*, pages 257–262, 2015. 1, 3, 25
- [GJS12] Oded Goldreich, Brendan Juba, and Madhu Sudan. A theory of goal-oriented communication. *J. ACM*, 59(2):8, 2012. 1
- [GKKS16] Badih Ghazi, Ilan Komargodski, Pravesh Kothari, and Madhu Sudan. Communication with contextual uncertainty. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM*

Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016, pages 2072–2085, 2016. [1](#), [2](#), [3](#), [4](#), [7](#), [8](#), [9](#), [18](#), [19](#), [25](#), [26](#)

- [GKS16] Badih Ghazi, Pritish Kamath, and Madhu Sudan. Communication complexity of permutation-invariant functions. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1902–1921, 2016. [3](#)
- [GLM⁺15] Mika Göös, Shachar Lovett, Raghu Meka, Thomas Watson, and David Zuckerman. Rectangles are nonnegative juntas. In *Proceedings of the 47th Symposium on Theory of Computing (STOC). ACM*, 2015. [8](#)
- [GPW15] Mika Göös, Toniann Pitassi, and Thomas Watson. The landscape of communication complexity classes. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 22, page 49, 2015. [8](#)
- [HS14] Elad Haramaty and Madhu Sudan. Deterministic compression with uncertain priors. In *Innovations in Theoretical Computer Science, ITCS*, pages 377–386, 2014. [1](#)
- [JKKS11] Brendan Juba, Adam Tauman Kalai, Sanjeev Khanna, and Madhu Sudan. Compression without a common prior: an information-theoretic justification for ambiguity in language. In *Innovations in Computer Science, ICS*, pages 79–86, 2011. [1](#)
- [JRS03] Rahul Jain, Jaikumar Radhakrishnan, and Pranab Sen. A direct sum theorem in communication complexity via message compression. *Lecture notes in computer science*, 2719:300–315, 2003. [19](#), [24](#)
- [JS08] Brendan Juba and Madhu Sudan. Universal semantic communication I. In *40th Annual ACM Symposium on Theory of Computing*, pages 123–132, 2008. [1](#)
- [JS11] Brendan Juba and Madhu Sudan. Efficient semantic communication via compatible beliefs. In *Innovations in Computer Science, ICS*, pages 22–31, 2011. [1](#)
- [JW13] Brendan Juba and Ryan Williams. Massive online teaching to bounded learners. In *Innovations in Theoretical Computer Science, ITCS*, pages 1–10, 2013. [1](#)
- [KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? *SIAM Journal on Computing*, 37(1):319–357, 2007. [30](#)
- [Kla01] Hartmut Klauck. Lower bounds for quantum communication complexity. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 288–297. IEEE, 2001. [8](#)
- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997. [12](#)
- [Lin92] Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992. [7](#), [11](#), [12](#)

- [LSS08] Troy Lee, Adi Shraibman, and Robert Spalek. A direct product theorem for discrepancy. In *Computational Complexity, 2008. CCC'08. 23rd Annual IEEE Conference on*, pages 71–80. IEEE, 2008. [9](#), [18](#), [19](#), [25](#)
- [LZ10] Troy Lee and Shengyu Zhang. Composition theorems in communication complexity. In *Automata, Languages and Programming*, pages 475–489. Springer, 2010. [8](#)
- [MORS10] Kevin Matulef, Ryan O’Donnell, Ronitt Rubinfeld, and Rocco A Servedio. Testing halfspaces. *SIAM Journal on Computing*, 39(5):2004–2047, 2010. [30](#)
- [MWY15] Marco Molinaro, David P Woodruff, and Grigory Yaroslavtsev. Amplification of one-way information complexity via codes and noise sensitivity. In *Automata, Languages, and Programming*, pages 960–972. Springer, 2015. [8](#)
- [New91] Ilan Newman. Private vs. common random bits in communication complexity. *Information processing letters*, 39(2):67–71, 1991. [3](#)
- [O’D14] Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014. Cambridge Books Online. [17](#)
- [She99] WF Sheppard. On the application of the theory of error to cases of normal distribution and normal correlation. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, pages 101–531, 1899. [10](#)
- [She08] Alexander A Sherstov. The pattern matrix method for lower bounds on quantum communication. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 85–94. ACM, 2008. [8](#)
- [SZ07] Yaoyun Shi and Yufan Zhu. Quantum communication complexity of block-composed functions. *arXiv preprint arXiv:0710.0095*, 2007. [8](#)
- [VW08] Emanuele Viola and Avi Wigderson. Norms, xor lemmas, and lower bounds for polynomials and protocols. *Theory of Computing*, 4(1):137–168, 2008. [18](#)
- [Wit75] Hans S Witsenhausen. On sequences of pairs of dependent random variables. *SIAM Journal on Applied Mathematics*, 28(1):100–113, 1975. [3](#)
- [Woo07] David P Woodruff. *Efficient and private distance approximation in the communication and streaming models*. PhD thesis, Citeseer, 2007. [9](#), [19](#), [23](#), [24](#)
- [Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *11h Annual ACM Symposium on Theory of Computing*, pages 209–213, 1979. [1](#), [9](#)

A Two-Dimensional Berry-Esseen Theorem for Independent Random Variables

In this section, we state a two-dimensional Berry-Esseen theorem for independent (but not necessarily identically distributed) binary random variables that we used in the proofs in Section [4](#)

(namely, in the proofs of Part (i) of Theorem 1.1 and of Lemma 4.9). It follows from a known multi-dimensional Berry-Esseen theorem and the argument is very similar to that of [MORS10], the only exceptions being that in our case the random variables are not necessarily identically distributed, and each of them takes values in $\{-1, 0, +1\}$ instead of $\{\pm 1\}$.

Theorem A.1 (Two-dimensional Berry-Esseen). *Consider the linear form $\ell(z) \triangleq k^{-1/2} \cdot \sum_{i \in [k]} z_i$ where $z \in \{-1, 0, +1\}^k$. Let $(x, y) \in \{-1, 0, +1\}^k \times \{-1, 0, +1\}^k$ such that independently for each $i \in [k]$, (x_i, y_i) is a pair of zero-mean random variables with covariance matrix Σ_i . Let $\Sigma \triangleq k^{-1} \cdot \sum_{i \in [k]} \Sigma_i$ and denote by λ the smallest eigenvalue of Σ . Then, for any intervals $I_1, I_2 \subseteq \mathbb{R}$, it holds that*

$$|\Pr[(\ell(x), \ell(y)) \in I_1 \times I_2] - \Pr[(X, Y) \in I_1 \times I_2]| \leq O\left(\frac{1}{\lambda^{3/2} \cdot \sqrt{k}}\right),$$

where (X, Y) is a pair of zero-mean Gaussians with covariance matrix Σ .

Proof. The following statement appears as Theorem 16 in [KKMO07] and as Corollary 16.3 in [BR86].

Theorem A.2. *Let X_1, \dots, X_k be independent random variables taking values in \mathbb{R}^d and satisfying:*

- $\mathbb{E}[X_j]$ is the all-zero vector for every $j \in \{1, \dots, k\}$.
- $w^{-1} \sum_{j=1}^w \text{Cov}[X_j] = \Sigma$ where Cov denotes the covariance matrix.
- λ is the smallest eigenvalue of Σ and Λ is the largest eigenvalue of Σ .
- $\rho_3 = k^{-1} \sum_{j=1}^k \mathbb{E}[|X_j|^3] < \infty$.

Let Q_k denote the distribution of $k^{-1/2}(X_1 + \dots + X_k)$, let $\Phi_{0,V}$ denote the distribution of the d -dimensional Gaussian with mean 0 and covariance matrix Σ , and let $\eta = C\lambda^{-3/2}\rho_3 k^{-1/2}$, where C is a certain universal constant. Then, for any Borel set A ,

$$|Q_n(A) - \Phi_{0,V}(A)| \leq \eta + B(A),$$

where $B(A)$ is the following measure of the boundary of A : $B(A) = 2 \sup_{y \in \mathbb{R}^d} \Phi_{0,V}((\partial A)^{\eta'} + y)$, $\eta' = \Lambda^{1/2}\eta$ and $(\partial A)^{\eta'}$ denotes the set of points within distance η' of the topological boundary of A .

We now apply Theorem A.2 with $d = 2$ in order to complete the proof of Theorem A.1. We are given that for every $i \in \{1, \dots, k\}$, $\mathbb{E}[X_i] = \mathbb{E}[Y_i] = 0$ and $\text{Cov}[(X_i, Y_i)] = \Sigma_i$. Thus, $k^{-1} \cdot \sum_{j=1}^k \text{Cov}[(X_j, Y_j)] = \sum_{j=1}^k \Sigma_j = \Sigma$. Note that the largest eigenvalue of Σ is $\Lambda = O(1)$. Moreover, since each coordinate of our random variables is $\{-1, 0, +1\}$ -valued, for every $j \in \{1, \dots, k\}$, $\mathbb{E}[|X_j|^3] \leq 2^{3/2}$. Thus, $\rho_3 \leq 2^{3/2}$. Hence, $\eta = O(\lambda^{-3/2}k^{-1/2})$. As in [KKMO07, MORS10], one can check that the topological boundary of any set of the form $I_1 \times I_2$ (where $I_1, I_2 \subseteq \mathbb{R}$ are intervals) is $O(\eta')$. Since $\eta' = \Lambda^{1/2}\eta = O(\eta)$, Theorem A.1 follows. \square