

Communication with Imperfectly Shared Randomness

Clément L. Canonne* Venkatesan Guruswami† Raghu Meka‡ Madhu Sudan§

November 14, 2014

Abstract

The communication complexity of many fundamental problems reduces greatly when the communicating parties share randomness that is independent of the inputs to the communication task. Natural communication processes (say between humans) however often involve large amounts of shared correlations among the communicating players, but rarely allow for perfect sharing of randomness. Can the communication complexity benefit from shared correlations as well as it does from shared randomness? This question was considered mainly in the context of simultaneous communication by Bavarian et al. [1]. In this work we study this problem in the standard interactive setting and give some general results. In particular, we show that every problem with communication complexity of k bits with perfectly shared randomness has a protocol using imperfectly shared randomness with complexity $\exp(k)$ bits. We also show that this is best possible by exhibiting a promise problem with complexity k bits with perfectly shared randomness which requires $\exp(k)$ bits when the randomness is imperfectly shared. Along the way we also highlight some other basic problems such as compression, and agreement distillation, where shared randomness plays a central role and analyze the complexity of these problems in the imperfectly shared randomness model.

The technical highlight of this work is the lower bound that goes into the result showing the tightness of our general connection. This result builds on the intuition that communication with imperfectly shared randomness needs to be less sensitive to its random inputs than communication with perfectly shared randomness. The formal proof invokes results about the small-set expansion of the noisy hypercube and an invariance principle to convert this intuition to a proof, thus giving a new application domain for these fundamental results.

*Columbia University. Email: cannonne@cs.columbia.edu. Research supported in part by NSF CCF-1115703 and NSF CCF-1319788. Some of this work was done when the author was an intern at Microsoft Research New England.

†Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA. venkatg@cs.cmu.edu. Some of this work was done when the author was visiting Microsoft Research New England. Research supported in part by NSF CCF-0963975.

‡Microsoft Research, 1288 Pear Avenue, Mountain View, CA 94043, USA. meka@microsoft.com.

§Microsoft Research, 1 Memorial Drive, Cambridge, MA 02142, USA. madhu@mit.edu.

Contents

1	Introduction	1
2	Model, Formal Description of Results and Main Ideas	3
2.1	Model	3
2.2	Problems, Results and Techniques	4
2.2.1	Compression	4
2.2.2	Agreement distillation	4
2.2.3	General relationships between perfect and imperfect sharing	5
3	Compression	7
4	Agreement Distillation	8
5	General connection between perfect and imperfect shared randomness	9
5.1	Communication Strategies: Inner Products and Convexity	9
5.2	Upper bound on ISR in terms of PSR	13
5.3	ISR lower bound for SPARSEGAPINNERPRODUCT	15
6	Proof of Theorem 5.8	16
6.1	Proof setup	16
6.2	Overview of proof of Theorem 5.8	17
6.3	Background on influence of variables	17
6.4	Proof of Theorem 5.8	19
7	Low-influence communication strategies	22
7.1	Lower bound for Gaussian Inner Product	24
7.2	Putting things together and proof of Theorem 6.8	25
8	Conclusions	26
A	Proofs from Section 7	29

1 Introduction

The availability of shared randomness can lead to enormous savings in communication complexity when computing some basic functions whose inputs are spread out over different communicating players. A basic example of this is Equality Testing, where two players Alice and Bob have inputs $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^n$ and need to determine if $x = y$. Deterministically this takes n bits of communication. This reduces to $\Theta(\log n)$ bits if Alice and Bob can toss coins and they are allowed some error. But if they share some randomness $r \in \{0, 1\}^*$ independent of x and y then the communication cost drops to $O(1)$. (See, for instance, [10]).

A more prevalent example of a communication problem is compression with uncertain priors. Here Alice has a distribution P on a universe $[N] = \{1, \dots, N\}$, and a message $m \in [N]$ chosen according to the distribution P . Alice is allowed to send some bits to Bob and Bob should output m and the goal is to minimize the expected number of bits that Alice sends Bob (over the random choice of m). If Bob knows the distribution P exactly then this is the classical compression problem, solved for example by Huffman coding. In most forms of natural communication (e.g., think about the next email you are about to send), Alice and Bob are not perfectly aware of the underlying context to their exchange, but have reasonably good ideas about each other. One way to model this is to say that Bob has a distribution Q that is *close* to the distribution P that Alice is working with, but is not identical to P . Compressing information down to its entropy in the presence of such uncertainty (i.e., $P \neq Q$) turns out to be possible if Alice and Bob share randomness that is independent of (P, Q, m) as shown by Juba et al. [8]. However it remains open as to whether such compression can be effected deterministically, without the shared randomness — the best known schemes can only achieve a compression length of roughly $O(H(P) + \log \log N)$, where $H(P) = \sum_{i \in [N]} P(i) \log 1/P(i)$ denotes the entropy of P .¹

In both examples above it is natural to ask the question: can the (presumed) savings in communication be achieved in the absence of perfect sharing of randomness? The question especially makes sense in the latter context where the essential motivation is that Alice and Bob are not in perfect synchrony with each other: If Alice and Bob are not perfectly aware of the distributions P and Q , why should their randomness be identical?

The question of communication with imperfectly shared randomness was considered recently in the work of Bavarian et al. [1]. They consider the setting where Alice and Bob have randomness r and s respectively, with some known correlation between r and s , and study the implications of correlated randomness in the simultaneous message communication model (where a referee gets messages from Alice and Bob and computes some joint function of their inputs). Their technical focus is on the different kinds of correlations possible between r and s , but among basic results they show that equality testing has a $O(1)$ communication complexity protocol with correlated shared randomness.

In this work we are concerned with the setting of general communication protocols, where Alice and Bob interact to determine the value of some function. From some perspectives, this setting does not seem to offer a major difference between “private randomness” and “perfectly shared randomness” — Newman [13] shows that the communication complexity in the former setting can be larger by at most an additive $\log n$ term, where n is the input size. “Imperfectly shared randomness” being in between the

¹We stress that the setting of uncertain compression is completely different from that of compression with the “wrong distribution”, a well-studied question in information theory. In the “wrong distribution problem” (see, for instance, [3, Theorem 5.4.3]) the sender and receiver agree on the distribution, say P , but both have it wrong and the distribution the message comes from is R . This leads to a compression length of $\mathbb{E}_{m \sim R}[\log(1/P(m))] \approx H(R) + D(R||P)$. The important aspect here is that while the compression is not as good, there is no confusion between sender and receiver; and the latter is the focus of our problem.

two models cannot therefore be too far from them either. However, problems like compression above highlight a different perspective. There N is the size of the universe of all possible messages, and compression to $\log N$ bits of communication is trivial and uninteresting. Even a solution with $\log \log N$ bits of communication is not completely satisfactory. The real target is $O(H(P))$ bits of communication, which may be a constant independent of the universe size N (and for natural communication, the set of possible messages could be thought of as an infinitely large set). Thus the gap between the communication complexity with perfectly shared randomness and imperfectly shared randomness remains a very interesting question, which we explore in this paper.

We provide a formal description of our models and results in the following section, and here give an informal preview. We consider communication complexity in a simplified setting of imperfectly shared randomness: Alice has a uniform binary string $r \in \{0, 1\}^m$ and Bob has a string s obtained by flipping each bit of r independently with some tiny probability. (While this setting is not the most general possible, it seems to capture the most interesting aspects of the “lack of prior agreement” between Alice and Bob.) Our main contributions in this work are the introduction of some new problems of interest in the context of communication complexity, and a comparison of their communication complexity with/without perfect sharing of randomness.

The first problem we study is the complexity of *compression with uncertain priors*. We show that any distribution P can be compressed to $O(H(P))$ bits even when the randomness is not perfectly shared. As in the analogous result of Juba et al. [8] this protocol sheds some light on natural communication processes, and introduces an error-correcting element that was not previously explained.

The next problem we mention is that of *agreement distillation*. Here Alice and Bob try to agree on a small random string using little communication. This is a natural problem to study in the context of communication complexity with imperfect randomness, since an efficient solution for this problem would allow Alice and Bob to convert any protocol using perfectly shared randomness into one that relies only on imperfectly shared randomness. It turns out that the zero-communication version of this question, where Alice and Bob are not allowed to communicate at all with each other, was studied by Bogdanov and Mossel [2]. They give a very strong negative result for this problem, showing that the probability that Alice and Bob can agree on a k -bit string is exponentially small in k . By a simple reduction we show that this implies that $o(k)$ bits of communication are insufficient to get agreement on k bits. Conversely, we also show that Alice and Bob can get a constant factor advantage — so they can communicate αk bits for some $\alpha < 1$. Such a result seems implicit in [2].

Returning to our work, we next attempt to get a general conversion of communication protocols from the perfectly-shared setting to the imperfectly-shared setting. We introduce a complete promise problem `GAPINNERPRODUCT` which captures two-way communication, and use it to show that any problem with k bits of communication with perfectly shared randomness also has a $\min\{\exp(k), k + \log n\}$ bit (one-way) protocol with imperfectly shared randomness. While the protocol is simple, we feel its existence is somewhat surprising; and indeed it yields a very different protocol for equality testing when compared with Bavarian et al. [1].

Lastly, our *main technical result* is a matching lower bound giving a parameterized family of promise problems, `SPARSEGAPINNERPRODUCT`, where the k 'th problem can be solved with k bits of communication with perfect randomness, but requires $\exp(\Omega(k))$ bits with imperfect sharing. This result builds a new connection between influence of variables and communication complexity, which may be of independent interest. Finally we conclude with a variety of open questions.

2 Model, Formal Description of Results and Main Ideas

Throughout the paper, we denote by \mathbb{Z}^+ the set of positive integers, and by $[n]$ the set $\{1, \dots, n\}$. Unless specified otherwise, all logarithms are in base 2. We also recall, for $x \in [0, 1]$, the definition of the binary entropy function $h(x) = -x \log x - (1-x) \log(1-x)$; furthermore, for any $p \in [0, 1]$, we will write $\text{Bern}(p)$ for the Bernoulli distribution on $\{0, 1\}$ with parameter p , and $\text{Bern}^n(p)$ for the product distribution on $\{0, 1\}^n$ of n independent Bernoulli random variables. For a distribution P over a domain Ω , we write $H(P) = \sum_{x \in \Omega} P(x) \log(1/P(x))$ for its entropy, and $x \sim P$ to indicate that x is drawn from P . \mathcal{U}_Ω denotes the uniform distribution over Ω .

Finally, for two elements $x, y \in \{+1, -1\}^n$, their *Hamming distance* $\text{dist}(x, y)$ is defined as the number of coordinates in which they differ (and similarly for $x, y \in \{0, 1\}^n$).

2.1 Model

We use the familiar model of communication complexity, augmented by the notion of correlated shared randomness. Recall that in the standard model, two players, Alice and Bob, have access to inputs x and y respectively. A protocol Π specifies the interaction between Alice and Bob (who speaks when and what), and concludes with Alice and Bob producing outputs w_A and w_B respectively. A communication problem P is (informally) specified by conditions on the inputs and outputs (x, y, w_A, w_B) . In usual (promise) problems this is simply a relationship on the 4-tuple. In sampling problems, this may be given by requirements on the distribution of this output given x and y . For functional problems, $P = (f_A, f_B)$ and the conditions require that $w_A = f_A(x, y)$ and $w_B = f_B(x, y)$. A randomized protocol is said to solve a functional problem P if the outputs are correct with probability at least $2/3$. The (worst-case) complexity of a protocol Π , denoted $\text{cc}(\Pi)$ is the maximum over all x, y of the expected number of bits communicated by Π . This is the main complexity measure of interest to us, although distributional complexity will also be considered, as also any mix. (For instance, the most natural measure in compression is a max-average measure.)

We will be considering the setting where Alice and Bob have access to an arbitrarily long sequence of correlated random bits. For this definition it will be convenient to let a random bit be an element of $\{+1, -1\}$. For $\rho \in [-1, +1]$, we say a pair of bits (a, b) are ρ -correlated (uniform) bits if $\mathbb{E}[a] = \mathbb{E}[b] = 0$ and $\mathbb{E}[ab] = \rho$. We will consider the performance of protocols when given access to sequences (r, r') where each coordinate pair (r_i, r'_i) are ρ -correlated uniform bits chosen independently for each i . We shall write $r \sim_\rho r'$ for such ρ -correlated pairs.

The *communication complexity of a problem P* with access to ρ -correlated bits, denoted² $\text{isr-cc}_\rho(P)$ is the minimum over all protocols Π that solve P with access to ρ -correlated bits of $\text{cc}(\Pi)$. For integer k , we let $\text{ISR-CC}_\rho(k)$ denote the collections of problems P with $\text{isr-cc}_\rho(P) \leq k$. The one-way communication complexity and simultaneous message complexities are defined similarly (by restricting to appropriate protocols) and denoted $\text{isr-cc}_\rho^{\text{ow}}(P)$ and $\text{isr-cc}_\rho^{\text{sm}}(P)$ respectively. The corresponding complexity classes are denoted similarly by $\text{ISR-CC}_\rho^{\text{ow}}(k)$ and $\text{ISR-CC}_\rho^{\text{sm}}(k)$.

Note that when $\rho = 1$ we get the standard model of communication with shared randomness. We denote this measure by $\text{psr-cc}(P) = \text{isr-cc}_1(P)$, and write $\text{PSR-CC}(k)$ for the corresponding complexity class. Similarly, when $\rho = 0$ we get communication complexity with private randomness $\text{private-cc}(P) = \text{isr-cc}_0(P)$. We note that $\text{isr-cc}_\rho(P)$ is non-increasing in ρ . Combined with Newman's Theorem [13], we obtain:

²All throughout "isr" stands for *imperfect shared randomness*, while *psr* refers to *perfect* shared randomness.

Proposition 2.1. *For every problem P with inputs $x, y \in \{0, 1\}^n$ and $0 \leq \rho \leq \rho' \leq 1$ we have*

$$\text{psr-cc}(P) \leq \text{isr-cc}_{\rho'}(P) \leq \text{isr-cc}_{\rho}(P) \leq \text{private-cc}(P) \leq \text{psr-cc}(P) + O(\log n).$$

The proposition also holds for one-way communication, and (except for the last inequality) simultaneous messages.

2.2 Problems, Results and Techniques

We now define some of the new problems we consider in this work and describe our main results.

2.2.1 Compression

Definition 2.2 (Uncertain Compression). For $\delta > 0, \Delta \geq 0$ and integers ℓ, n , the *uncertain compression problem* $\text{COMPRESS}_{\Delta, \delta}^{\ell, n}$ is a promise problem with Alice getting as input the pair (P, m) , where $P = (P_1, \dots, P_n)$ is a probability distribution on $[n]$ and $m \in [n]$. Bob gets a probability distribution Q on $[n]$. The promises are that $H(P) \leq \ell$ and for every $i \in [n]$, $|\log(P_i/Q_i)| \leq \Delta$. The goal is for Bob to output m , i.e., $w_B = m$ with probability at least $1 - \delta$. The measure of interest here is the maximum, over (P, Q) satisfying the promise, of the expected one-way communication complexity when m is sampled according to P .

When $\Delta = 0$, this is the classical compression problem and Huffman coding achieves a compression length of at most $\ell + 1$; and this is optimal for “prefix-free” compressions. For larger values of Δ , the work of [8] gives an upper bound of $\ell + 2\Delta + O(1)$ in the setting of perfectly shared randomness (to get constant error probability). In the setting of deterministic communication or private randomness, it is open if this communication complexity can be bounded by a function of ℓ and Δ alone (without dependence on n). (The work of [5] studies the deterministic setting.) Our first result shows that the bound of [8] can be extended naturally to the setting of imperfectly shared randomness.

Theorem 2.3. *For every $\varepsilon, \delta > 0$ and $0 < \rho \leq 1$ there exists $c = c_{\varepsilon, \delta, \rho}$ such that for every ℓ, n , we have*

$$\text{isr-cc}_{\rho}^{\text{ow}}\left(\text{COMPRESS}_{\Delta, \delta}^{\ell, n}\right) \leq \frac{1+\varepsilon}{1-h((1-\rho)/2)}(H(P) + 2\Delta + c).$$

We stress that the notation $\text{isr-cc}_{\rho}^{\text{ow}}\left(\text{COMPRESS}_{\Delta, \delta}^{\ell, n}\right)$ describes the *worst-case* complexity over P with entropy $H(P) \leq \ell$ of the *expected* compression length when $m \leftarrow P$. The protocol that achieves this bound is a simple modification of the protocol of [8]. Roughly, Alice and Bob use their correlated randomness to define a “redundant and ambiguous dictionary” with words of every length for every message. Alice communicates using a word of appropriate length given the distribution P , and Bob decodes using maximum likelihood decoding given Q . The main difference in our case is that Alice and Bob work knowing their dictionaries do not match exactly (as if they spelled the same words differently) and so use even longer words during encoding and decoding with some error-correction to allow for spelling errors. Details can be found in [Section 3](#).

2.2.2 Agreement distillation

Next we turn to a very natural problem in the context of imperfect sharing of randomness. Can Alice and Bob communicate to distill a few random bits from their large collection r and r' (of correlated random bits), bits on which they can agree perfectly?

Definition 2.4 (Agreement distillation). In the AGREEMENT-DISTILLATION $_{\gamma}^k$ problem, Alice and Bob have no inputs. Their goal is to output w_A and w_B satisfying the following properties:

- (i) $\Pr[w_A = w_B] \geq \gamma$;
- (ii) $H_{\infty}(w_A) \geq k$; and
- (iii) $H_{\infty}(w_B) \geq k$

where $H_{\infty}(X) = \min_x \log \frac{1}{\Pr[X=x]}$.

The (slightly) special case of this problem where Alice and Bob are not allowed to communicate at all was considered by Bogdanov and Mossel [2]. The setting where some communication is allowed is closely related but we describe the results in our language anyway.

A trivial way to distill randomness would be for Alice to toss random coins and send their outcome to Bob. This would achieve $\gamma = 1$ and communication complexity of k for k bits of entropy. Our first proposition says that with non-trivial correlation, some savings can always be achieved over this naive protocol.

Proposition 2.5. *For every $\rho > 0$, we have $\text{isr-cc}_{\rho}^{\text{ow}}(\text{AGREEMENT-DISTILLATION}_{\gamma}^k) = (h(\frac{1-\rho}{2}) + o_k(1)) \cdot k$ with $\gamma = 1 - o_k(1)$. In particular for every $\rho > 0$ there exists $\alpha < 1$ such that for every sufficiently large k $\text{isr-cc}_{\rho}^{\text{ow}}(\text{AGREEMENT-DISTILLATION}_{1/2}^k) \leq \alpha k$.*

We prove this proposition in Section 4. We note that this proposition is similar in spirit to Theorem 4 of [2] with the difference that they use their “better strategy” to improve γ with zero-communication, but γ remains exponentially small. We use communication to increase γ close to 1.

Our next theorem says that these linear savings are the best possible: one cannot get away with $o(k)$ communication unless $\rho = 1$. This theorem follows immediately from Theorem 1 of [2] (restated as Lemma 4.1 here) and a simple reduction that converts protocols with communication to zero-communication protocols with a loss in γ .

Theorem 2.6. $\forall \rho < 1, \exists \varepsilon > 0$ such that $\text{isr-cc}_{\rho}(\text{AGREEMENT-DISTILLATION}_{\gamma}^k) \geq \varepsilon k - \log \frac{1}{\gamma}$.

Section 4 contains details of this proof.

2.2.3 General relationships between perfect and imperfect sharing

Our final target in this work is to get some general relationships for communication complexity in the settings of perfect and imperfectly shared randomness. Our upper bounds for communication complexity are obtained by considering a natural promise problem, that we call GAPINNERPRODUCT, which is a “hard problem” for communication complexity. We use a variant, SPARSEGAPINNERPRODUCT, for our lower bounds. We define both problems below.

Definition 2.7 (GAPINNERPRODUCT $_{c,s}^n$, SPARSEGAPINNERPRODUCT $_{q,c,s}^n$). The GAPINNERPRODUCT $_{c,s}^n$ problem has parameters $n \in \mathbb{Z}^+$ (dimension), and $c > s \in [0, 1]$ (completeness and soundness). Both yes- and no-instances of this problem have inputs $x, y \in \{0, 1\}^n$. An instance (x, y) is a yes-instance if $\langle x, y \rangle \geq cn$, and a no-instance if $\langle x, y \rangle < sn$. The SPARSEGAPINNERPRODUCT $_{q,c,s}^n$ is a restriction of GAPINNERPRODUCT $_{c,s}^n$ where both the yes- and the no-instances are sparse, i.e., $\|x\|_2^2 \leq n/q$.

In Proposition 5.5 we show that GAPINNERPRODUCT $_{c,s}^n$ is “hard” for PSR-CC(k) with $c = (2/3)2^{-k}$ and $s = (1/3)2^{-k}$. Then in Lemma 5.6 we show that this problem is in ISR-CC $_{\rho}^{\text{ow}}(\text{poly}(1/(c-s)))$.

Putting the two results together we get the following theorem giving a general upper bound on $\text{isr-cc}_\rho^{\text{ow}}(P)$ in terms of $\text{psr-cc}(P)$ for any promise problem P .

Theorem 2.8. $\forall \rho > 0, \exists c < \infty$ such that $\forall k$, we have $\text{PSR-CC}(k) \subseteq \text{ISR-CC}_\rho^{\text{ow}}(c^k)$.

We prove this theorem in [Section 5.2](#).

[Theorem 2.8](#) is obviously tight already because of known gaps between one-way and two-way communication complexity. For instance, it is well known that the “indexing” problem (where Alice gets a vector $x \in \{0, 1\}^n$ and Bob an index $i \in [n]$ and they wish to compute x_i) has one-way communication complexity of $\Omega(n)$ with perfectly shared randomness, while its deterministic two-way communication complexity is at most $\log n + 2$. However one could hope for tighter results capturing promise problems P with low $\text{psr-cc}^{\text{ow}}(P)$, or to give better upper bounds on $\text{isr-cc}(P)$ for P with low $\text{psr-cc}(P)$. Our next theorem rules out any further improvements to [Theorem 2.8](#) when n is sufficiently large (compared to k). We do so by focusing on the problem `SPARSEGAPINNERPRODUCT`. In [Proposition 5.7](#) we show that $\text{psr-cc}^{\text{ow}}(\text{SPARSEGAPINNERPRODUCT}_{q,c,s}^n) = O(\text{poly}(\frac{1}{q^{c-s}}) \log q)$ for every q, n and $c > s$. In particular if say $c = 1/(2q)$ and $s = 1/(4q)$ the one-way communication complexity with perfectly shared randomness reduces to $O(\log q)$, in contrast to the $\text{poly}(q)$ upper bound on the one-way communication complexity with imperfectly shared randomness from [Lemma 5.6](#).

Our main technical theorem shows that this gap is necessary for every $\rho < 1$. Specifically in [Theorem 5.8](#) we show that $\text{isr-cc}_\rho(\text{SPARSEGAPINNERPRODUCT}_{q,c=.9/q,s=.6/q}^n) = \Omega(\sqrt{q})$. Putting the two together we get a strong converse to [Theorem 2.8](#), stated below.

Theorem 2.9. *For every k , there exists a promise problem $P = (P_n)_{n \in \mathbb{Z}^+}$ such that $\text{psr-cc}^{\text{ow}}(P) \leq k$, but for every $\rho < 1$ it is the case that $\text{isr-cc}_\rho(P) = 2^{\Omega_\rho(k)}$.*

Remarks on the proofs. [Theorem 2.8](#) and [Theorem 2.9](#) are the technical highlights of this paper and we describe some of the ideas behind them here.

[Theorem 2.8](#) gives an upper bound for $\text{isr-cc}_\rho^{\text{ow}}$ for problems with low psr-cc . As such this ought to be somewhat surprising in that for known problems with low probabilistic communication complexity (notably, equality testing), the known solutions are very sensitive to perturbations of the randomness. But the formulation in terms of `GAPINNERPRODUCT` suggests that any such problem reduces to an approximate inner product calculation; and the theory of metric embeddings, and examples such as locality sensitive hashing, suggest that one can reduce the dimensionality of the problems here significantly and this may lead to some reduced complexity protocols that are also robust to the noise of the ρ -correlated vectors. This leads us to the following idea: To estimate $\langle x, y \rangle$, where $x, y \in \{0, 1\}^n$, Alice can compute $a = \langle g_1, x \rangle$ where g_1 is a random n -dimensional spherical Gaussian and send a (or the most significant bits of a) to Bob. Bob can compute $b = \langle g_2, y \rangle$ and $a \cdot b$ is an unbiased estimator (up to normalization) of $\langle x, y \rangle$ if $g_1 = g_2$. This protocol can be easily shown to be robust in that if g_2 is only ρ -correlated with g_1 , $a \cdot b$ is still a good estimator, with higher variance. And it is easy to convert a collection of ρ -correlated bits to ρ -correlated Gaussians, so it is possible for Alice and Bob to generate the g_1 and g_2 as desired from their imperfectly shared randomness. A careful analysis (of a variant of this protocol) shows that to estimate $\langle x, y \rangle$ to within an additive error $\epsilon \|x\|_2 \|y\|_2$, it suffices for Alice to send about $1/\epsilon^2$ bits to Bob, and this leads to a proof of [Theorem 2.8](#).

Next we turn to the proof of [Theorem 2.9](#), which shows a roughly matching lower bound to [Theorem 2.8](#) above. The insight to this proof comes from examining the “Gaussian protocol” above carefully and contrasting it with the protocol used in the perfect randomness setting. In the latter case Alice uses

the randomness to pick one (or few) coordinates of x and sends some function of these bits to Bob achieving a communication complexity of roughly $\log(1/\varepsilon)$, using the fact that only $O(\varepsilon n)$ bits of x are non-zero. In the Gaussian protocol Alice sends a very “non-junta”-like function of x to Bob; this seems robust to the perturbations of the randomness, but leads to $1/\varepsilon^2$ bits of communication. This difference in behavior suggests that perhaps functions where variables have low “influence” cannot be good strategies in the setting of perfect randomness, and indeed we manage to prove such a statement in [Theorem 6.8](#). The proof of this theorem uses a variant of the invariance principle that we prove (see [Theorem 7.1](#)), which shows that if a communication protocol with low-influences works in a “product-distributional” setting, it will also work with inputs being Gaussian and with the same moments. This turns out to be a very useful reduction. The reason that SPARSEGAPINNERPRODUCT has nice $\text{psr-cc}^{\text{ow}}$ protocols is the asymmetry between the inputs of Alice and the inputs of Bob — inputs of Alice are sparse! But with the Gaussian variables there is no notion of sparsity and indeed Alice and Bob have symmetric inputs and so one can now reduce the “disjointness” problem from communication complexity (where now Alice and Bob hold sets $A, B \subseteq [1/\varepsilon]$, and would like to distinguish $|A \cap B| = 0$ from $|A \cap B| = 1$) to the Gaussian inner product problem. Using the well-known lower bound on disjointness, we conclude that $\Omega(1/\varepsilon)$ bits of communication are necessary and this proves [Theorem 6.8](#).

Of course, all this rules out only one part of the solution space for the communication complexity problem, one where Alice and Bob use functions of low-influence. To turn this into a general lower bound we note that if Alice and Bob use functions with some very influential variables, then they should agree on which variable to use (given their randomness r and r'). Such agreement on the other hand cannot happen with too high a probability by our lower bound on AGREEMENT-DISTILLATION (from [Theorem 2.6](#)). Putting all these ingredients together gives us a proof of [Theorem 2.9](#) (see [Section 5.3](#) for more details).

Organization of the rest of the paper The rest of the paper contains details and proofs of the theorems mentioned in this section. In the next section ([Section 3](#)), we prove our isr upper bound for the “Uncertain Compression” problem, namely [Theorem 2.3](#). We then turn, in [Section 4](#), to the matching upper and lower bounds for “Agreement Distillation” as described in [Proposition 2.5](#) and [Theorem 2.6](#). [Section 5](#) contains the details of our main results relating communication with perfect and imperfect shared randomness, [Theorem 2.8](#) and [Theorem 2.9](#): we first describe an alternate characterization of communication strategies in [Section 5.1](#), which allows us to treat them as vectors in (carefully defined) convex sets. This enables us to use ideas and machinery from Gaussian analysis: in particular, our lower bound on isr presented in [Section 6](#) relies on a new invariance theorem, [Theorem 7.1](#), that we prove in [Section 7](#).

3 Compression

In this section, we prove [Theorem 2.3](#), restated below:

Theorem 2.3. *For every $\varepsilon, \delta > 0$ and $0 < \rho \leq 1$ there exists $c = c_{\varepsilon, \delta, \rho}$ such that for every ℓ, n , we have*

$$\text{isr-cc}_{\rho}^{\text{ow}} \left(\text{COMPRESS}_{\Delta, \delta}^{\ell, n} \right) \leq \frac{1+\varepsilon}{1-h((1-\rho)/2)} (H(P) + 2\Delta + c).$$

Proof of [Theorem 2.3](#). Let $\mu = (1 - \rho)/2$ and $\varepsilon' > 0$ be such that $1/(1 - h(\mu + \varepsilon')) = (1 + \varepsilon)/(1 - h(\mu))$. Let $c = O(\frac{1}{\varepsilon'^2} \ln(1/\delta))$.

We interpret the random strings r and r' as two “dictionaries”, i.e., as describing words $\{w_{i,j} \in \{-1, +1\}^j\}_{j \in [n]}$ and $\{w'_{i,j} \in \{-1, +1\}^j\}_{j \in [n]}$, with the property that for every i, j and coordinate $k \in [j]$, the k th coordinates of $w_{i,j}$ and $w'_{i,j}$ are ρ -correlated.

On input P, m Alice sends $X = w_{m,j}$ to Bob where $j = \max\{c, \frac{1+\varepsilon}{1-h(\mu)} (\log(1/P(m)) + 2\Delta + \log(1/\delta))\}$. On input Q and on receiving X from Alice, Bob computes $j = |X|$ and the set

$$S_X = \{ \tilde{m} : \text{dist}(w'_{\tilde{m},j}, X) \leq (\mu + \varepsilon')j \},$$

where dist denotes the Hamming distance between strings. Bob then outputs $\arg\max_{\tilde{m} \in S_X} \{Q(\tilde{m})\}$ (so it outputs the most likely message after some error-correction).

It is clear from construction that the expected length of the communication when $m \sim P$ is at most

$$\begin{aligned} \mathbb{E}_{m \sim P} \left[\frac{1+\varepsilon}{1-h(\mu)} (\log(1/P(m)) + 2\Delta + c) \right] &= \\ \frac{1+\varepsilon}{1-h(\mu)} (\mathbb{E}_{m \sim P} [\log(1/P(m))] + 2\Delta + c) &= \frac{1+\varepsilon}{1-h(\mu)} (H(P) + 2\Delta + c). \end{aligned}$$

We finally turn to correctness, i.e., to show that Bob’s output $\tilde{m} = m$ with probability at least $1 - \delta$. First note that the probability that $m \in S_X$ is at least $(1 - \delta/2)$ (by a simple application of Chernoff bounds and the fact that j is sufficiently large compared to ε' and δ). Now let $T_m = \{m' \neq m : P(m') \geq P(m)/4^\Delta\}$. Note that $|T_m| \leq 4^\Delta/P(m)$. For any fixed $m' \in T_m$, we have that the probability (over the choice of $w'_{m',j}$) that $m' \in S_X$ is at most $2^{-(1-h(\mu+\varepsilon'))j}$. Taking the union bound over $m' \in T_m$ and plugging in our choice of j , we have that with probability at least $1 - \delta/2$, $T_m \cap S_X = \emptyset$. With probability at least $1 - \delta$ both events above happen and when they do we have $\tilde{m} = m$. \square

4 Agreement Distillation

In this section we give proofs of [Proposition 2.5](#) and [Theorem 2.6](#) which respectively give upper and lower bounds on the one-way communication complexity of randomness distillation.

We start with the upper bound, which relies on the existence of linear error-correcting codes, capable of correcting $\mu \triangleq \frac{1-\rho}{2}$ fraction errors. The fact that such codes have rate approaching $1 - h(\mu)$ yields the result that agreement distillation requires $(1 + o_k(1)) \cdot h(\mu) \cdot k$ communication for $\gamma \rightarrow 1$. Details below.

Proof of [Proposition 2.5](#). Let $\varepsilon > 0$ be any positive constant and let $\text{Bern}^k(\mu)$ be the distribution on $\{0, 1\}^k$ where each bit is independent and is 1 with probability μ . Let $\ell \in \mathbb{Z}^+$ be such that there exists a matrix $H \in \{0, 1\}^{\ell \times k}$ such that

$$\Pr_{e \sim \text{Bern}^k(\mu)} [\exists e' \neq e \text{ s.t. } \text{wt}(e') \leq (\mu + \varepsilon)k \text{ and } H \cdot e' = H \cdot e] \leq \delta/2.$$

Note a random matrix satisfies this condition for $\ell = h(\mu + \varepsilon)k$ with probability tending to 1 as k goes to ∞ .

Given ρ correlated strings $r, r' \in \{0, 1\}^k$, Alice’s output is $w_A = r$. She communicates $y = H \cdot r$ to Bob. Bob’s output is $w_B = \tilde{r}$ such that (i) $H \cdot \tilde{r} = y$ and (ii) $\text{dist}(\tilde{r}, r') \leq (\mu + \varepsilon)k$, provided \tilde{r} with these properties exists and is unique. Else he outputs r' .

It follows that unless $\text{dist}(\tilde{r}, r') > (\mu + \varepsilon)k$ or if $\exists e' \neq e \triangleq r - r'$ such that $\text{wt}(e') \leq (\mu + \varepsilon)k$ and $H \cdot e' = H \cdot e$, we have $\tilde{r} = r$. The probability of either event above is small (by Chernoff bound for the first, and by the condition on H for the second). \square

We now turn towards a proof of [Theorem 2.6](#). We first consider the setting of *zero* communication, i.e., when Alice and Bob are not allowed to communicate at all. Here we use the following lemma due to [\[2\]](#) which shows that the agreement probability γ is exponentially small in k .

Lemma 4.1 ([\[2, Theorem 1\]](#)). $\forall \rho < 1, \exists \varepsilon > 0$ such that for every zero-communication protocol for $\text{AGREEMENT-DISTILLATION}(\gamma, k)$, we have $\gamma \leq 2^{-\varepsilon k}$. (Furthermore, one can take $\varepsilon = 1 - O(\rho)$).

We now derive [Theorem 2.6](#) as an easy corollary of [Lemma 4.1](#).

Proof of [Theorem 2.6](#). Suppose Π is a c -bit communication protocol for $\text{AGREEMENT-DISTILLATION}_{\gamma}^k$. We can convert Π to a zero-bit communication protocol where Bob simply guesses the bits Alice would have sent him and Alice guesses the bits that Bob would have sent her. For each bit, the guess is correct with probability $1/2$ and so all guesses are correct with probability 2^{-c} . Conditioned on Alice and Bob guessing all bits correctly they succeed in outputting $w_B = w_A$ with probability at least γ , giving a net success probability of $2^{-c} \cdot \gamma$. Applying [Lemma 4.1](#), we get $2^{-c} \gamma \leq 2^{-\varepsilon k}$ and thus $c \geq \varepsilon k - \log(1/\gamma)$ as desired. \square

5 General connection between perfect and imperfect shared randomness

In this section we present proofs of [Theorem 2.8](#) and [Theorem 2.9](#). Key to both our upper bound on $\text{isr-cc}^{\text{ow}}(P)$ in terms of $\text{psr-cc}(P)$, and our lower bound on $\text{isr-cc}(\text{SPARSEGAPINNERPRODUCT})$ is a representation of communication strategies as vectors, where the success probability of an interaction is proportional to the inner product of these vectors. We describe this representation in [Section 5.1](#) below. We then use this representation to show that GAPINNERPRODUCT is hard for $\text{PSR-CC}(k)$ in [Section 5.2](#). We also give a one-way isr protocol for GAPINNERPRODUCT in the same section thus giving a proof of [Theorem 2.8](#). Finally in [Section 5.3](#) we give a one-way psr protocol for $\text{SPARSEGAPINNERPRODUCT}$, and then state our main technical result — an exponentially higher lower bound for it in the two-way isr setting (with the proof deferred to [Section 6](#) modulo an invariance principle which is established in [Section 7](#)). The lower bound uses the fact that the space of strategies in the vector representation forms a bounded convex set.

5.1 Communication Strategies: Inner Products and Convexity

We start by formalizing deterministic and probabilistic (private-coin) two-way communication strategies for Alice and Bob. By “strategy” we mean what Alice would do given her input and randomness, as a function of different messages that Bob may send her, and vice versa. We restrict our attention to canonical protocols in which Alice and Bob strictly alternate and communicate in bits; and the eventual outcome is a Boolean one, determined after k rounds of communication. (So the only problems that can be solved this way are “promise problems”.) Without loss of generality we also assume that the last bit communicated is the output of the communication protocol.

The natural way to define strategies would be in terms of a triple (f_A, f_B, v) where $f_A = (f_A^{2i} : \{0, 1\}^{2i} \rightarrow \{0, 1\})_{0 \leq i < k/2}$ is a sequence of functions and so is $f_B = (f_B^{2i+1} : \{0, 1\}^{2i+1} \rightarrow \{0, 1\})_{0 \leq i < k/2}$ and $v : \{0, 1\}^k \rightarrow \{0, 1\}$. The function $f_A^{2i}(h)$ determines Alice's message bit after $2i$ rounds of communication, with $h \in \{0, 1\}^{2i}$ being the transcript of the interaction thus far. Similarly the functions $f_B^{2i+1}(h)$ determine Bob's message bit after $2i + 1$ rounds of communication. Finally, v denotes the verdict function. Since we assumed that the last bit transmitted is the output, we have $v(\ell_1, \dots, \ell_k) = \ell_k$. Thus the output of an interaction is given by $v(\ell)$ where $\ell = (\ell_1, \dots, \ell_k)$ is given by $\ell_{2i+1} = f_A^{2i}(\ell_1, \dots, \ell_{2i})$ and $\ell_{2i+2} = f_B^{2i+1}(\ell_1, \dots, \ell_{2i+1})$ for $0 \leq i \leq k/2$. The interpretation is that Alice can determine the function f_A from her input and Bob can determine f_B from his input, and this allows both to determine the output after k rounds of interaction.

We will be moving on to the vector representation of strategies shortly, but first we describe probabilistic interactions, where Alice and Bob have private randomness. Such an interaction is also described by a triple (f_A, f_B, v) except that now $f_A = (f_A^{2i} : \{0, 1\}^{2i} \rightarrow [0, 1])_{0 \leq i < k/2}$ and $f_B = (f_B^{2i+1} : \{0, 1\}^{2i+1} \rightarrow [0, 1])_{0 \leq i < k/2}$. The outcome is now the random variable $v(\ell)$ where $\ell = (\ell_1, \dots, \ell_k)$ is the random variable determined inductively by letting $\ell_{2i+1} = 1$ with probability $f_A^{2i}(\ell_1, \dots, \ell_{2i})$ and $\ell_{2i+2} = 1$ with probability $f_B^{2i+1}(\ell_1, \dots, \ell_{2i+1})$ for $0 \leq i \leq k/2$.

Our vector representation of deterministic interactions is obtained by considering the set of "plausible final transcripts" that a player might see given their own strategy. Recall that the transcript of an interaction is a k -bit string and there are 2^k possible transcripts. In the new representation, we represent Alice's strategy (i.e., the functions f_A) by a vector $x \in \{0, 1\}^{2^k}$ where $x(\ell) = 1$ if and only if $\ell \in \{0, 1\}^k$ is a transcript *consistent* with Alice's strategy. (We give a more formal description shortly.) For probabilistic communication strategies (corresponding to Alice and Bob working with private randomness), we represent them by vectors x and y in $[0, 1]^{2^k}$. We formalize the set of such strategies, and verdicts, below.

In what follows we describe sets $K_A^{(k)}, K_B^{(k)} \subseteq [0, 1]^{2^k}$ that are supposed to describe the strategy space for Alice and Bob. Roughly, we wish to allow $x = (x(i_1, \dots, i_k))_{i_1, \dots, i_k \in \{0, 1\}}$ to be an "Alice strategy" (i.e., a member of K_A) if for every i_1, \dots, i_k there exists a Bob strategy such that Alice reaches the transcript i_1, \dots, i_k with probability $x(i_1, \dots, i_k)$. To describe this set explicitly we introduce auxiliary variables $x_A(i_1, \dots, i_j)$ for every $0 \leq j \leq k$ and $i_1, \dots, i_j \in \{0, 1\}$ where $x_A(i_1, \dots, i_j)$ denotes the probability (again maximized over Bob strategies) of reaching the partial transcript i_1, \dots, i_j . In what follows we first show that the auxiliary variables are linear forms in x and then show the conditions that the auxiliary variables satisfy. (We warn the reader that the first step — showing that the $x_A(\dots)$'s are linear forms in x — relies on the constraints imposed later and so some of the definition may be slightly non-intuitive.) Together the two steps allows us to show that the space of strategies is a (closed) convex set.

Definition 5.1. For vector $x \in [0, 1]^{2^k}$ and $i_1, \dots, i_j \in \{0, 1\}$ let $x_A(i_1, \dots, i_j)$ and $x_B(i_1, \dots, i_j)$ be defined as follows:

$$x_A(i_1, \dots, i_j) = \begin{cases} x(i_1, \dots, i_k) & \text{if } j = k \\ x_A(i_1, \dots, i_j, 0) + x_A(i_1, \dots, i_j, 1) & \text{if } j \text{ is even.} \\ \frac{1}{2}(x_A(i_1, \dots, i_j, 0) + x_A(i_1, \dots, i_j, 1)) & \text{if } j \text{ is odd.} \end{cases}$$

$$x_B(i_1, \dots, i_j) = \begin{cases} x(i_1, \dots, i_k) & \text{if } j = k \\ \frac{1}{2}(x_B(i_1, \dots, i_j, 0) + x_B(i_1, \dots, i_j, 1)) & \text{if } j \text{ is even.} \\ x_B(i_1, \dots, i_j, 0) + x_B(i_1, \dots, i_j, 1) & \text{if } j \text{ is odd.} \end{cases}$$

Define

$$\bar{K}_A = \bar{K}_A^{(k)} = \left\{ x \in [0, 1]^{2^k} : x_A() = 1 \text{ and } \forall \text{ odd } j, \forall i_1, \dots, i_j \in \{0, 1\}, x_A(i_1, \dots, i_j, 0) = x_A(i_1, \dots, i_j, 1) \right\},$$

and

$$\bar{K}_B = \bar{K}_B^{(k)} = \left\{ x \in [0, 1]^{2^k} : x_B() = 1 \text{ and } \forall \text{ even } j, \forall i_1, \dots, i_j \in \{0, 1\}, x_B(i_1, \dots, i_j, 0) = x_B(i_1, \dots, i_j, 1) \right\}.$$

Let $K_A = \{x * v : x \in \bar{K}_A\}$, where $v \in \{0, 1\}^{2^k}$ is given by $v_{i_1, \dots, i_k} = i_k$ (and $a * b$ denotes coordinate-wise multiplication of vectors a and b). Let $\bar{S}_A = \bar{K}_A \cap \{0, 1\}^{2^k}$, $\bar{S}_B = \bar{K}_B \cap \{0, 1\}^{2^k}$, $S_A = K_A \cap \{0, 1\}^{2^k}$, and $S_B = K_B \cap \{0, 1\}^{2^k}$.




Figure 1: Illustration of the constraints on x_A (Definition 5.1).

In what follows we first focus on deterministic communication strategies and show that \bar{S}_A , \bar{S}_B correspond to the space of deterministic communication strategies for Alice and Bob, while S_A and S_B correspond to outputs computed by such strategies. This step is not strictly needed for this paper since our main focus is on probabilistic strategies and the convex sets K_A and K_B , but the analysis of the deterministic strategies clarifies the probabilistic case.

Proposition 5.2. *\bar{S}_A and \bar{S}_B correspond to the set of deterministic communication strategies with k bits. For every strategy f_A of Alice there exists vectors $\bar{x} \in \bar{S}_A$ and $x \in S_A$ and for every strategy f_B of Bob there exist vectors $y \in \bar{S}_B$ and $y \in S_B$ such that if $\ell \in \{0, 1\}^k$ is the transcript of the interaction between Alice and Bob under strategies f_A and f_B , then ℓ is the unique sequence satisfying $\bar{x}(\ell) = \bar{y}(\ell) = 1$ and $\langle x, y \rangle = 1$ if the interaction accepts and 0 otherwise.*

Conversely every vector $x \in S_A$ corresponds to a strategy f_A for Alice (and similarly for Bob) such that Alice and Bob accept the interaction iff $\langle x, y \rangle = 1$.

Proof. Given f_A to construct \bar{x} , we let $\bar{x}(\ell) = 1$ if there exists $f_{B,\ell}$ such that the final transcript of the interaction given by f_A and $f_{B,\ell}$ is ℓ . Furthermore let $\bar{x}_A(i_1, \dots, i_j) = 1$ if there exists a Bob strategy f_{B,i_1, \dots, i_j} such that i_1, \dots, i_j is the partial transcript of the interaction between Alice and Bob. It is now straightforward to verify that the $\bar{x}_A(i_1, \dots, i_j)$ satisfy the conditions of the definition of \bar{x}_A and the conditions required for membership in \bar{K}_A . In particular we have the following three conditions: (1) $\bar{x}_A() = 1$ since the empty transcript is a legal partial transcript. (2) If j is an even index (and so Alice speaks in round $j+1$) and $x_A(i_1, \dots, i_j) = 0$ (so the partial transcript i_1, \dots, i_j is not reachable given Alice's strategy), then we must have $x_A(i_1, \dots, i_j, 0) = x_A(i_1, \dots, i_j, 1) = 0$ (no extension is reachable either). If $x_A(i_1, \dots, i_j) = 1$ then exactly one of the extensions must be reachable (based on Alice's message at

this stage) and so again we have $x_A(i_1, \dots, i_j) = x_A(i_1, \dots, i_j, 0) + x_A(i_1, \dots, i_j, 1)$. (3) If j is odd and it is Bob's turn to speak, then again if $x_A(i_1, \dots, i_j) = 0$ we have $x_A(i_1, \dots, i_j, 0) = x_A(i_1, \dots, i_j, 1) = 0$. On the other hand if $x_A(i_1, \dots, i_j) = 1$ then for each extension there exists a strategy of Bob that permits this extension and so we have $x_A(i_1, \dots, i_j, 0) = x_A(i_1, \dots, i_j, 1) = 1$ satisfying the condition for odd j . The above three conditions verify membership in \bar{K}_A and since \bar{x} is a 0/1 vector, we also have $\bar{x} \in \bar{S}_A$. The vector $x = \bar{x} * v$ gives the corresponding vector in S_A .

For the converse, the main steps are to show that a vector $x \in S_A$ corresponds to a unique vector $\bar{x} \in \bar{S}_A$ and the quantities $\bar{x}_A(i_1, \dots, i_j)$ are also in $\{0, 1\}$ where the latter is shown by induction. For the former, note that if $\bar{x} \in \bar{K}_A$ and k is even then $\bar{x}(i_1, \dots, i_{k-1}, 0) = \bar{x}(i_1, \dots, i_{k-1}, 1) = x(i_1, \dots, i_{k-1}, 1)$ and this defines the unique $\bar{x} \in \bar{K}_A$ corresponding to $x \in S_A$. On the other hand if k is odd, we first compute $x_A(i_1, \dots, i_j)$ for every $j \in \{0, \dots, k\}$ (in decreasing order of j). We then use these counts as lower bounds on $\bar{x}_A(i_1, \dots, i_j)$ and assign $\bar{x}_A(i_1, \dots, i_j)$ starting with $j = 0$ as follows. We set $\bar{x}_A() = 1$. For all larger values of j , if $\bar{x}_A(i_1, \dots, i_{j-1}) = 0$ or j is even set $\bar{x}_A(i_1, \dots, i_{j-1}, 0) = \bar{x}_A(i_1, \dots, i_{j-1}, 1) = \bar{x}_A(i_1, \dots, i_{j-1})$. If j is odd and $\bar{x}_A(i_1, \dots, i_{j-1}) = 1$ then if $x_A(i_1, \dots, i_{j-1}, 1) > 0$ then we set $\bar{x}_A(i_1, \dots, i_{j-1}, i_j) = i_j$ else we set $\bar{x}_A(i_1, \dots, i_{j-1}, i_j) = 1 - i_j$. It can be verified that this assignment leads to a $\bar{x} \in \bar{S}_A$ (and this is essentially unique except in settings where Alice rejects all paths in some subtree.)

For the latter property, we first note that for $\bar{x}_A(i_1, \dots, i_j)$ the ‘‘averaging’’ steps (j odd) are actually just equalities. i.e., if j is odd, then membership in \bar{K}_A implies that $\bar{x}_A(i_1, \dots, i_j, 0) = \bar{x}_A(i_1, \dots, i_j, 1)$ and so $\bar{x}_A(i_1, \dots, i_j) = \bar{x}_A(i_1, \dots, i_j, 0) = \bar{x}_A(i_1, \dots, i_j, 1)$. Thus by induction on $j = k$ down to 0, we get $\bar{x}_A(i_1, \dots, i_j) \in \{0, 1\}$. Using this the strategy f_A can be derived naturally: For any j , $f_A^{2j}(i_1, \dots, i_{2j}) = i$ for the unique i such that $x_A(i_1, \dots, i_{2j}, i) = 1$. \square

More significantly for us, the above equivalence also holds for probabilistic communication (i.e., with private randomness). Here the fact that the set of strategies forms a convex space is important to us.

Proposition 5.3. *K_A and K_B are closed convex sets that correspond to the set of probabilistic communication (and decision) strategies with k bits. More precisely, for every probabilistic strategy f_A of Alice there exists a vector $\bar{x} \in \bar{K}_A$ and $x \in K_A$ and for every strategy f_B of Bob there exists a vector $\bar{y} \in \bar{K}_B$ and $y \in K_B$ such that $\bar{x}(\ell) \cdot \bar{y}(\ell)$ is the probability that $\ell \in \{0, 1\}^k$ is the transcript of the interaction between Alice and Bob under strategies f_A and f_B and $\langle x, y \rangle$ is the acceptance probability of the interaction. Conversely every vector $x \in K_A$ corresponds to a probabilistic strategy f_A for Alice (and similarly for Bob, with $\langle x, y \rangle$ being the acceptance probability of the protocol).*

Proof. The fact that K_A and K_B are closed and convex sets is straightforward from their definition.

The conversion of f_A and f_B into vectors is similar to the conversion in the proof of [Proposition 5.2](#). In particular to get $\bar{x} \in \bar{K}_A$ from f_A we let $\bar{x}(i_1, \dots, i_k)$ be the maximum probability of arriving at the transcript i_1, \dots, i_k over strategies of Bob. (We omit the analysis which repeats steps of the proof of [Proposition 5.2](#).) It can also be verified (by induction on the length of partial transcripts) that for this conversion, for any pair of strategies that convert to \bar{x} for Alice and \bar{y} for Bob and for any transcript ℓ the probability of generating the transcript ℓ is exactly $\bar{x}(\ell) \cdot \bar{y}(\ell)$. It follows that the acceptance probability equals $\sum_{\ell} \bar{x}(\ell) \cdot \bar{y}(\ell) \cdot v(\ell) = \langle x, y \rangle$.

In the reverse direction, given $x \in K_A$, we first construct $\bar{x} \in \bar{K}_A$ as in the proof of [Proposition 5.2](#). Then from \bar{x} we construct the auxiliary variables $\bar{x}_A(i_1, \dots, i_j)$ for all i_1, \dots, i_j . Finally we let $f_A^{2j}(i_1, \dots, i_{2j}) = \bar{x}_A(i_1, \dots, i_{2j}, 1) / \bar{x}_A(i_1, \dots, i_{2j})$. Similarly we convert $y \in K_B$ into a strategy f_B for Bob. It can be veri-

fied that this conversion again satisfies the condition for every accepting leaf ℓ , and $x \in K_A$ and $y \in K_B$, the resulting strategies reach the leaf ℓ with probability $x(\ell) \cdot y(\ell)$ and so indeed $\langle x, y \rangle$ is the accepting probability of the resulting strategies. \square

5.2 Upper bound on ISR in terms of PSR

In this section we prove [Theorem 2.8](#). Our first step is to prove that the `GAPINNERPRODUCT` problem (with the right parameters) is hard for all problems with communication complexity k . But first we define what it means for a promise problem to be hard for some class of communication problems.

Recall that a promise problem $P = (P_n)_n$ is given by a collection of yes-instances $P_n^{\text{yes}} \subseteq \{0, 1\}^n \times \{0, 1\}^n$ and no-instances $P_n^{\text{no}} \subseteq \{0, 1\}^n \times \{0, 1\}^n$ with $P_n^{\text{yes}} \cap P_n^{\text{no}} = \emptyset$. We define below what it means for a promise problem P to reduce to a promise problem Q .

Definition 5.4. For promise problems $P = (P_n)_n$ and $Q = (Q_n)_n$ we say that P reduces to Q if there exist functions $\ell: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ and $f_n, g_n: \{0, 1\}^n \rightarrow \{0, 1\}^{\ell(n)}$ such that if $(x, y) \in P_n^{\text{yes}}$ then $(f_n(x), g_n(y)) \in Q_{\ell(n)}^{\text{yes}}$ and if $(x, y) \in P_n^{\text{no}}$ then $(f_n(x), g_n(y)) \in Q_{\ell(n)}^{\text{no}}$. We say Q is hard for a class \mathcal{C} if for every $P \in \mathcal{C}$ we have that P reduces to Q .

In other words Alice can apply f_n to her input, and Bob can apply g_n to his input and get a new pair that is an instance of the Q -problem. In particular if Q has communication complexity k , then so does P . This can be extended to functions $k(n)$ also: if Q has communication complexity $k(n)$, then P has complexity $k(\ell(n))$.

Since we are mostly interested in k being an absolute constant, we do not strictly care about the length stretching function ℓ . However, we note that in the following proposition we only need a polynomial blowup (so ℓ is a polynomial).

Proposition 5.5. *For every positive integer k , `GAPINNERPRODUCT` _{$(2/3)2^{-k}, (1/3)2^{-k}$} is hard for `PSR-CC`(k).*

Proof. Specifically we show that for any problem P with inputs of length n and $\text{psr-cc}(P) \leq k$, there exist $N = \text{poly}(n)$ and transformations f_n and g_n such that (x, y) is a yes-instance of P if and only if $(f_n(x), g_n(y))$ is a yes-instance of `GAPINNERPRODUCT` _{$(2/3)2^{-k}, (1/3)2^{-k}$} ^{N} .

Given $x \in \{0, 1\}^n$ and random string R , let $X_R \in S_A^{(k)}$ describe the communication strategy of Alice with input x and randomness R . Similarly let Y_R denote the strategy of Bob. Recall that $\langle X_R, Y_R \rangle = 1$ if the interaction accepts on randomness R and $\langle X_R, Y_R \rangle = 0$ otherwise. Let $f_n(x) = X$ be the concatenation of the strings $\{X_R\}_R$ and let $g_n(y) = Y$ be the concatenation of $\{Y_R\}_R$. By Newman's Theorem we have that the number of random strings R that we need to consider is some polynomial $N' = \text{poly}(n)$. Letting $N = 2^k \cdot N'$, we get that $X, Y \in \{0, 1\}^N$ and $\langle X, Y \rangle \geq (2/3)N' = (2/3) \cdot 2^{-k} \cdot N$ if (x, y) is a yes-instance of P and $\langle X, Y \rangle \leq (1/3)N' = (1/3) \cdot 2^{-k} \cdot N$ if (x, y) is a no-instance of P . This gives the desired reduction. \square

Next we give an upper bound on $\text{isr-cc}(\text{GAPINNERPRODUCT})$. In fact we give an upper bound on $\text{isr-cc}^{\text{ow}}(\text{GAPINNERPRODUCT})$.

Lemma 5.6. *For all $0 \leq s < c \leq 1$ and $\rho > 0$, $\text{isr-cc}_\rho^{\text{ow}}(\text{GAPINNERPRODUCT}) = O(1/\rho^2(c-s)^2)$.*

Proof. Let $X \in \{0, 1\}^n$ and $Y \in \{0, 1\}^n$ be the inputs to Alice and Bob. Recall that Alice and Bob want to distinguish the case $\langle X, Y \rangle \geq c \cdot n$ from the case $\langle X, Y \rangle \leq s \cdot n$.

We shall suppose without loss of generality that Alice and Bob have access to a source of ρ -correlated random spherical Gaussian vectors $g, g' \in \mathbb{R}^n$. We can enforce this in the limit by sampling several ρ -correlated random bit vectors $r_i, r'_i \in \{0, 1\}^n$ for $i \in [N]$ and setting $g = \sum_{i=1}^N r_i / \sqrt{N}$ and $g' = \sum_{i=1}^N r'_i / \sqrt{N}$. We leave out the details for this technical calculation (involving an appropriate use of the central limit theorem) here.

Let t be a parameter to be chosen later and let $(g_1, g'_1), (g_2, g'_2), \dots, (g_t, g'_t)$ be t independent ρ -correlated spherical Gaussian vectors chosen from the source as above. By the rotational invariance of the Gaussian distribution, we can assume without loss of generality that $g'_i = \rho g_i + \sqrt{1 - \rho^2} g''_i$, where the g''_i 's are independent spherical Gaussian vectors.

As g_1, \dots, g_t are independent spherical Gaussians, by standard tail bounds (e.g., see Ledoux and Talagrand [11]), with probability at least $1 - 1/6$,

$$\max_{i \in [t]} \langle X, g_i \rangle = (\alpha \sqrt{\log t} \pm O(1)) \cdot \sqrt{\langle X, X \rangle}$$

for some universal constant α .

The protocol then proceeds as follows:

- Alice computes $\ell = \arg \max_{i \in [t]} \langle X, g_i \rangle$ and m such that $\langle X, X \rangle \in ((m-1) \cdot \frac{(c-s)}{100} n, m \cdot \frac{(c-s)}{100} n]$ and sends (ℓ, m) to Bob (note that this implies $m = O(1/(c-s))$).
- Bob accepts if $m \geq \frac{100c}{c-s}$ and $\langle Y, g'_\ell \rangle \geq \alpha \rho \sqrt{\log t} \cdot \frac{(c+s)n}{2\sqrt{m(c-s)(n/100)}}$ and rejects otherwise.

Now, write $Y = aX + bX^\perp$ for some vector X^\perp with $a\langle X, X \rangle = \langle X, Y \rangle$ and $\langle X, X^\perp \rangle = 0$. Then,

$$\langle Y, g'_\ell \rangle = a\rho \langle X, g_\ell \rangle + b\rho \langle X^\perp, g_\ell \rangle + \sqrt{1 - \rho^2} \langle Y, g''_\ell \rangle.$$

As $\langle X, g_\ell \rangle$ is independent of $\langle X^\perp, g_\ell \rangle$ and $\langle Y, g''_\ell \rangle$, it follows from a simple tail bound for univariate Gaussians that with probability at least $1 - 1/6$, $|\langle X^\perp, g_\ell \rangle|, |\langle Y, g''_\ell \rangle| = O(\sqrt{n})$. By combining the above inequalities, we get that with probability at least $2/3$,

$$\langle Y, g'_\ell \rangle = \alpha \rho \sqrt{\log t} \langle X, Y \rangle / \sqrt{\langle X, X \rangle} \pm O(\sqrt{n}).$$

To finish the proof observe that for yes-instances, $\langle X, X \rangle \geq cn$ (so that $m \geq \frac{100c}{c-s}$) and $\langle X, Y \rangle / \sqrt{\langle X, X \rangle} \geq \beta_1 \triangleq c \cdot n / \sqrt{m(c-s)(n/100)}$; while for no-instances, $\langle X, Y \rangle \leq \beta_2 \triangleq s \cdot n / \sqrt{(m-1)(c-s)(n/100)}$. Hence, the protocol works correctly if $\alpha \rho \sqrt{\log t} (\beta_1 - \beta_2) \gg O(\sqrt{n})$.

It follows from the settings of parameters that this indeed happens for some $\log t = \Theta(1/(\rho^2(c-s)^2))$. In particular, we have

$$\beta_1 - \beta_2 = \frac{cn - sn}{\sqrt{m(c-s)(n/100)}} - \frac{sn}{\sqrt{(c-s)(n/100)}} \left(\frac{1}{\sqrt{m-1}} - \frac{1}{\sqrt{m}} \right).$$

By the condition $m \geq \frac{100c}{c-s}$ we have $\frac{1}{\sqrt{m-1}} - \frac{1}{\sqrt{m}} \leq \frac{c-s}{2s}$ and thus

$$\beta_1 - \beta_2 \geq \frac{1}{2} \frac{cn - sn}{\sqrt{m(c-s)(n/100)}} = \frac{\sqrt{25(c-s)n}}{\sqrt{m}}.$$

And so when $\log t \gg 1/\Omega(1/(\alpha^2 \rho^2 (c-s)^2))$ we find $\alpha \rho \sqrt{\log t} (\beta_1 - \beta_2) \gg O(\sqrt{n})$ as required. \square

The above lemma along with the hardness of GAPINNERPRODUCT gives us [Theorem 2.8](#).

Proof of Theorem 2.8. By [Proposition 5.5](#), for every promise problem P such that $\text{psr-cc}(P) \leq k$, P reduces to $\text{GAPINNERPRODUCT}_{c,s}$ with $c = (2/3)2^{-k}$ and $s = (1/3)2^{-k}$. By [Lemma 5.6](#) we get that the reduced instance of $\text{GAPINNERPRODUCT}_{c,s}$ has a one-way isr communication protocol of with $O_\rho(1/(c-s)^2) = O_\rho(2^{2k})$ bits of communication. The theorem follows. \square

5.3 ISR lower bound for SPARSEGAPINNERPRODUCT

In this section, we consider the promise problem $\text{SPARSEGAPINNERPRODUCT}_{.99q, .9q^{-1}, .6q^{-1}}^n$ and show that it has a one-way psr protocol with $O(\log q)$ bits of communication, and then give a two-way isr lower bound of $q^{\Omega(1)}$ for this problem. Together this proves [Theorem 2.9](#).

Proposition 5.7. $\forall c > s$ and $\forall q, n$, we have

$$\text{psr-cc}^{\text{ow}}(\text{SPARSEGAPINNERPRODUCT}_{q,c,s}^n) \leq O\left(\frac{1}{q^2(c-s)^2} \left(\log \frac{1}{c} + \log \frac{1}{q(c-s)} + \log \log \frac{c}{c-s}\right)\right).$$

Proof (Sketch). We first show that there exists an atomic one-way communication protocol for the problem $\text{SPARSEGAPINNERPRODUCT}_{q,c,s}^n$ with the following features (where $\gamma = \Theta((c-s)/c)$):

1. the length of communication is $O(\log 1/c + \log 1/(q(c-s)) + \log \log 1/\gamma)$.
2. yes-instances are accepted with probability at least $(1-\gamma) \cdot \frac{c}{c-s} \cdot p$ and no-instances with probability at most $\frac{s}{c-s} \cdot \frac{100}{m-1}$ for some $m = \Omega(c/(c-s))$ known by both parties. In particular, the difference between completeness and soundness is $\Omega(1/m)$.

The atomic protocol lets the shared randomness determine a sequence of $t \stackrel{\text{def}}{=} -\log(1/\gamma)/\log(1-c)$ indices i_1, i_2, \dots, i_t in $[n]$. Alice first computes $m = O(1/(c-s))$ such that $\|x\|_2^2 \in ((m-1) \cdot \frac{(c-s)}{100}n, m \cdot \frac{(c-s)}{100}n]$, and picks the smallest index ℓ such that $x_{i_\ell} \neq 0$. Then she sends (ℓ, m) to Bob, or $(0, 0)$ if no such index was found. (Note that by sparsity of x , we have $m = O(1/(q(c-s)))$). Bob outputs 0 if he received $(0, 0)$ or if $m < \frac{100c}{c-s}$, and y_{i_ℓ} otherwise.

The completeness follows from the fact that, for yes-instances, $\|x\|_2^2 \geq cn$ (implying $m \geq \frac{100c}{c-s}$) and one expects an index ℓ such that $x_{i_\ell} \neq 0$ among the first roughly $1/c$ choices of ℓ ; conditioned on this, y_{i_ℓ} is 1 with probability at least $\frac{cn}{\|x\|_2^2} \geq \frac{c}{c-s} \frac{100}{m}$. As for the soundness, observe that a no-instance for which Alice does not send 0 to Bob will have $y_{i_\ell} = 1$ with probability at most $\frac{sn}{\|x\|_2^2} < \frac{s}{c-s} \cdot \frac{100}{m-1}$. Now, since $m \geq \frac{100c}{c-s}$, $\frac{100s}{c-s} \left(\frac{1}{m-1} - \frac{1}{m}\right) \leq \frac{100}{3m}$; and by choice of $\gamma \leq c/3(c-s)$ we also have $\frac{c}{c-s} \frac{100}{m} \leq \frac{100}{3m}$. This implies the difference in acceptance probability between completeness and soundness is at least $\frac{100}{3m}$. Repeating this protocol $O(m^2) = O(1/(q^2(c-s)^2))$ times and thresholding yields the final result. \square

We now state our main lower bound theorem.

Theorem 5.8. *There exists $\varepsilon > 0$ such that $\forall q$ and for every sufficiently large n , we have*

$$\text{isr-cc}(\text{SPARSEGAPINNERPRODUCT}_{.99q, .9q^{-1}, .6q^{-1}}^n) \geq q^\varepsilon.$$

(Furthermore, one can take $\varepsilon = 1/2$.)

We prove [Theorem 5.8](#) in [Section 6](#), but we now note that [Theorem 2.9](#) follows immediately from [Proposition 5.7](#) and [Theorem 5.8](#).

Proof of [Theorem 2.9](#). The promise problem is $P = \text{SPARSEGAPINNERPRODUCT}_{.99, 2^k, .9, 2^{-k}, .6, 2^{-k}}$. By [Proposition 5.7](#) we have $\text{psr-cc}^{\text{ow}}(P) \leq O(k)$ and by [Theorem 5.8](#) we have $\text{isr-cc}(P) \geq 2^{\Omega(k)}$. \square

6 Proof of [Theorem 5.8](#)

Our goal for this section is to prove, modulo some technical theorems, that `SPARSEGAPINNERPRODUCT` has high communication complexity in the imperfect shared randomness setting. Before jumping into the proof we give some overview first.

6.1 Proof setup

To prove [Theorem 5.8](#), we will show that for every “strategy” of Alice and Bob, there is a pair of distributions \mathcal{Y} and \mathcal{N} supported (mostly) on yes and no instances, respectively, such that the strategies do not have much “success” in distinguishing them. We note that in contrast to typical lower bounds for perfectly shared randomness, we cannot hope to fix a distribution that works against every strategy. Indeed for every pair of distributions, by virtue of the protocol given in [Proposition 5.7](#) and the Yao min-max principle we have even a deterministic strategy (let alone randomized strategy with imperfect sharing) that succeeds in distinguishing them with high probability. So instead we have to fix the strategies first and then give a pair of distributions that does not work for that strategy. We define the notion of strategy and success more formally below, and then work towards the proof of [Theorem 5.8](#).

Strategy: We now use [Section 5.1](#) to formalize what it would mean to have a k -bit communication protocol for any communication problem. For aesthetic reasons we view Alice and Bob’s strategies as probabilistic ones. Recall, by [Proposition 5.3](#), that k -bit probabilistic communication strategies for Alice can be described by elements of $K_A^{(k)} \in [0, 1]^{2^k}$ and similarly by elements of $K_B^{(k)} \in [0, 1]^{2^k}$ for Bob. So, on randomness r we have that Alice’s communication strategy can be described by a function $f^{(r)}: \{0, 1\}^n \rightarrow K_A^{(k)}$. Similarly for randomness s , Bob’s communication strategy can be described by a function $g^{(s)}: \{0, 1\}^n \rightarrow K_B^{(k)}$.

Thus, a *strategy* for a game is a pair of sets of functions $\mathcal{F} = (f^{(r)})_r, \mathcal{G} = (g^{(s)})_s$, where

$$\begin{aligned} f^{(r)} &: \{0, 1\}^n \rightarrow K_A^{(k)} \\ \text{and } g^{(s)} &: \{0, 1\}^n \rightarrow K_B^{(k)}. \end{aligned}$$

We consider a pair of distributions $D = (\mathcal{Y}, \mathcal{N})$ to be *valid* if \mathcal{Y} is mostly (say with probability .9) supported on yes-instances and \mathcal{N} mostly on no-instances. For valid D , we define

$$\begin{aligned} \text{succ}_D(f, g) &\stackrel{\text{def}}{=} \mathbb{E}_{(x,y) \sim \mathcal{Y}} [\langle f(x), g(y) \rangle] - \mathbb{E}_{(x,y) \sim \mathcal{N}} [\langle f(x), g(y) \rangle] \\ \text{succ}_{D,\rho}(\mathcal{F}, \mathcal{G}) &\stackrel{\text{def}}{=} \left| \mathbb{E}_{(r,s) \sim \rho} \left[\text{succ}_D(f^{(r)}, g^{(s)}) \right] \right| \\ \text{succ}_\rho(\mathcal{F}, \mathcal{G}) &\stackrel{\text{def}}{=} \min_{\text{valid } D} \text{succ}_{D,\rho}(\mathcal{F}, \mathcal{G}) \end{aligned}$$

We note that any strategy that distinguishes yes-instances of SPARSEGAPINNERPRODUCT from no-instances with probability ε must have success $\varepsilon - .1$ on every valid distribution as well (with the difference of $.1$ coming up due to the fact that valid distributions are not entirely supported on the right instances). In what follows we will explain why strategies (with small k) do not have sufficiently positive success.

6.2 Overview of proof of [Theorem 5.8](#).

To prove [Theorem 5.8](#) we need to show that if a pair of strategies $(\mathcal{F}, \mathcal{G})$ achieves $\text{succ}_\rho(\mathcal{F}, \mathcal{G}) > .01$ then k must be large. Roughly our strategy for showing this is as follows: We first define two simple distributions \mathcal{Y} and \mathcal{N} (independent of the strategy $(\mathcal{F}, \mathcal{G})$) and show that any fixed pair of functions (f, g) that are successful in distinguishing \mathcal{Y} from \mathcal{N} must have a few influential variables and furthermore at least one of these variables must be common to both f and g (see [Theorem 6.8](#)). Our proof of this theorem, is based on the “invariance principle” [[12](#)] and [Theorem 6.8](#) is a variant of it which is particular suited for use in communication complexity. The proof of this theorem is deferred to [Section 7](#).

We use this theorem to design agreement distillation strategies for two new players Charlie and Dana as follows: Given shared random pair (r, s) , Charlie picks a random influential variable x_i of the function $f^{(r)}$ used by Alice on random string r and outputs the index $i \in [n]$. Dana similarly picks a random influential variable y_j of the function $g^{(s)}$ used by Bob and outputs j . [Theorem 6.8](#) assures us that with non-trivial probability $i = j$ and this gives an agreement protocol.

If we could argue that $i = j$ has high min-entropy, then we would be done (using [Lemma 4.1](#) which asserts that it is not possible to distill agreement with high-entropy and high probability). But this step is not immediate (and should not be since we have not crafted a distribution specific to $(\mathcal{F}, \mathcal{G})$). To show that this strategy produces indices of high min-entropy, we consider the distribution of indices that is produced by Charlie as we vary r and let BAD_C denote the indices that are produced with too high a probability. Similarly we let BAD_D denote the indices that are produced with too high a probability by Dana. We now consider a new distribution \mathcal{Y}' supported on yes-instances of the SPARSEGAPINNERPRODUCT problem. In \mathcal{Y}' the (x, y) pairs are chosen so that when restricted to coordinates in $\text{BAD}_C \cup \text{BAD}_D$ they look like they come from \mathcal{N} while when restricted to coordinates outside $\text{BAD}_C \cup \text{BAD}_D$ they look like they come from \mathcal{Y} (see [Definition 6.13](#) below for a precise description). Since $\text{BAD}_C \cup \text{BAD}_D$ is small, the distribution \mathcal{Y}' remains supported mostly on yes-instances, but strategies that depend mainly on coordinates from $\text{BAD}_C \cup \text{BAD}_D$ would not have much success in distinguishing \mathcal{Y}' from \mathcal{N}' (which remains the original \mathcal{N}).

We use this intuition to argue formally in [Lemma 6.14](#) that a slightly modified sampling protocol of Charlie and Dana, where they discard i, j from $\text{BAD}_C \cup \text{BAD}_D$, leads to agreement with noticeably high probability on a high-entropy random variable, yielding the desired contradiction.

In the rest of this section we first present the main definitions needed to state [Theorem 6.8](#). We then prove [Theorem 5.8](#) assuming [Theorem 6.8](#). We prove the latter in [Section 7](#), along with the main technical ingredient it relies on, the invariance principle of [Theorem 7.1](#).

6.3 Background on influence of variables

We now turn to defining the notion of influential variables for functions and related background material for functions defined on product probability spaces.

Recall that a finite probability space is given by a pair (Ω, μ) where Ω is a finite set and μ is a probability measure on Ω . We will begin with the natural probabilistic definition of influence of a variable on functions defined on product spaces, and then relate it to a more algebraic definition which is needed for the notion of low-degree influence.

Definition 6.1 (Influence and variance). Let (Ω, μ) be a finite probability space, and let $h: \Omega^n \rightarrow \mathbb{R}$ be a function on product probability space. The *variance* of h , denoted $\text{Var}(h)$, is defined as the variance of the random variable $h(x)$ for $x \in \Omega^n \sim \mu^{\otimes n}$, i.e., $\text{Var}(h) = \mathbb{E}_x[h(x)^2] - (\mathbb{E}_x[h(x)])^2$.

For $i \in [n]$, the *i -th influence* of h is defined as

$$\text{Inf}_i(h) = \mathbb{E}_{x^{(-i)} \sim \mu^{\otimes(n-1)}} \left[\text{Var}_{x_i \sim \mu} [h(x)] \right]$$

where $x^{(-i)}$ denotes all coordinates of x except the i 'th coordinate.

To define the notion of *low-degree* influence, we need to work with a multilinear representation of functions $h: \Omega^n \rightarrow \mathbb{R}$. Let $b = |\Omega|$ and $\mathcal{B} = \{\chi_0, \chi_1, \dots, \chi_{b-1}\}$ be a basis of real-valued functions over Ω . Then, every function $h: \Omega^n \rightarrow \mathbb{R}$ has a unique multilinear expansion of the form

$$h(x) = \sum_{\sigma=(\sigma_1, \dots, \sigma_n) \in \{0, 1, \dots, b-1\}^n} \hat{h}_\sigma \chi_\sigma(x) \quad (1)$$

for some real coefficients \hat{h}_σ , where χ_σ is given by $\chi_\sigma(x) \stackrel{\text{def}}{=} \prod_{i \in [n]} \chi_{\sigma_i}(x_i)$.

When the ensemble \mathcal{B} is a collection of *orthonormal* random variables, namely $\chi_0 = 1$ and $\mathbb{E}_{a \sim \mu}[\chi_{j_1}(a)\chi_{j_2}(a)] = \delta_{j_1 j_2}$, it is easy to check that $\text{Var}(h) = \sum_{\sigma \neq 0} \hat{h}_\sigma^2$ and also that

$$\text{Inf}_i(h) = \sum_{\sigma: \sigma_i \neq 0} \hat{h}_\sigma^2.$$

One can also take the above as the algebraic definition of influence, noting that it is independent of the choice of the orthonormal basis \mathcal{B} and thus well-defined. The degree of a multi-index σ is defined as $|\sigma| = |\{i : \sigma_i \neq 0\}|$, and this leads to the definition of low-degree influence.

Definition 6.2 (Low-degree influence). For a function $h: \Omega^n \rightarrow \mathbb{R}$ with multilinear expansion as in (1) with respect to any orthonormal basis, the *i -th degree d influence* of h is the influence of the truncated multilinear expansion of h at degree d , that is

$$\text{Inf}_i^d(h) \stackrel{\text{def}}{=} \sum_{\substack{\sigma: \sigma_i \neq 0 \\ |\sigma| \leq d}} \hat{h}_\sigma^2.$$

Remark 6.3 (Functions over size 2 domain). When $|\Omega| = 2$, and $\{1, \chi\}$ is an orthonormal basis of real-valued functions over Ω , the expansion (1) becomes the familiar Fourier expansion $h(x) = \sum_{S \subseteq [n]} \hat{h}(S) \prod_{i \in S} \chi(x_i)$, and we have $\text{Inf}_i(h) \stackrel{\text{def}}{=} \sum_{S \ni i} \hat{h}(S)^2$ and $\text{Inf}_i^d(h) \stackrel{\text{def}}{=} \sum_{\substack{S \ni i \\ |S| \leq d}} \hat{h}(S)^2$.

We will make use of the following simple upper bound on the number of low-degree influential coordinates (which follows immediately, for instance, from [12, Proposition 3.8])

Proposition 6.4. *For every $\tau > 0$ and $d \in \mathbb{Z}^+$ there exists $t = t(\tau, d)$ such that for all n and all functions $h: \Omega^n \rightarrow [-1, 1]$, we have $\left| \left\{ i \in [n] : \text{Inf}_i^d(h) > \tau \right\} \right| \leq t$. (Furthermore, one can take $t = d/\tau$.)*

For the invariance principle, we will understand the behavior of a function when its domain is replaced by a different probability space with matching second moments. For this purpose, we will view functions as multilinear polynomials as follows.

Definition 6.5 (Functions on product spaces as multilinear polynomials). The multilinear polynomial associated with a function $h: \Omega^n \rightarrow \mathbb{R}$ with respect to a basis $\mathcal{B} = \{\chi_0, \chi_1, \dots, \chi_{b-1}\}$ of real-valued functions over Ω is a polynomial in indeterminates $\mathbf{z} = \{z_{i,j} : i \in [n], j \in \{0, 1, \dots, b-1\}\}$ given by

$$H(\mathbf{z}) = \sum_{\sigma \in \{0, 1, \dots, m-1\}^n} \hat{h}_\sigma \mathbf{z}_\sigma,$$

\mathbf{z}_σ stands for the monomial $\prod_{i=1}^n z_{i, \sigma_i}$ and the coefficients \hat{h}_σ are given by the multilinear expansion (1) of f w.r.t. \mathcal{B} .

Above, we saw how a function can be viewed as a multilinear polynomial w.r.t. a basis of random variables. Conversely, one can view multilinear polynomials as functions by substituting random variables for its indeterminates.

Definition 6.6 (Multilinear polynomials as random variables on product spaces). Given a collection of random variables $\mathcal{X} = \{\chi_0, \dots, \chi_{m-1}\}$ over a probability space (Ω, μ) , one can view a multilinear polynomial P in indeterminates $\mathbf{z} = \{z_{i,j} : i \in [n], j \in \{0, 1, \dots, m-1\}\}$ given by

$$P(\mathbf{z}) = \sum_{\sigma \in \{0, 1, \dots, m-1\}^n} \hat{P}_\sigma \mathbf{z}_\sigma,$$

where \mathbf{z}_σ stands for the monomial $\prod_{i=1}^n z_{i, \sigma_i}$, as a random variable $P(\mathcal{X}^n)$ over the probability space $(\Omega^n, \mu^{\otimes n})$ mapping $x = (x_1, \dots, x_n)$ to

$$\sum_{\sigma \in \{0, 1, \dots, m-1\}^n} \hat{P}_\sigma \prod_{i=1}^n \chi_{\sigma_i}(x_i). \quad (2)$$

6.4 Proof of Theorem 5.8

We start by introducing a few definitions, in particular of the central distributions and the extraction strategy. We begin with the description of the basic distributions \mathcal{Y} and \mathcal{N} .

Definition 6.7 (\mathcal{Y} , \mathcal{N}). We define two distributions B_N and B_Y on $\{0, 1\} \times \{0, 1\}$ below. The distributions \mathcal{Y} and \mathcal{N} will be product distributions on $(\{0, 1\} \times \{0, 1\})^n$, given by $\mathcal{Y} = B_Y^{\otimes n}$ and $\mathcal{N} = B_N^{\otimes n}$.

- A pair (x, y) is drawn from B_N by setting $x \sim \text{Bern}(1/q)$ and $y \in \{0, 1\}$ uniformly at random. Note that x, y are independent, and $\mathbb{E}[xy] = \frac{1}{2q}$.
- A pair (x, y) is drawn from B_Y by setting

$$(x, y) = \begin{cases} (0, 1) & \text{w.p. } \frac{1}{2} \left(1 - \frac{1.95}{q}\right) \\ (0, 0) & \text{w.p. } \frac{1}{2} \left(1 - \frac{0.05}{q}\right) \\ (1, 1) & \text{w.p. } \frac{1.95}{2q} \\ (1, 0) & \text{w.p. } \frac{0.05}{2q} \end{cases}$$

so that the marginals of x, y in B_Y match those of B_N , and $\mathbb{E}[xy] = \frac{1.95}{2q}$.

A straightforward application of tail inequalities for independent, identically distributed (i.i.d.) random variables tells us that \mathcal{Y} is mostly supported on yes-instances of $\text{SPARSEGAPINNERPRODUCT}_{.99q, 0.9q, 0.6q}^n$ with high probability for sufficiently large n . Similarly \mathcal{N} is mostly supported on no-instances.

Our main technical result is the following theorem showing any fixed pair of vector-valued functions (f, g) (corresponding to strategies for Alice and Bob) that succeed in distinguishing \mathcal{Y} from \mathcal{N} must share an influential variable (with non-trivially high influence of non-trivially low-degree).

Theorem 6.8. *There exist functions $k_0 \geq \Omega_\varepsilon(\sqrt{q})$, $d(q, \varepsilon) < \infty$, and $\tau(q, \varepsilon) > 0$, defined for $q \in \mathbb{Z}^+$, and $\varepsilon > 0$, such that the following holds: For every $\varepsilon > 0$ and $k, q \in \mathbb{Z}^+$ and every sufficiently large n , if $k < k_0(q, \varepsilon)$ and $f: \{0, 1\}^n \rightarrow K_A^{(k)}$ and $g: \{0, 1\}^n \rightarrow K_B^{(k)}$ are functions such that $\text{succ}_{(\mathcal{Y}, \mathcal{N})}(f, g) \geq \varepsilon$, then there exists $i \in [n]$ such that*

$$\min \left\{ \max_{j \in [2^k]} \text{Inf}_i^{d(q, \varepsilon)}(f_j), \max_{j \in [2^k]} \text{Inf}_i^{d(q, \varepsilon)}(g_j) \right\} \geq \tau(q, \varepsilon).$$

(Here, the influence of f_j is w.r.t. to the $\text{Bern}(1/q)$ distribution on $\{0, 1\}$, and that of g_j is w.r.t. the uniform distribution on $\{0, 1\}$.)

This theorem is proved in [Section 7](#). Building on this theorem, we can try to build agreement distillation protocols $(\text{Ext}_C, \text{Ext}_D)$ that exploit the success of the strategies $(\mathcal{F}, \mathcal{G})$ to distill common randomness. We start by first identifying coordinates that may be influential for too many pairs (r, s) (and thus may be produced with too high a probability by a naive distillation protocol).

For the rest of the section we fix $q \in \mathbb{Z}^+$ and $\varepsilon > 0$ and let $d = d(q, \varepsilon)$ and $\tau = \tau(q, \varepsilon)$ where $d(\cdot, \cdot)$ and $\tau(\cdot, \cdot)$ are the functions from [Theorem 6.8](#).

Definition 6.9 ($\text{BAD}_C, \text{BAD}_D$). Let $\delta = 1/(100 \cdot 2^{k_0} t)$ where $t = t(\tau, d)$ as given by [Proposition 6.4](#), and $k_0 = k_0(q, \varepsilon)$ is given by [Theorem 6.8](#). Define

$$\text{BAD}_C \stackrel{\text{def}}{=} \left\{ i \in [n] : \Pr_r \left[\max_{j \in [2^k]} \text{Inf}_i^d(f_j^{(r)}) > \tau \right] > \frac{1}{\delta n} \right\} \quad \text{and}$$

$$\text{BAD}_D \stackrel{\text{def}}{=} \left\{ i \in [n] : \Pr_s \left[\max_{j \in [2^k]} \text{Inf}_i^d(g_j^{(s)}) > \tau \right] > \frac{1}{\delta n} \right\},$$

where r, s denote the randomness available to Alice and Bob, $f_j^{(r)}$ denotes the j 'th component function for Alice's strategy on randomness r , and similarly for $g_j^{(s)}$.

Directly from this definition, we get

Proposition 6.10. $|\text{BAD}_C|, |\text{BAD}_D| \leq 2^k \cdot t \cdot \delta \cdot n \leq n/100$.

Next, we define the extraction distillation protocols for Charlie and Dana:

Definition 6.11 ($(\text{Ext}_C, \text{Ext}_D)$). For $r \in \{0, 1\}^*$, let

$$S_r \stackrel{\text{def}}{=} \left\{ i \in [n] \setminus \text{BAD}_C : \max_{j \in [2^k]} \text{Inf}_i^d(f_j^{(r)}) > \tau \right\} \quad \text{and} \quad T_s \stackrel{\text{def}}{=} \left\{ i \in [n] \setminus \text{BAD}_D : \max_{j \in [2^k]} \text{Inf}_i^d(g_j^{(s)}) > \tau \right\}.$$

Then, $\text{Ext}_C(r)$ is defined as follows:

if $S_r = \emptyset$ output $i \sim \mathcal{U}_{[n]}$; otherwise output $i \sim \mathcal{U}_{S_r}$.

$\text{Ext}_D(s)$ is defined similarly:

if $T_s = \emptyset$ output $j \sim \mathcal{U}_{[n]}$; otherwise output $j \sim \mathcal{U}_{T_s}$.

Proposition 6.12. $H_\infty(\text{Ext}_C(r)) \geq \log n - \log(1 + 1/\delta)$.

Proof. Fix $i \in [n] \setminus (\text{BAD}_C \cup \text{BAD}_D)$. We have

$$\Pr[i \text{ is output}] \leq \Pr[i \in S_r \text{ and } i \text{ is output}] + \Pr[i \text{ is output} \mid S_r = \emptyset] \leq 1/(\delta n) + 1/n.$$

The proposition follows. \square

Finally we turn to proving that Ext_C and Ext_D do agree with non-trivial probability. To do so we need to consider a new distribution on yes-instances, defined next:

Definition 6.13 (\mathcal{Y}'). The distribution \mathcal{Y}' is a product distribution on $(\{0, 1\} \times \{0, 1\})^n$, where $(x_i, y_i) \sim B_N$ if $i \in \text{BAD}_C \cup \text{BAD}_D$ and $(x_i, y_i) \sim B_Y$ otherwise.

Using **Proposition 6.10** above we have that $\mathbb{E}_{i,x,y}[x_i y_i] \geq .93/q$ and so by standard tail inequalities we still have that \mathcal{Y}' is mostly supported on yes-instances. Our main lemma for this section is that if $(\mathcal{F}, \mathcal{G})$ are successful in distinguishing \mathcal{Y}' and \mathcal{N} and k is small, then Ext_C and Ext_D are likely to agree with noticeable probability (which would contradict **Lemma 4.1**).

Lemma 6.14. Let $k_0 = k_0(q, \varepsilon)$, $d = d(q, \varepsilon)$ and $\tau = \tau(q, \varepsilon)$ be as given in **Theorem 6.8**, and let $t = t(\tau, d)$ as given by **Proposition 6.10**. If $\text{succ}_{(\mathcal{Y}', \mathcal{N}), \rho}(\mathcal{F}, \mathcal{G}) \geq 2\varepsilon$, and $k < k_0$ then

$$\Pr_{(r,s) \sim \rho} [\text{Ext}_C(r) = \text{Ext}_D(s)] \geq \varepsilon / (2t^2).$$

Proof. Expanding the definition of $\text{succ}(\cdot, \cdot)$, we have

$$\mathbb{E}_{(r,s)} \left[\mathbb{E}_{(x,y) \sim \mathcal{Y}'} \left[\langle f^{(r)}(x), g^{(s)}(y) \rangle \right] \right] - \mathbb{E}_{(x,y) \sim \mathcal{N}} [\langle f_r(x), g_s(y) \rangle] \geq 2\varepsilon.$$

Say that a pair (r, s) is GOOD if

$$\mathbb{E}_{(x,y) \sim \mathcal{Y}'} \left[\langle f^{(r)}(x), g^{(s)}(y) \rangle \right] - \mathbb{E}_{(x,y) \sim \mathcal{N}} \left[\langle f^{(r)}(x), g^{(s)}(y) \rangle \right] \geq \varepsilon.$$

By a Markov argument we thus have

$$\Pr_{(r,s)} [(r,s) \text{ is GOOD}] \geq \varepsilon.$$

For any fixed GOOD (r, s) we now prove that there exists $i \in (S_r \cap T_s) \setminus (\text{BAD}_C \cup \text{BAD}_D)$. Note that once we have such an i , we have that $\Pr[\text{Ext}_C(r) = \text{Ext}_C(r') = i]$ with probability at least $1/t(\tau, d)^2$. Combining this with the probability that (r, s) is good, we have $\Pr_{(r,s)}[\text{Ext}_C(r) = \text{Ext}_D(s)] \geq \varepsilon/t(\tau, d)^2$ which yields the lemma. So we turn to this claim.

To simplify notation, assume without loss of generality that $\text{BAD}_C \cup \text{BAD}_D = \{m+1, \dots, n\}$. Define functions $f_1: \{0, 1\}^m \rightarrow K_A^{(k)}$ and $g_1: \{0, 1\}^m \rightarrow K_B^{(k)}$ by letting

$$f_1(x) = \mathbb{E}_{z \sim \text{Bern}^{n-m}(1/q)} [f^{(r)}(x \cdot z)] \quad \text{and} \quad g_1(y) = \mathbb{E}_{w \sim \mathcal{U}(\{0,1\}^{n-m})} [g^{(s)}(y \cdot w)].$$

Note that the success of (f_r, g_s) in distinguishing \mathcal{Y}' from \mathcal{N} turns into the success of (f_1, g_1) in distinguishing \mathcal{Y}_m from \mathcal{N}_m (where $\mathcal{Y}_m = B_Y^{\otimes m}$ and $\mathcal{N}_m = B_N^{\otimes m}$) — this is immediate since $(x \cdot z, y \cdot w) \sim \mathcal{Y}'$ if $(x, y) \sim \mathcal{Y}_m$ and $(x \cdot z, y \cdot w) \sim \mathcal{N}$ if $(x, y) \sim \mathcal{N}_m$.

So we have $\text{succ}_{(\mathcal{Y}_m, \mathcal{N}_m)}(f_1, g_1) \geq \varepsilon$. Since $k < k_0$ we have that there must exist a variable $i \in [m]$ and indices $j, j' \in [2^k]$ with $\text{Inf}_i^d(f_{1,j}) > \tau$ and $\text{Inf}_i^d(g_{1,j'}) > \tau$. (Here $f_{1,j}$ is the j 'th component function of f_1 , and similarly for $g_{1,j'}$.) But $\text{Inf}_i^d(f_j^{(r)}) \geq \text{Inf}_i^d(f_{1,j})$ and $\text{Inf}_i^d(g_j^{(r')}) \geq \text{Inf}_i^d(g_{1,j'})$. To see this, note that $\widehat{f_{1,j}}(S) = \widehat{f_j^{(r)}}(S)$ for $S \subseteq [m]$ and so

$$\begin{aligned} \text{Inf}_i^d(f_j^{(r)}) &= \sum_{i \in S \subseteq [n], |S| \leq d} \widehat{f_j^{(r)}}(S)^2 \\ &\geq \sum_{i \in S \subseteq [m], |S| \leq d} \widehat{f_j^{(r)}}(S)^2 \\ &= \sum_{i \in S \subseteq [m], |S| \leq d} \widehat{f_{1,j}}(S)^2 \\ &= \text{Inf}_i^d(f_{1,j}). \end{aligned}$$

We thus conclude that $i \in S_r \cap T_s \cap [m]$ and this concludes the claim, and thus the lemma. \square

Proof of Theorem 5.8. The proof follows easily from Lemma 4.1 and Lemma 6.14. Assume for contradiction that there is a protocol for SPARSEGAPINNERPRODUCT $_{.99q, .9q, .6q}^n$ with communication complexity less than $k_0(.05, q) = \Omega(\sqrt{q})$ that on access to $r \sim_{\rho} s$ accepts yes-instances with probability at least $2/3$ and no-instances with probability at most $1/3$. This implies that there exist strategies $(\mathcal{F} = \{f^{(r)}\}_r, \mathcal{G} = \{g^{(s)}\}_s)$ such that for every pair of distributions $(\mathcal{Y}, \mathcal{N})$ supported mostly (i.e., with probability .9) on yes and no instances respectively, we have $\text{succ}_{(\mathcal{Y}, \mathcal{N}), \rho}(\mathcal{F}, \mathcal{G}) > .1$. In particular, this holds for the distribution \mathcal{Y}' as defined in Definition 6.13 and \mathcal{N} as defined in Definition 6.7.

Let $\text{Ext}_C, \text{Ext}_D$ be strategies for AGREEMENT-DISTILLATION as defined in Definition 6.11. By Proposition 6.12 we get that $H_{\infty}(\text{Ext}_C(r)), H_{\infty}(\text{Ext}_D(s)) \geq \log n - O(1)$. By Lemma 6.14 we also have $\Pr_{r \sim_{\rho} s}[\text{Ext}_C(r) = \text{Ext}_D(s)] \geq \Omega_q(1)$. But this contradicts Lemma 4.1 which asserts (in particular) that protocols extracting $\omega_n(1)$ bits can agree with probability $o_n(1)$. \square

7 Low-influence communication strategies

The following theorem states that the expected inner product between two multidimensional Boolean functions without common low-degree influential variables when applied to correlated random strings, is well approximated by the expected inner product of two related functions, this time applied to similarly correlated Gaussians. As, per Section 5.1, the former quantity captures the behavior of communication protocols, this invariance principle enables one to transfer the study to the (more manageable) Gaussian setting. (For convenience, in this section we switch to the equivalent view of Boolean functions as being defined on $\{+1, -1\}^n$).

We denote by $N_{p_1, p_2, \theta}$ the distribution on $\{+1, -1\} \times \{+1, -1\}$ such that the marginals of $(x, y) \sim N_{p_1, p_2, \theta}$ have expectations respectively p_1 and p_2 , and correlation θ (see Definition A.1 for an explicit definition).

Theorem 7.1. Fix any two parameters $p_1, p_2 \in (-1, 1)$. For all $\varepsilon \in (0, 1]$, $\ell \in \mathbb{Z}^+$, $\theta_0 \in [0, 1)$ and closed convex sets $K_1, K_2 \subseteq [0, 1]^\ell$, there exist $n_0 \in \mathbb{Z}^+$, $d \in \mathbb{Z}^+$ and $\tau \in (0, 1)$ such that the following holds. For all $n \geq n_0$, there exist mappings

$$\begin{aligned} T_1 &: \{f: \{+1, -1\}^n \rightarrow K_1\} \rightarrow \{F: \mathbb{R}^n \rightarrow K_1\} \\ T_2 &: \{g: \{+1, -1\}^n \rightarrow K_2\} \rightarrow \{G: \mathbb{R}^n \rightarrow K_2\} \end{aligned}$$

such that for all $\theta \in [-\theta_0, \theta_0]$, if f, g satisfy

$$\max_{i \in [n]} \min \left(\max_{j \in [\ell]} \text{Inf}_i^d(f_j), \max_{j \in [\ell]} \text{Inf}_i^d(g_j) \right) \leq \tau \quad (3)$$

then, for $F = T_1(f)$ and $G = T_2(g)$, we have

$$\left| \mathbb{E}_{(x,y) \sim N^{\otimes n}}[\langle f(x), g(y) \rangle] - \mathbb{E}_{(X,Y) \sim \mathcal{G}^{\otimes n}}[\langle F(X), G(Y) \rangle] \right| \leq \varepsilon. \quad (4)$$

where $N = N_{p_1, p_2, \theta}$ and \mathcal{G} is the Gaussian distribution which matches the first and second-order moments of N , i.e. $\mathbb{E}[x_i] = \mathbb{E}[X_i]$, $\mathbb{E}[x_i^2] = \mathbb{E}[X_i^2]$ and $\mathbb{E}[x_i y_i] = \mathbb{E}[X_i Y_i]$.

The theorem follows in a straightforward manner from [Lemma 7.2](#) and [Theorem 7.3](#):

Proof of Theorem 7.1. For $\varepsilon \in (0, 1]$, $\ell \in \mathbb{Z}^+$ and $\theta_0 \in (0, 1)$ as above, let $\tau_1 \stackrel{\text{def}}{=} \tau(\varepsilon/2, \ell, \theta_0)$ as in [Theorem 7.3](#). Define the operators T_1, T_2 as

$$T_1 = T_1^{(2)} \circ T_1^{(1)}, \quad T_2 = T_2^{(2)} \circ T_2^{(1)}$$

where $T_1^{(1)}, T_1^{(2)}$ are the operators from [Lemma 7.2](#) (for $\varepsilon/2, \ell, \theta_0$ and τ_1 as above, which yield the values of τ, d and n_0) and $T_1^{(2)}, T_2^{(2)}$ are the (non-linear) ones from [Theorem 7.3](#) (with parameters ℓ, θ_0 and $\varepsilon/2$). The result follows. \square

The first step towards proving the theorem is to convert the expected inner product of Boolean functions with no shared low-degree influential variables into expected inner product of Boolean functions with no influential variables at all.

Lemma 7.2. Fix any two parameters $p_1, p_2 \in (-1, 1)$. For all $\varepsilon \in (0, 1]$, $\ell \in \mathbb{Z}^+$, $\tau \in (0, 1)$, $\theta_0 \in [0, 1)$ and convex sets $K_1, K_2 \subseteq [0, 1]^\ell$, there exist $n_0 \in \mathbb{Z}^+$, $d \in \mathbb{Z}^+$ and $\tau' \in (0, 1)$ such that the following holds. For all $n \geq n_0$ there exist operators

$$\begin{aligned} T_1^{(1)} &: \{f: \{+1, -1\}^n \rightarrow K_1\} \rightarrow \{\tilde{f}: \{+1, -1\}^n \rightarrow K_1\} \\ T_2^{(1)} &: \{g: \{+1, -1\}^n \rightarrow K_2\} \rightarrow \{\tilde{g}: \{+1, -1\}^n \rightarrow K_2\} \end{aligned}$$

such that for all $\theta \in [-\theta_0, \theta_0]$, if f, g satisfy

$$\max_{i \in [n]} \min \left(\max_{j \in [\ell]} \text{Inf}_i^d(f_j), \max_{j \in [\ell]} \text{Inf}_i^d(g_j) \right) \leq \tau' \quad (5)$$

then, for $\tilde{f} = T_1^{(1)} f$ and $\tilde{g} = T_2^{(1)} g$,

$$\max_{i \in [n]} \max \left(\max_{j \in [\ell]} \text{Inf}_i(\tilde{f}_j), \max_{j \in [\ell]} \text{Inf}_i(\tilde{g}_j) \right) \leq \tau \quad (6)$$

and

$$\left| \mathbb{E}_{(x,y) \sim N^{\otimes n}} \langle f(x), g(y) \rangle - \mathbb{E}_{(x,y) \sim N^{\otimes n}} \langle \tilde{f}(x), \tilde{g}(y) \rangle \right| \leq \varepsilon. \quad (7)$$

where $N = N_{p_1, p_2, \theta}$.

Proof. The proof uses Lemmas 6.1 and 6.7 in [12] applied to each pair of functions (f_i, g_i) , for $i \in [\ell]$ applied with parameter θ_0 and ε/ℓ ; using when applying the first lemma the fact that the correlation of these $N_{p_1, p_2, \theta}$ is bounded away from 1. The operators given in Lemmas 6.1 and 6.7 in [12] are simple averaging operators (averaging the value of f over some neighborhood of x to get its new value at x) and by the convexity of K_1 we have that the averaged value remains in K_1 . Similarly for g and K_2 . We omit the details. \square

The last ingredient needed is the actual invariance principle, which will take us from the Boolean, low-influence setting to the Gaussian one.

Theorem 7.3. *Fix any two parameters $p_1, p_2 \in (-1, 1)$. For all $\varepsilon \in (0, 1]$, $\ell \in \mathbb{Z}^+$, $\theta_0 \in [0, 1)$, and closed convex sets $K_1, K_2 \subseteq [0, 1]^\ell$ there exist $\tau > 0$ and mappings*

$$\begin{aligned} T_1^{(2)} &: \{f: \{+1, -1\}^n \rightarrow K_1\} \rightarrow \{F: \mathbb{R}^n \rightarrow K_1\} \\ T_2^{(2)} &: \{g: \{+1, -1\}^n \rightarrow K_2\} \rightarrow \{G: \mathbb{R}^n \rightarrow K_2\} \end{aligned}$$

such that for all $\theta \in [-\theta_0, \theta_0]$, if $f: \{+1, -1\}^n \rightarrow K_1$ and $g: \{+1, -1\}^n \rightarrow K_2$ satisfy

$$\max_{i \in [n]} \max \left(\max_{j \in [\ell]} \text{Inf}_i(f_j), \max_{j \in [\ell]} \text{Inf}_i(g_j) \right) \leq \tau$$

then for $F = T_1^{(2)}(f): \mathbb{R}^n \rightarrow K_1$ and $G = T_2^{(2)}(g): \mathbb{R}^n \rightarrow K_2$

$$\left| \mathbb{E}_{(x,y) \sim N^{\otimes n}} [\langle f(x), g(y) \rangle] - \mathbb{E}_{(X,Y) \sim \mathcal{G}^{\otimes n}} [\langle F(X), G(Y) \rangle] \right| \leq \varepsilon,$$

where $N = N_{p_1, p_2, \theta}$ and \mathcal{G} is the Gaussian distribution which matches the first and second-order moments of N .

Proof. Deferred to [Appendix A](#). \square

7.1 Lower bound for Gaussian Inner Product

We now deduce a lower bound on k , the communication complexity of the strategies captured by the range of f and g , needed to achieve sizeable advantage in distinguishing between ξ -correlated and uncorrelated Gaussian inputs. Hereafter, \mathcal{G}_ρ denotes the bivariate normal Gaussian distribution with correlation ρ .

Lemma 7.4. *Let $\xi \in (0, 1/2)$, $\gamma > 0$. There exists a function $k_1(\xi, \gamma) \geq \Omega_\gamma(1/\xi)$ such that for every n the following holds: if there are functions $F: \mathbb{R}^n \rightarrow K_A^{(k)}$ and $G: \mathbb{R}^n \rightarrow K_B^{(k)}$ such that*

$$\left| \mathbb{E}_{(x,y) \sim \mathcal{G}_\xi^{\otimes n}} [\langle F(x), G(y) \rangle] - \mathbb{E}_{(x,y) \sim \mathcal{G}_0^{\otimes n}} [\langle F(x), G(y) \rangle] \right| \geq \gamma,$$

then $k \geq k_1(\xi, \gamma)$.

We will prove the above theorem by translating the above question to a communication lower bound question.

GAUSSIANCORRELATION $_{\xi,n}$: In this (promise) communication game, Alice holds $x \in \mathbb{R}^n$ and Bob holds $y \in \mathbb{R}^n$ from one of two distributions:

- μ_{yes} : each (x_i, y_i) is an independent pair of ξ -correlated standard normal variables.
- μ_{no} : each (x_i, y_i) is an independent copy of uncorrelated standard normal variables.

The goal is for Alice and Bob to communicate with each other, with shared randomness, and distinguish between the two cases with good advantage.

Note that if (X, Y) denotes the random variable each pair (x_i, y_i) , estimating $\mathbb{E}[XY]$ within accuracy $< \xi/2$ suffices to solve the above problem. If Alice sends the values of x_i (suitably discretized) for the first $O(1/\xi^2)$ choices of i , then by standard Chebyshev tail bounds Bob can estimate $\mathbb{E}[XY]$ to the desired accuracy, and so this problem can be solved with $O(1/\xi^2)$ bits of (one-way) communication. We now show that $\Omega(1/\xi)$ is a lower bound.

Lemma 7.5. *Let $\xi \in (0, 1/2)$ and n be sufficiently large. Suppose there is a k -bit communication protocol for **GAUSSIANCORRELATION** (ξ, n) that distinguishes between μ_{yes} and μ_{no} with advantage $\gamma > 0$. Then $k \geq \Omega_{\gamma}(1/\xi)$.*

Before we prove the result, note that **Lemma 7.4** follows immediately with $k_1(\xi, \gamma) = \Omega_{\gamma}(1/\xi)$, since by **Proposition 5.3** the functions $F: \mathbb{R}^n \rightarrow K_A^{(k)}$ and $G: \mathbb{R}^n \rightarrow K_B^{(k)}$ simply correspond to strategies for an k -bit two-way communication protocol with acceptance probability given by $\mathbb{E}_{X,Y}[\langle F(X), G(Y) \rangle]$.

Proof of Lemma 7.5. The lower bound is proved by reducing the **DISJOINTNESS** problem (in particular a promise version of it) to the **GAUSSIANCORRELATION** problem.

Specifically we consider the promise **DISJOINTNESS** problem with parameter m , where Alice gets a vector $u \in \{0, 1\}^m$ and Bob gets $v \in \{0, 1\}^m$. The yes-instances satisfy $\langle u, v \rangle = 1$ while the no-instances satisfy $\langle u, v \rangle = 0$, where the inner product is over the reals. Kalyanasundaram and Schnitger [9] show that distinguishing yes-instances from no-instances requires $\Omega(m)$ bits of communication, even with shared randomness.

We reduce **DISJOINTNESS $_m$** to **GAUSSIANCORRELATION** with $\xi = 1/m$ as follows: Alice and Bob share mn independent standard Gaussians $\{G_{ij} : i \in [n], j \in [m]\}$. Alice generates $x = (x_1, \dots, x_n)$ by letting $x_i = \frac{1}{\sqrt{m}} \sum_{j=1}^m u_j \cdot G_{ij}$ and Bob generates $y = (y_1, \dots, y_n)$ by letting $y_i = \frac{1}{\sqrt{m}} \sum_{j=1}^m v_j \cdot G_{ij}$. It can be verified that x_i and y_i are standard Gaussians with $\mathbb{E}[x_i y_i] = \frac{1}{m} \langle u, v \rangle$. Thus yes-instances of **DISJOINTNESS** map to yes-instances of **GAUSSIANCORRELATION** drawn according to μ_{yes} with $\xi = 1/m$, and no-instances map to no-instances drawn according to μ_{no} . The communication lower bound of $\Omega(m)$ for **DISJOINTNESS** thus translates to a lower bound of $\Omega(1/\xi)$ for **GAUSSIANCORRELATION**. \square

7.2 Putting things together and proof of **Theorem 6.8**

We now combine the results from the previous two sections to prove **Theorem 6.8**.

Proof of Theorem 6.8. Postponing the precise setting of parameters for now, the main idea behind the proof is the following. Suppose the conclusion of the theorem does not hold and f, g do not have a common influential variable so that

$$\max_{i \in [n]} \min \left\{ \text{Inf}_i^d(f), \text{Inf}_i^d(g) \right\} \leq \tau \quad (8)$$

for parameters d, τ that can be picked with an arbitrary dependence on q, ε .

We now associate the domains of f and g with $\{+1, -1\}^n$ in the natural way by mapping $x \in \{0, 1\} \rightarrow 2x - 1 \in \{+1, -1\}$. This defines us functions $f': \{+1, -1\}^n \rightarrow K_A^{(k)}$ and $g': \{+1, -1\}^n \rightarrow K_B^{(k)}$ which satisfy the same conditions on influence as f . Further, under this mapping, the distribution B_Y is mapped to $N_{\mathcal{Y}} \equiv N_{2/q-1, 0, 1, 9/q}$ and B_N is mapped to $N_{\mathcal{N}} \equiv N_{2/q-1, 0, 0}$ (for $N_{p_1, p_2, \theta}$ as defined in Theorem 7.1). Let $\mathcal{G}_{\mathcal{Y}}$ and $\mathcal{G}_{\mathcal{N}}$ denote bivariate Gaussian distributions whose first two moments match those of $N_{\mathcal{Y}}$ and $N_{\mathcal{N}}$ respectively.

Since the ranges of f', g' are closed and convex (from Proposition 5.3) we get, by applying Theorem 7.1 to functions f', g' and distributions $N_{\mathcal{Y}}, \mathcal{G}_{\mathcal{Y}}$ and $N_{\mathcal{N}}, \mathcal{G}_{\mathcal{N}}$ respectively, that there exist functions $F: \mathbb{R}^n \rightarrow K_A^{(k)}$ and $G: \mathbb{R}^n \rightarrow K_B^{(k)}$ such that

$$\begin{aligned} \left| \mathbb{E}_{(x,y) \sim N_{\mathcal{Y}}^{\otimes n}}[\langle f'(x), g'(y) \rangle] - \mathbb{E}_{(X,Y) \sim \mathcal{G}_{\mathcal{Y}}^{\otimes n}}[\langle F(X), G(Y) \rangle] \right| &\leq \frac{\varepsilon}{3} \\ \left| \mathbb{E}_{(x,y) \sim N_{\mathcal{N}}^{\otimes n}}[\langle f'(x), g'(y) \rangle] - \mathbb{E}_{(X,Y) \sim \mathcal{G}_{\mathcal{N}}^{\otimes n}}[\langle F(X), G(Y) \rangle] \right| &\leq \frac{\varepsilon}{3}. \end{aligned} \quad (9)$$

Combining the above equations with the hypothesis that $\text{succ}_{(\mathcal{Y}, \mathcal{N})}(f, g) \geq \varepsilon$, we get

$$\left| \mathbb{E}_{(X,Y) \sim \mathcal{G}_{\mathcal{Y}}^{\otimes n}}[\langle F(X), G(Y) \rangle] - \mathbb{E}_{(X,Y) \sim \mathcal{G}_{\mathcal{N}}^{\otimes n}}[\langle F(X), G(Y) \rangle] \right| \geq \frac{\varepsilon}{3}.$$

To finish the argument, we shall appeal to Lemma 7.4. Let $p = 1/q$ and $\theta = .95p/\sqrt{p-p^2} = \Theta(1/\sqrt{q})$. Let $\phi: \mathbb{R} \rightarrow \mathbb{R}$ be defined by $\phi(z) = 2\sqrt{p-p^2} \cdot z + (2p-1)$. It is easy to check that for $(z, w) \sim \mathcal{G}_{\theta}$, $(\phi(z), w) \sim \mathcal{G}_{\mathcal{Y}}$ and for $(z, w) \sim \mathcal{G}_0$, $(\phi(z), w) \sim \mathcal{G}_{\mathcal{N}}$. Therefore, if we define $f': \mathbb{R}^n \rightarrow \Delta([\ell])$ by $f'(X) = F(\phi(X_1), \dots, \phi(X_n))$, then the above equation is equivalent to

$$\left| \mathbb{E}_{(X,Y) \sim \mathcal{G}_{\theta}^{\otimes n}}[\langle f'(X), G(Y) \rangle] - \mathbb{E}_{(X,Y) \sim \mathcal{G}_0^{\otimes n}}[\langle f'(X), G(Y) \rangle] \right| \geq \frac{\varepsilon}{3}.$$

We can now conclude from Lemma 7.4 that $k \geq \Omega_{\varepsilon}(1/\theta) = \Omega_{\varepsilon}(\sqrt{q})$. To complete the proof of theorem by a contradiction we set the parameters as follows: choose d, τ in Equation 8 so as to deduce Equation 9 from Theorem 7.1 (with $\varepsilon/3$ playing role of ε) and set $k_0 = k_1(\theta, \varepsilon/3)$ for k_1 as given by Lemma 7.4. □

8 Conclusions

In this paper we carried out an investigation of the power of imperfectly shared randomness in the context of communication complexity. There are two important aspects to the perspective that motivated our work: First, the notion that in many forms of natural communication, the communicating parties understand each other (or “know” things about each other) fairly well, but never perfectly. This

imperfection in knowledge/understanding creates an obstacle to many of the known solutions and new solutions have to be devised, or new techniques need to be developed to understand whether the obstacles are barriers. Indeed for the positive results described in this paper, classical solutions do not work and the solutions that ended up working are even “provably” different from classical solutions. (In particular they work hard to preserve “low influence”).

However, we also wish to stress a second aspect that makes the problems here interesting in our view, which is an aspect of scale. Often in communication complexity our main motivation is to compute functions with sublinear communication, or prove linear lower bounds. Our work, and natural communication in general, stresses the setting where inputs are enormous, and the communication complexity one is considering is tiny. This models many aspects of natural communication where there is a huge context to any conversation which is implicit. If this context were known exactly to sender and receiver, then it would play no significant mathematical role. However in natural communication this context is not exactly known, and resolving this imperfection of knowledge before communicating the relevant message would be impossibly hard. Such a setting naturally motivates the need to study problems of input length n , but where any dependence on n in the communication complexity would be impracticable.

We note that we are not at the end of the road regarding questions of this form: Indeed a natural extension to communication complexity might be where Alice wishes to compute $f_A(x,y)$ and Bob wishes to compute $f_B(x,y)$ but Alice does not know f_B and Bob does not know f_A (or have only approximate knowledge of these functions). If x and y are n -bits strings, f_A and f_B might require 2^n bits to describe and this might be the real input size. There is still a trivial upper bound of $2n$ bits for solving any such communication problem, but it would be interesting to study when, and what form of, approximate knowledge of f_A and f_B helps improve over this trivial bound.

Turning to the specific questions studied in this paper a fair number of natural questions arise that we have not been able to address in this work. For instance, we stuck to a specific and simple form of correlation in the randomness shared by Alice and Bob. One could ask what general forms of randomness (r, r') are equally powerful. In particular if the distribution of (r, r') is known to both Alice and Bob, can they convert their randomness to some form of correlation in the sense used in this paper (in product form with marginals being uniform)?

In [Section 4](#) we considered the AGREEMENT-DISTILLATION problem where the goal was for Alice and Bob to agree perfectly on some random string. What if their goal is only to generate more correlated bits than they start with? What is possible here and what are the limits?

In the study of perfectly shared randomness, Newman’s Theorem [13] is a simple but powerful tool, showing that $O(\log n)$ bits of randomness suffice to deal with problems on n bit inputs. When randomness is shared imperfectly, such a randomness reduction is not obvious. Indeed for the problem of equality testing, the protocol of [1] uses 2^n bits of randomness, and our Gaussian protocol (which can solve this with one-way communication) uses $\text{poly}(n)$ bits. Do $O(\log n)$ bits of imperfectly shared randomness suffice for this problem? How about for general problems?

Finally almost all protocols we give for imperfectly shared randomness lead to two-sided error. This appears to be an inherent limitation (with some philosophical implications) but we do not have a proof. It would be nice to show that one-sided error with imperfectly shared randomness cannot lead to any benefits beyond that offered by private randomness.

Acknowledgments

We thank Brendan Juba for his helpful notes [7] on the invariance principle.

References

- [1] Mohammad Bavarian, Dmitry Gavinsky, and Tsuyoshi Ito. On the role of shared randomness in simultaneous communication. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *ICALP (1)*, volume 8572 of *Lecture Notes in Computer Science*, pages 150–162. Springer, 2014. ISBN 978-3-662-43947-0. (document), 1, 8
- [2] Andrej Bogdanov and Elchanan Mossel. On extracting common random bits from correlated sources. *CoRR*, abs/1007.2315, 2010. URL <http://arxiv.org/abs/1007.2315>. 1, 2.2.2, 2.2.2, 4, 4.1
- [3] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley Publishing, New York, 1991. 1
- [4] Venkatesan Guruswami, Johan Håstad, Rajsekar Manokaran, Prasad Raghavendra, and Moses Charikar. Beating the random ordering is hard: Every ordering CSP is approximation resistant. *SIAM J. Comput.*, 40(3):878–914, 2011. A
- [5] Elad Haramaty and Madhu Sudan. Deterministic compression with uncertain priors. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, ITCS '14, pages 377–386, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2698-8. doi: 10.1145/2554797.2554832. URL <http://doi.acm.org/10.1145/2554797.2554832>. 2.2.1
- [6] Marcus Isaksson and Elchanan Mossel. Maximally stable Gaussian partitions with discrete applications. *Israel Journal of Mathematics*, 189:347–396, June 2012. A
- [7] Brendan Juba. 18.177 course project: Invariance principles, 2009. URL <http://people.seas.harvard.edu/~bjuba/papers/18177-report.pdf>. 8
- [8] Brendan Juba, Adam Tauman Kalai, Sanjeev Khanna, and Madhu Sudan. Compression without a common prior: an information-theoretic justification for ambiguity in language. In Bernard Chazelle, editor, *ICS*, pages 79–86. Tsinghua University Press, 2011. ISBN 978-7-302-24517-9. 1, 2.2.1, 2.2.1
- [9] Bala Kalyanasundaram and Georg Schnitger. The probabilistic communication complexity of set intersection. *SIAM J. Discrete Math.*, 5(4):545–557, 1992. doi: 10.1137/0405044. URL <http://dx.doi.org/10.1137/0405044>. 7.1
- [10] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, New York, NY, USA, 2006. ISBN 9780521029834. URL <http://books.google.com/books?id=dHH7rdhKwzsC>. 1
- [11] Michel Ledoux and Michel Talagrand. *Probability in Banach spaces: Isoperimetry and processes*. Springer, Berlin, 1991. ISBN 3540520139. 5.2

- [12] Elchanan Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010. ISSN 1016-443X. doi: 10.1007/s00039-010-0047-x. URL <http://dx.doi.org/10.1007/s00039-010-0047-x>. 6.2, 6.3, 7, A
- [13] Ilan Newman. Private vs. common random bits in communication complexity. *Inf. Process. Lett.*, 39(2):67–71, July 1991. ISSN 0020-0190. doi: 10.1016/0020-0190(91)90157-D. URL [http://dx.doi.org/10.1016/0020-0190\(91\)90157-D](http://dx.doi.org/10.1016/0020-0190(91)90157-D). 1, 2.1, 8

A Proofs from Section 7

Our goal in this section is to prove the needed invariance principle, as stated in [Theorem 7.3](#), that allows us to pass from a correlated distribution on $\{+1, -1\}^2$ to a two-dimensional Gaussian distribution with matching moments. We first formally define the discrete distribution of interest to us.

Definition A.1. For parameters $p_1, p_2, \theta \in [-1, 1]$, let the distribution $N_{p_1, p_2, \theta}$ on $\{+1, -1\} \times \{+1, -1\}$ be defined as follows:³

$$(x, y) = \begin{cases} (+1, +1) & \text{with probability } \frac{1+\theta}{4} + \frac{p_1+p_2}{4} \\ (+1, -1) & \text{with probability } \frac{1-\theta}{4} + \frac{p_1-p_2}{4} \\ (-1, +1) & \text{with probability } \frac{1-\theta}{4} - \frac{p_1-p_2}{4} \\ (-1, -1) & \text{with probability } \frac{1+\theta}{4} - \frac{p_1+p_2}{4} \end{cases}$$

so that $\mathbb{E}[x] = p_1$, $\mathbb{E}[y] = p_2$ and $\mathbb{E}[xy] = \theta$.

The proof of [Theorem 7.3](#) relies on two general ingredients. The first is that replacing f and g by their *smoothed* versions $T_{1-\eta}f$ and $T_{1-\eta}g$ (obtained by applying the Bonami–Beckner noise operator, defined below) does not change the inner product $\langle f(x), g(y) \rangle$ much, due to the fact that the components (x_j, y_j) are sampled independently from a bounded correlation space (namely $N_{p_1, p_2, \theta}$ for $\theta < 1$). The second is a multi-dimensional invariance principle asserting that these smoothed functions behave similarly on Gaussian inputs that have matching moments, with respect to Lipschitz test functions. We then apply this to the Lipschitz function which is the inner product of appropriately rounded versions of inputs, thereby yielding $\Delta([\ell])$ and $[0, 1]^\ell$ valued functions in the Gaussian domain with inner product close to $\langle f(x), g(y) \rangle$.

Definition A.2 (Bonami–Beckner $T_{1-\eta}$ operator). Let (Ω, μ) be a finite probability space, and $\eta \in (0, 1)$. For a function $h: \Omega^n \rightarrow \mathbb{R}$, the function $T_{1-\eta}h$ is defined as $T_{1-\eta}h(x) = \mathbb{E}_y[h(y)]$, where each coordinate y_i is sampled independently as follows:

- with probability $(1 - \eta)$ set $y_i = x_i$; and
- with probability η , pick $y_i \in \Omega$ as a fresh sample according to μ .

For a vector-valued function, $T_{1-\eta}$ acts component-wise, i.e., if $f = (f_1, \dots, f_\ell): \Omega^n \rightarrow \mathbb{R}^\ell$, we define $T_{1-\eta}f = (T_{1-\eta}f_1, \dots, T_{1-\eta}f_\ell)$.

³We assume that the parameters p_1, p_2, θ are such that each of the probabilities is in $[0, 1]$.

A useful property of the $T_{1-\eta}$ operator for us is that if h has convex range $K \subseteq [0, 1]^\ell$ then so does $T_{1-\eta}h$. As stated below, the action of $T_{1-\eta}$ has a particularly nice form when a function is expanded in an orthonormal basis, but this will not be important for us.

Fact A.3. *If a function $h: \Omega^n \rightarrow \mathbb{R}$ has multilinear expansion $h(x) = \sum_{\sigma} \hat{h}_{\sigma} \prod_{i=1}^n \chi_{\sigma_i}(x_i)$ w.r.t. an orthonormal ensemble $\mathcal{L} = (\chi_0, \dots, \chi_{b-1})$ of random variables over Ω , then the multilinear expansion of $T_{1-\eta}h$ is given by $\sum_{\sigma} \hat{h}_{\sigma} (1-\eta)^{|\sigma|} \prod_{i=1}^n \chi_{\sigma_i}(x_i)$.*

We next state the multi-dimensional invariance principle that we rely on. A version similar to the following is stated formally in [4, Theorem 10.1] (we have renamed some variables to avoid conflict with other uses in this paper) and it follows from Theorem 3.6 in the work of Isaksson and Mossel [6].

Theorem A.4. *Let (Ω, μ) be a finite probability space with the least non-zero probability of an atom being at least $\alpha \leq 1/2$. Let $b = |\Omega|$ and let $\mathcal{L} = \{\chi_0 = 1, \chi_1, \chi_2, \dots, \chi_{b-1}\}$ be a basis for random variables over Ω . Let $\Upsilon = \{\xi_0 = 1, \xi_1, \dots, \xi_{b-1}\}$ be an ensemble of real-valued Gaussian random variables with first and second moments matching those of the χ_i 's; specifically:*

$$\mathbb{E}[\chi_i] = \mathbb{E}[\xi_i] \quad \mathbb{E}[\chi_i^2] = \mathbb{E}[\xi_i^2] \quad \mathbb{E}[\chi_i \chi_j] = \mathbb{E}[\xi_i \xi_j] \quad \forall i, j \in \{1, \dots, b-1\}$$

Let $h = (h_1, h_2, \dots, h_t): \Omega^n \rightarrow \mathbb{R}^t$ be a vector-valued function such that $\text{Inf}_i(h_\ell) \leq \tau$ and $\text{Var}(h_\ell) \leq 1$ for all $i \in [n]$ and $\ell \in [t]$. For $\eta \in (0, 1)$, let H_ℓ , $\ell = 1, 2, \dots, t$, be the multilinear polynomial associated with $T_{1-\eta}h_\ell$ with respect to the basis \mathcal{L} , as per [Definition 6.5](#).

If $\Psi: \mathbb{R}^t \rightarrow \mathbb{R}$ is a Lipschitz-continuous function with Lipschitz constant Λ (with respect to the L_2 -norm), then

$$\left| \mathbb{E} \left[\Psi(H_1(\mathcal{L}^n), \dots, H_t(\mathcal{L}^n)) \right] - \mathbb{E} \left[\Psi(H_1(\Upsilon^n), \dots, H_t(\Upsilon^n)) \right] \right| \leq C(t) \cdot \Lambda \cdot \tau^{(\eta/18) \log(1/\alpha)} = o_\tau(1) \quad (10)$$

for some constant $C(t)$ depending on t , where $H_\ell(\mathcal{L}^n)$ and $H_\ell(\Upsilon^n)$, $\ell \in [t]$, denote random variables as in [Definition 6.6](#).

Armed with the above invariance principle, we now turn to the proof of [Theorem 7.3](#), restated below.

Theorem 7.3. *Fix any two parameters $p_1, p_2 \in (-1, 1)$. For all $\varepsilon \in (0, 1)$, $\ell \in \mathbb{Z}^+$, $\theta_0 \in [0, 1)$, and closed convex sets $K_1, K_2 \subseteq [0, 1]^\ell$ there exist $\tau > 0$ and mappings*

$$\begin{aligned} T_1^{(2)} &: \{f: \{+1, -1\}^n \rightarrow K_1\} \rightarrow \{F: \mathbb{R}^n \rightarrow K_1\} \\ T_2^{(2)} &: \{g: \{+1, -1\}^n \rightarrow K_2\} \rightarrow \{G: \mathbb{R}^n \rightarrow K_2\} \end{aligned}$$

such that for all $\theta \in [-\theta_0, \theta_0]$, if $f: \{+1, -1\}^n \rightarrow K_1$ and $g: \{+1, -1\}^n \rightarrow K_2$ satisfy

$$\max_{i \in [n]} \max \left(\max_{j \in [\ell]} \text{Inf}_i(f_j), \max_{j \in [\ell]} \text{Inf}_i(g_j) \right) \leq \tau$$

then for $F = T_1^{(2)}(f): \mathbb{R}^n \rightarrow K_1$ and $G = T_2^{(2)}(g): \mathbb{R}^n \rightarrow K_2$

$$\left| \mathbb{E}_{(x,y) \sim N^{\otimes n}} [\langle f(x), g(y) \rangle] - \mathbb{E}_{(X,Y) \sim \mathcal{G}^{\otimes n}} [\langle F(X), G(Y) \rangle] \right| \leq \varepsilon,$$

where $N = N_{p_1, p_2, \theta}$ and \mathcal{G} is the Gaussian distribution which matches the first and second-order moments of N .

Proof of Theorem 7.3. Let $\Omega = \{+1, -1\} \times \{+1, -1\}$ with the measure $N := N_{p_1, p_2, \theta}$. Define the basis $\mathcal{L} = \{\chi_0, \chi_1, \chi_2, \chi_3\}$ of functions on Ω as:

- $\chi_0 = 1$,
- $\chi_1((w_1, w_2)) = w_1$ (where $w_1, w_2 \in \{+1, -1\}$),
- $\chi_2((w_1, w_2)) = w_2$, and
- $\chi_3((w_1, w_2)) = w_1 w_2$.

We will apply the above invariance principle [Theorem A.4](#) with $t = 2\ell$, $h_j = f_j$ and $h_{\ell+j} = g_j$ for $j \in [\ell]$. We note that while $f_j, j \in [\ell]$ are functions on $\{+1, -1\}^n$, we can view them as functions on Ω^n by simply ignoring the second coordinate. (Thus, for $(x, y) \sim \Omega^n$, $f_j(x, y) = f_j(x)$.) The multilinear expansion of f_j w.r.t. \mathcal{L} will only involve χ_0 and χ_1 . Similarly, the functions h_j 's only depend on the second coordinate of Ω and have a multilinear expansion depending only on χ_0, χ_2 . The function $\Psi: \mathbb{R}^{2\ell} \rightarrow \mathbb{R}$ is defined as

$$\Psi(\mathbf{a}, \mathbf{b}) = \langle \text{Round}_{K_1}(\mathbf{a}), \text{Round}_{K_2}(\mathbf{b}) \rangle$$

for $\mathbf{a}, \mathbf{b} \in \mathbb{R}^\ell$, where for a *closed convex* set $K \subset \mathbb{R}^\ell$, $\text{Round}_K: \mathbb{R}^\ell \rightarrow \mathbb{R}^\ell$ maps a point to its (unique) closest point (in Euclidean distance) in K – in particular, it is the identity map on K . It is easy to see that by the convexity of K , Round_K is a 1-Lipschitz function,⁴ and it follows that the function Ψ is $O(\sqrt{\ell})$ -Lipschitz. Also, since $T_{1-\eta}f$ is K_1 -valued and $T_{1-\eta}g$ is K_2 -valued on $\{+1, -1\}^n$, the Round functions act as the identity on their images, and hence

$$\mathbb{E} \left[\Psi(H_1(\mathcal{L}^n), \dots, H_t(\mathcal{L}^n)) \right] = \mathbb{E}_{(x,y)} \left[\langle T_{1-\eta}f(x), T_{1-\eta}g(y) \rangle \right], \quad (11)$$

where (x, y) is distributed according to $N_{p_1, p_2, \theta}^{\otimes n}$.

For $j \in [\ell]$, define real-valued functions $\tilde{F}_j = H_j(\Upsilon^n)$ and $\tilde{G}_j = H_{\ell+j}(\Upsilon^n)$. Note that as the multilinear expansion of $T_{1-\eta}f_j$ (resp. $T_{1-\eta}h_j$) only involves χ_0, χ_1 (resp. χ_0, χ_2), the multilinear expansion of \tilde{F}_j (resp. \tilde{G}_j) only involves ξ_0, ξ_1 (resp. ξ_0, ξ_2). As $\xi_0 = 1$, the functions \tilde{F}_j (resp. \tilde{G}_j) are defined on \mathbb{R}^n under a product measure with coordinates distributed as Gaussians with mean p_1 (resp. mean p_2) and second moment 1.

Let $\tilde{F} = (\tilde{F}_1, \dots, \tilde{F}_\ell)$ and $\tilde{G} = (\tilde{G}_1, \dots, \tilde{G}_\ell)$, and finally let $F: \mathbb{R}^n \rightarrow K_1$ be $F(X) = \text{Round}_{K_1}(\tilde{F}(X))$ and $G: \mathbb{R}^n \rightarrow K_2$ be $G(Y) = \text{Round}_{K_2}(\tilde{G}(Y))$. Note that F (resp. G) depends only on $f = (f_1, \dots, f_\ell)$ (resp. $g = (g_1, \dots, g_\ell)$) as required in the statement of [Theorem 7.3](#). By construction, it is clear that

$$\mathbb{E} \left[\Psi(H_1(\Upsilon^n), \dots, H_t(\Upsilon^n)) \right] = \mathbb{E}_{(X,Y)} \left[\langle F(X), G(Y) \rangle \right], \quad (12)$$

for $(X, Y) \sim (\xi_1, \xi_2)^{\otimes n} = \mathcal{G}^{\otimes n}$ where \mathcal{G} is the Gaussian distribution which matches the first and second moments of $N = N_{p_1, p_2, \theta}$.

Combining (11) and (12) with the guarantee (10) of [Theorem A.4](#), we get that

$$\left| \mathbb{E}_{(x,y) \sim N^{\otimes n}} \left[\langle T_{1-\eta}f(x), T_{1-\eta}g(y) \rangle \right] - \mathbb{E}_{(X,Y) \sim \mathcal{G}^{\otimes n}} \left[\langle F(X), G(Y) \rangle \right] \right| \leq \varepsilon/2 \quad (13)$$

⁴ To see why, let a, b be two arbitrary points and $a' = \text{Round}_K(a)$, $b' = \text{Round}_K(b)$. Without loss of generality, we can change the coordinates so that $a' = (0, \dots, 0)$ and $b' = (c, 0, \dots, 0)$ for some $c > 0$: by convexity, the segment $[a'b']$ lies within K . Now, by virtue of a' (resp. b') being the closest point to a (resp. b), this implies the first coordinate of a must be non-positive and the first coordinate of b must be at least c ; but this in turn means the distance between a and b is at least c .

for $\tau > 0$ chosen small enough (as a function of $\varepsilon, \ell, p_1, p_2, \theta_0$ and η). We are almost done, except that we would like to be close to the inner product $\langle f(x), g(y) \rangle$ of the original functions, and we have the noised versions in (13) above. However, as the correlation of the space $N_{p_1, p_2, \theta}$ is bounded away from 1, applying Lemma 6.1 of [12] implies that for small enough $\eta > 0$ (as a function of $\varepsilon, \ell, p_1, p_2, \theta_0$),

$$\left| \mathbb{E}_{(x,y) \sim N^{\otimes n}} [\langle T_{1-\eta} f(x), T_{1-\eta} g(y) \rangle] - \mathbb{E}_{(x,y) \sim N^{\otimes n}} [\langle f(x), g(y) \rangle] \right| \leq \varepsilon/2 .$$

Combining this with (13), the proof of **Theorem 7.3** is complete. □