

Communication with Contextual Uncertainty

Badih Ghazi ^{*} Ilan Komargodski [†] Pravesh Kothari [‡] Madhu Sudan [§]

July 19, 2015

Abstract

We introduce a simple model illustrating the role of context in communication and the challenge posed by uncertainty of knowledge of context. We consider a variant of distributional communication complexity where Alice gets some information x and Bob gets y , where (x, y) is drawn from a known distribution, and Bob wishes to compute some function $g(x, y)$ (with high probability over (x, y)). In our variant, Alice does not know g , but only knows some function f which is an approximation of g . Thus, the function being computed forms the context for the communication, and knowing it imperfectly models (mild) uncertainty in this context.

A naive solution would be for Alice and Bob to first agree on some common function h that is close to both f and g and then use a protocol for h to compute $h(x, y)$. We show that any such agreement leads to a large overhead in communication ruling out such a universal solution.

In contrast, we show that if g has a one-way communication protocol with complexity k in the standard setting, then it has a communication protocol with complexity $O(k \cdot (1 + I))$ in the uncertain setting, where I denotes the mutual information between x and y . In the particular case where the input distribution is a product distribution, the protocol in the uncertain setting only incurs a constant factor blow-up in communication and error.

Furthermore, we show that the dependence on the mutual information I is required. Namely, we construct a class of functions along with a non-product distribution over (x, y) for which the communication complexity is a single bit in the standard setting but at least $\Omega(\sqrt{n})$ bits in the uncertain setting.

^{*}Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge MA 02139. Supported in part by NSF STC Award CCF 0939370 and NSF Award CCF-1217423. badih@mit.edu.

[†]Weizmann Institute of Science, Israel. Email: ilan.komargodski@weizmann.ac.il. Work done while an intern at MSR New England. Supported in part by a grant from the I-CORE Program of the Planning and Budgeting Committee, the Israel Science Foundation, BSF and the Israeli Ministry of Science and Technology.

[‡]UT Austin, USA. Email: kothari@cs.utexas.edu. Work done while an intern at MSR New England.

[§]Microsoft Research, One Memorial Drive, Cambridge, MA 02142, USA. Email: madhu@mit.edu.

1 Introduction

Most forms of communication involve communicating players that share a large common context and use this context to compress communication. In natural settings, the context may include understanding of language, and knowledge of the environment and laws. In designed (computer-to-computer) settings, the context includes knowledge of the operating system, communication protocols, and encoding/decoding mechanisms. Remarkably, especially in the natural setting, context can seemingly be used to compress communication, even when it is not shared perfectly. This ability to communicate despite a major source of uncertainty has led to a series of works attempting to model various forms of communication amid uncertainty, starting with Goldreich, Juba and Sudan [JS08, GJS12] followed by [JKKS11, JS11, JW13, HS14, CGMS15]. This current work introduces a new theme to this series of works by introducing a functional notion of uncertainty and studying this model. We start by describing our model and results below and then contrast our model with some of the previous works.

Model. Our model builds upon the classical setup of communication complexity due to Yao [Yao79], and we develop it here. The classical model considers two interacting players Alice and Bob each possessing some private information x and y with x known only to Alice and y to Bob. They wish to compute some joint function $g(x, y)$ and would like to do so while exchanging the minimum possible number of bits. In this work, we suggest that the function g is the *context* of the communication and consider a setting where it is shared imperfectly. Specifically, we say that Bob knows the function g and Alice knows some approximation f to g (with f not being known to Bob). This leads to the question: when can Alice and Bob interact to compute $g(x, y)$ with limited communication ?

It is clear that if $x \in \{0, 1\}^n$, then n bits of communication suffice — Alice can simply ignore f and send x to Bob. We wish to consider settings that improve on this. To do so correctly on every input, a necessary condition is that g must have low communication complexity in the standard model. However, this necessary condition does not seem to be sufficient — since Alice only has an approximation f to g . Thus, we settle for a weaker goal: determining g correctly only on most inputs. This puts us in a distributional communication complexity setting. A necessary condition now is that g must have a low-error low-communication protocol in the standard setting. The question is then: can g be computed with low error and low communication when Alice only knows an approximation f to g (with f being unknown to Bob) ?

More precisely, in this setting, the input to Alice is a pair (f, x) and the input to Bob is a pair (g, y) . The functions (f, g) are adversarially chosen subject to the restrictions that they are close to each other (under some distribution μ on the inputs) and that g (and hence f) has a low-error low-communication protocol. The pair (x, y) is drawn from the distribution μ (independent of the choice of f and g). The players both know μ in addition to their respective inputs.

Results. In order to describe our results, we first introduce some notation. Let $\delta_\mu(f, g)$ denote the (weighted and normalized) Hamming distance between f and g with respect to the distribution μ . Let $\text{CC}_\epsilon^\mu(f)$ denote the minimum communication complexity of a protocol computing f correctly on all but an ϵ fraction of the inputs. Let $\text{owCC}_\epsilon^\mu(f)$ denote the corresponding *one-way* communication complexity of f . Given a family \mathcal{F} of pairs of functions (f, g) , we denote the uncertain complexity $\text{CCU}_\epsilon^\mu(\mathcal{F})$ to be the minimum over all public-coin protocols Π of the maximum over $(f, g) \in \mathcal{F}$, (x, y) in the support of μ and settings of public coins, of the communication cost of Π , subject to

the condition that for every $(f, g) \in \mathcal{F}$, Π outputs $g(x, y)$ with probability $1 - \epsilon$ over the choice of (x, y) and the shared randomness. That is,

$$\text{CCU}_\epsilon^\mu(\mathcal{F}) \triangleq \min_{\{\Pi \mid \forall (f, g) \in \mathcal{F}: \delta_\mu(\Pi, g) \leq \epsilon\}} \max_{\{(f, g) \in \mathcal{F}, (x, y) \in \text{supp}(\mu), \text{public coins}\}} \{\text{Comm. cost of } \Pi((f, x), (g, y))\}.$$

Similarly, let $\text{owCCU}_\epsilon^\mu(\mathcal{F})$ denote the one-way uncertain communication complexity of \mathcal{F} .

Our first result (Theorem 1.1) shows that if μ is a distribution on which f and g are close and each has a one-way protocol with communication k bits in the standard model, then the pair (f, g) has one-way uncertain communication complexity of at most $O(k \cdot (1 + I))$ bits with I being the mutual information¹ of $(x, y) \sim \mu$. More precisely, let $\text{ow}\mathcal{F}_{k, \epsilon, \delta}^\mu$ denote the family of all pairs of functions (f, g) with $\text{owCC}_\epsilon^\mu(f), \text{owCC}_\epsilon^\mu(g) \leq k$ and $\delta_\mu(f, g) \leq \delta$. We prove the following theorem.

Theorem 1.1. *There exists an absolute constant c such that for every pair of finite sets X and Y , every distribution μ over $X \times Y$ and every $\theta > 0$, it holds that*

$$\text{owCCU}_{\epsilon+2\delta+\theta}^\mu(\text{ow}\mathcal{F}_{k, \epsilon, \delta}^\mu) \leq \frac{c(k + \log(\frac{1}{\theta}))}{\theta^2} \cdot \left(1 + \frac{I(X; Y)}{\theta^2}\right). \quad (1)$$

In the special case where μ is a product distribution, then $I(X; Y) = 0$ and we obtain the following particularly interesting corollary of Theorem 1.1.

Corollary 1.2. *There exists an absolute constant c such that for every pair of finite sets X and Y , every product distribution μ over $X \times Y$ and every $\theta > 0$, it holds that*

$$\text{owCCU}_{\epsilon+2\delta+\theta}^\mu(\text{ow}\mathcal{F}_{k, \epsilon, \delta}^\mu) \leq \frac{c(k + \log(\frac{1}{\theta}))}{\theta^2}.$$

In words, Corollary 1.2 says that for product distributions and for constant error probabilities, communication in the uncertain model is only a constant factor larger than in the standard model.

Our result is significant in that it achieves (moderately) reliable communication despite uncertainty about the context, even when the uncertainty itself is hard to resolve. To elaborate on this statement, note that one hope for achieving a low-communication protocol for g would be for Alice and Bob to first agree on some function q that is close to f and g , and then apply some low-communication protocol for this common function q . This would be the “resolve the uncertainty first” approach. We prove (Theorem 3.2) that resolving the uncertainty can be very expensive (much more so than even the trivial protocol of sending x) and hence, this would not be a way to prove Theorem 1.1. Instead, we show a path around the inherent uncertainty to computing the desired function, and this leads to a proof of Theorem 1.1. To handle non-product distributions in Theorem 1.1, we in particular use a *one-way distributional* variant of the correlated sampling protocol of Braverman and Rao [BR11]. For a high-level overview of the proof of Theorem 1.1, we refer the reader to Section 4.1.

We now describe our lower bound. Given the upper bound in Theorem 1.1, a natural question is whether the dependence on $I(X; Y)$ in the right-hand side of Equation (1) is actually needed. In other words, is it also the case that for non-product distributions, contextual uncertainty can only cause a constant-factor blow-up in communication (for constant error probabilities)? Perhaps surprisingly, the answer to this question turns out to be negative. Namely, we show that a dependence of the communication in the uncertain setting on $I(X; Y)$ is required.

¹Given a distribution μ over a pair (X, Y) of random variables with marginals μ_X and μ_Y over X and Y respectively, the *mutual information* of X and Y is defined as $I(X; Y) \triangleq \mathbb{E}_{(x, y) \sim \mu}[\log(\frac{\mu(x, y)}{\mu_X(x)\mu_Y(y)})]$.

Theorem 1.3. *There exist a distribution μ and a function class $\mathcal{F} \subseteq \text{ow}\mathcal{F}_{1,0,\delta}^\mu$ such that for every $\epsilon > 0$,*

$$\text{CCU}_{\frac{1}{2}-\epsilon}^\mu(\mathcal{F}) \geq \Omega(\sqrt{\delta n}) - \log(1/\epsilon).$$

In particular, if δ is any small constant (e.g., $1/5$), then Theorem 1.3 asserts the existence of a distribution and a class of distance- δ functions for which the zero-error (one-way) communication complexity in the standard model is a single bit, but under contextual uncertainty, any two-way protocol (with an arbitrary number of rounds of interaction) having a noticeable advantage over random guessing requires $\Omega(\sqrt{n})$ bits of communication! We note that the distribution μ in Theorem 1.3 has mutual information $\approx n$, so Theorem 1.3 rules out improving the dependence on the mutual information in Equation (1) to anything smaller than $\sqrt{I(X;Y)}$. It is an interesting open question to determine the correct exponent of $I(X;Y)$ in Equation (1).

In order to prove Theorem 1.3, the function class \mathcal{F} will essentially consist of the set of all close-by pairs of parity functions and the distribution μ will correspond to the *noisy Boolean hypercube*. We are then able to reduce the problem of computing \mathcal{F} under μ with contextual uncertainty, to the problem of computing a related function in the standard distributional communication complexity model (i.e., *without* uncertainty) under a related distribution. We then use the *discrepancy method* to prove a lower bound on the communication complexity of the new problem. This task itself reduces to upper bounding the spectral norm of a certain communication matrix. The choice of our underlying distribution then implies a tensor structure for this matrix, which reduces the spectral norm computation to bounding the largest singular value of an explicit family of 4×4 matrices. For more details about the proof of Theorem 1.3, we refer the reader to Section 5.

Contrast with prior work. The first works to consider communication with uncertainty in a manner similar to this work were those of [JS08, GJS12]. Their goal was to model an extreme form of uncertainty, where Alice and Bob do not have any prior (known) commonality in context and indeed both come with their own “protocol” which tells them how to communicate. So communication is needed even to resolve this uncertainty. While their setting is thus very broad, the solutions they propose are much slower and typically involve resolving the uncertainty as a first step.

The later works [JKKS11, HS14, CGMS15] tried to restrict the forms of uncertainty to see when it could lead to more efficient communication solutions. For instance, Juba et al. [JKKS11] consider the compression problem when Alice and Bob do not completely agree on the prior. This introduces some uncertainty in the beliefs, and they provide fairly efficient solutions by restricting the uncertainty to a manageable form. Canonne et al. [CGMS15] were the first to connect this stream of work to communication complexity, which seems to be the right umbrella to study the broader communication problems. The imperfectness they study is however restricted to the randomness shared by the communicating parties, and does not incorporate any other elements. They suggest studying imperfect understanding of the function being computed as a general direction, though they do not suggest specific definitions, which we in particular do in this work.

Organization In Section 2, we carefully develop the uncertain communication complexity model after recalling the standard distributional communication complexity model. In Section 3, we prove the hardness of contextual agreement. In Section 4, we prove our main upper bound (Theorem 1.1). In Section 5, we prove our main lower bound (Theorem 1.3). For a discussion of some intriguing future directions that arise from this work, we refer the reader to the conclusion section 6.

2 The Uncertain Communication Complexity Model

We start by recalling the classical communication complexity model of Yao [Yao79] and then present our definition and measures.

2.1 Communication Complexity

We start with some basic notation. For an integer $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, \dots, n\}$. We use $\log x$ to denote a logarithm in base 2. For two sets A and B , we denote by $A \Delta B$ their symmetric difference. For a distribution μ , we denote by $x \sim \mu$ the process of sampling a value x from the distribution μ . Similarly, for a set X we denote by $x \sim X$ the process of sampling a value x from the uniform distribution over X . For any event E , let $\mathbb{1}(E)$ be the 0-1 indicator of E . For a probability distribution μ over $X \times Y$, we denote by μ_X the marginal of μ over X . By $\mu_{Y|x}$, we denote the conditional distribution of μ over Y conditioned on $X = x$.

Given a distribution μ supported on X and functions $f, g: X \rightarrow \Sigma$, we let $\delta_\mu(f, g)$ denote the (weighted and normalized) Hamming distance between f and g , i.e., $\delta_\mu(f, g) \triangleq \Pr_{x \sim \mu}[f(x) \neq g(x)]$. (Note that this definition extends naturally to probabilistic functions f and g – by letting $f(x)$ and $g(x)$ be sampled independently.)

We now turn to the definition of *communication complexity*. A more thorough introduction can be found in [KN97]. Let $f: X \times Y \rightarrow \{0, 1\}$ be a function and Alice and Bob be two parties. A protocol Π between Alice and Bob specifies how and what Alice and Bob communicate given their respective inputs and communication thus far. It also specifies when they stop and produce an output (that we require to be produced by Bob). A protocol is said to be one-way if it involves a single message from Alice to Bob, followed by Bob producing the output. The protocol Π is said to compute f if for every $(x, y) \in X \times Y$ it holds that $\Pi(x, y) = f(x, y)$. The communication complexity of Π is the number of bits transmitted during the execution of the protocol between Alice and Bob. The communication complexity of f is the minimal communication complexity of a protocol computing f .

It is standard to relax the above setting by introducing a distribution μ over the input space $X \times Y$ and requiring the protocol to succeed with high probability (rather than with probability 1). We say that a protocol Π ϵ -computes a function f under distribution μ if $\delta_\mu(\Pi(x, y), f(x, y)) \leq \epsilon$.

Definition 2.1 (Distributional Communication Complexity). *Let $f: X \times Y \rightarrow \{0, 1\}$ be a Boolean function and μ be a probability distribution over $X \times Y$. The distributional communication complexity of f under μ with error ϵ , denoted by $\text{CC}_\epsilon^\mu(f)$, is defined as the minimum over all protocols Π that ϵ -compute f over μ , of the communication complexity of Π . The one-way communication complexity $\text{owCC}_\epsilon^\mu(f)$ is defined similarly by minimizing over one-way protocols Π .*

We note that it is also standard to provide Alice and Bob with a shared random string which is independent of x, y and f . In the distributional communication complexity model, it is a known fact that any protocol with shared randomness can be used to get a protocol that does *not* use shared randomness without increasing its distributed communication complexity.

In this paper, unless stated otherwise, whenever we refer to a protocol, we think of the input pair (x, y) as coming from a distribution.

2.2 Uncertain Communication Complexity

We now turn to the central definition of this paper, namely *uncertain communication complexity*. Our goal is to understand how Alice and Bob can communicate when the function that Bob wishes to determine is not known to Alice. In this setting, we make the functions g (that Bob wants to compute) and f (Alice's estimate of g) explicitly part of the input to the protocol Π . Thus, in this setting a protocol Π specifies how Alice with input (f, x) and Bob with input (g, y) communicate, and how they stop and produce an output. We say that Π computes (f, g) if for every $(x, y) \in X \times Y$, the protocol outputs $g(x, y)$. We say that a (public-coin) protocol Π ϵ -computes (f, g) over μ if $\delta_\mu(g, \Pi) \leq \epsilon$.

Next, one may be tempted to define the communication complexity of a pair of functions (f, g) as the minimum over all protocols that compute (f, g) of their maximum communication. But this does not capture the uncertainty! (Rather, a protocol that works for the pair corresponds to both Alice and Bob knowing both f and g .) To model uncertainty, we have to consider the communication complexity of a whole class of pairs of functions, from which the pair (f, g) is chosen (in our case by an adversary).

Let $\mathcal{F} \subseteq \{f: X \times Y \rightarrow \{0, 1\}\}^2$ be a family of pairs of Boolean functions with domain $X \times Y$. We say that a public-coin protocol Π ϵ -computes \mathcal{F} over μ if for every $(f, g) \in \mathcal{F}$, we have that Π ϵ -computes (f, g) over μ . We are now ready to present our main definition.

Definition 2.2 (Contextually Uncertain Communication Complexity). *Let μ be a distribution on $X \times Y$ and $\mathcal{F} \subseteq \{f: X \times Y \rightarrow \{0, 1\}\}^2$. The communication complexity of \mathcal{F} under contextual uncertainty, denoted $\text{CCU}_\epsilon^\mu(\mathcal{F})$, is the minimum over all public-coin protocols Π that ϵ -compute \mathcal{F} over μ , of the maximum communication complexity of Π over $(f, g) \in \mathcal{F}$, (x, y) from the support of μ and settings of the public coins.*

As usual, the one-way contextually uncertain communication complexity $\text{owCCU}_\epsilon^\mu(\mathcal{F})$ is defined similarly.

We remark that while in the standard distributional model of Subsection 2.1, shared randomness can be assumed *without loss of generality*, this is not necessarily the case in Definition 2.2. This is because in principle, shared randomness can help fool the adversary who is selecting the pair $(f, g) \in \mathcal{F}$. Also, observe that in the special case where $\mathcal{F} = \{(f, g)\}$, Definition 2.2 boils down to the standard definition of distributional communication complexity (i.e., Definition 2.1) for the function g , and we thus have $\text{CCU}_\epsilon^\mu(\{(f, g)\}) = \text{CC}_\epsilon^\mu(g)$. Furthermore, the uncertain communication complexity is monotone, i.e., if $\mathcal{F} \subseteq \mathcal{F}'$ then $\text{CCU}_\epsilon^\mu(\mathcal{F}) \leq \text{CCU}_\epsilon^\mu(\mathcal{F}')$. Hence, we conclude that $\text{CCU}_\epsilon^\mu(\mathcal{F}) \geq \max_{\{g \mid \exists f \text{ s.t. } (f, g) \in \mathcal{F}\}} \{\text{CC}_\epsilon^\mu(g)\}$.

In this work, we attempt to identify a setting under which the above lower bound can be matched. If the set of functions $\Gamma(g) = \{f \mid (f, g) \in \mathcal{F}\}$ is not sufficiently informative about g , then it seems hard to conceive of settings where Alice can do non-trivially well. We thus pick a simple and natural restriction on $\Gamma(g)$, namely, that it contains functions that are close to g (in δ_μ -distance). This leads us to our main target classes. For parameters $k, \epsilon, \delta > 0$, define the sets of pairs of functions

$$\mathcal{F}_{k, \epsilon, \delta} \triangleq \{(f, g) \mid \delta_\mu(f, g) \leq \delta \ \& \ \text{CC}_\epsilon^\mu(f), \text{CC}_\epsilon^\mu(g) \leq k\}$$

and

$$\text{ow}\mathcal{F}_{k, \epsilon, \delta} \triangleq \{(f, g) \mid \delta_\mu(f, g) \leq \delta \ \& \ \text{owCC}_\epsilon^\mu(f), \text{owCC}_\epsilon^\mu(g) \leq k\}.$$

In words, $\mathcal{F}_{k,\epsilon,\delta}$ (resp. $\text{ow}\mathcal{F}_{k,\epsilon,\delta}$) considers all possible functions g with communication complexity (resp. one-way communication complexity) at most k with Alice being roughly under all possible uncertainties within distance δ of Bob.²

It is clear that $\text{owCC}_\epsilon^\mu(\text{ow}\mathcal{F}_{k,\epsilon,\delta}) \geq k$. Our first main result, Theorem 1.1, gives an upper bound on this quantity, which in the particular case of product distributions is comparable to k (up to a constant factor increase in the error and communication complexity). In Theorem 3.2 we show that a naive strategy that attempts to reduce the uncertain communication problem to a “function agreement problem” (where Alice and Bob agree on a function q that is close to f and g and then use a protocol for q) cannot work. Furthermore, our second main result, Theorem 1.3, shows that for general non-product distributions, $\text{CC}_\epsilon^\mu(\text{ow}\mathcal{F}_{k,\epsilon,\delta})$ can be much larger than k . More precisely, we construct a function class along with a distribution μ for which the one-way communication complexity in the standard model is a single bit whereas, under contextual uncertainty, the two-way communication complexity is at least $\Omega(\sqrt{n})!$

3 Hardness of Contextual Agreement

In this section, we show that even if both f and g have small one-way distributional communication complexity on some distribution μ , agreeing on a q such that $\delta_\mu(q, f)$ is small takes communication that is roughly the size of the bit representation of f (which is exponential in the size of the input). Thus, agreeing on q before simulating a protocol for q is exponentially costlier than even the trivial protocol where Alice sends her input x to Bob. Formally, we consider the following communication problem:

Definition 3.1 ($\text{AGREE}_{\delta,\gamma}(\mathcal{F})$). *For a family of pairs of functions $\mathcal{F} \subseteq \{f: X \times Y \rightarrow \{0, 1\}\}^2$, the \mathcal{F} -agreement problem with parameters $\delta, \gamma \geq 0$ is the communication problem where Alice gets f and Bob gets g such that $(f, g) \in \mathcal{F}$ and their goal is for Alice to output q_A and Bob to output q_B such that $\delta(q_A, f), \delta(q_B, g) \leq \delta$ and $\Pr[q_A = q_B] \geq \gamma$.*

Somewhat abusing notation, we will use $\text{AGREE}_{\delta,\gamma}(\mathcal{D})$ to denote the distributional problem where \mathcal{D} is a distribution on $\{f: X \times Y \rightarrow \{0, 1\}\}^2$ and the goal now is to get agreement with probability δ over the randomness of the protocol and the input.

If the agreement problem could be solved with low communication for the family $\mathcal{F}_{k,\delta,\epsilon}$ as defined at the end of Section 2, then this would turn into a natural protocol for $\text{CCU}(\mathcal{F}_{k,\epsilon',\delta'})$ for some positive ϵ' and δ' as well. Our following theorem proves that agreement is a huge overkill.

Theorem 3.2. *For every $\delta, \delta_2 > 0$, there exists $\alpha > 0$ and a family $\mathcal{F} \subseteq \mathcal{F}_{0,0,\delta}$ such that for every $\gamma > 0$, it holds that $\text{CC}(\text{AGREE}_{\delta_2,\gamma}(\mathcal{F})) \geq \alpha|Y| - \log(1/\gamma)$.*

In words, Theorem 3.2 says that there is a family of pairs of functions supported on functions of zero communication complexity (with zero error) for which agreement takes communication polynomial in the size of the domain of the functions. Note that this is exponentially larger than the trivial communication complexity for any function g , which is at most $\min\{1 + \log|Y|, \log|X|\}$.

We stress that while an agreement lower bound for zero communication functions may feel a lower bound for a toy problem, a lower bound for this setting is inherent in any separation between

²For the sake of symmetry, we insist that $\text{CC}_\epsilon^\mu(f) \leq k$ (resp. $\text{owCC}_\epsilon^\mu(f) \leq k$). We need not have insisted on it but since the other conditions anyhow imply that $\text{CC}_{\epsilon+\delta}^\mu(f) \leq k$ (resp. $\text{owCC}_{\epsilon+\delta}^\mu(f) \leq k$), we decided to include this stronger condition for aesthetic reasons.

agreement complexity for \mathcal{F} and communication complexity with uncertainty for \mathcal{F} . To see this, note that given any input to the $\text{CCU}(\mathcal{F})$ problem, Alice and Bob can execute any protocol for $\text{CCU}(\mathcal{F})$ pinning down the value of the function to be computed with high probability and low communication. If one considers the remaining challenge to agreement, it comes from a zero communication problem.

Our proof of Theorem 3.2 uses a lower bound on the communication complexity of *agreement distillation (with imperfectly shared randomness)* problem defined in [CGMS15], who in turn rely on a lower bound for randomness extraction from correlated sources due to Bogdanov and Mosser [BM11].

We describe their problem below and the result that we use. We note that their context is slightly different and our description below is a reformulation. First, we define the notion of ρ -perturbed sequences of bits. A pair of bits (a, b) is said to be a pair of ρ -perturbed uniform bits if a is uniform over $\{0, 1\}$, and $b = a$ with probability $1 - \rho$ and $b \neq a$ with probability ρ . A pair of sequences of bits (r, s) is said to be ρ -perturbed if $r = (r_1, \dots, r_n)$ and $s = (s_1, \dots, s_n)$ and each coordinate pair (r_i, s_i) is a ρ -perturbed uniform pair drawn independently of all other pairs. For a random variable W , we define its min-entropy as $H_\infty(w) \triangleq \min_{w \in \text{supp}(W)} \{-\log(\Pr[W = w])\}$.

Definition 3.3 ($\text{AGREEMENT-DISTILLATION}_{\gamma, \rho}^k$). *In this problem, Alice and Bob get as inputs r and s , where (r, s) form a ρ -perturbed sequence of bits. Their goal is to communicate deterministically and produce as outputs w_A (Alice's output) and w_B (Bob's output) with the following properties: (i) $H_\infty(w_A), H_\infty(w_B) \geq k$ and (ii) $\Pr_{(r, s)}[w_A = w_B] \geq \gamma$.*

Lemma 3.4 ([CGMS15, Theorem 2]). *For every $\rho > 0$, there exists $\epsilon > 0$ such that for every k and γ , it holds that every deterministic protocol Π that computes $\text{AGREE}_{\gamma, \rho}^k$ has communication complexity at least $\epsilon k - \log 1/\gamma$.*

We note that while the agreement distillation problem is very similar to our agreement problem, there are some syntactic differences. We are considering pairs of functions with low communication complexity, whereas the agreement-distillation problem considers arbitrary random sequences. Also, our output criterion is proximity to the input functions, whereas in the agreement-distillation problem, we need to produce high-entropy outputs. Finally, we want a lower bound for our agreement problem when Alice and Bob are allowed to share perfect randomness while the agreement-distillation bound only holds for deterministic protocols. Nevertheless, we are able to reduce to their setting quite easily as we will see shortly.

Our proof of Theorem 3.2 uses the standard Chernoff-Hoeffding tail inequality on random variables that we include below. Denote $\exp(x) \triangleq e^x$, where e is the base of the natural logarithm.

Proposition 3.5 (Chernoff bound). *Let $X = \sum_{i=1}^n X_i$ be a sum of identically distributed independent random variables $X_1, \dots, X_n \in \{0, 1\}$. Let $\mu = \mathbb{E}[X] = \sum_{i=1}^n \mathbb{E}[X_i]$. It holds that for $\delta \in (0, 1)$,*

$$\Pr[X < (1 - \delta)\mu] \leq \exp(-\delta^2\mu/2)$$

and

$$\Pr[X > (1 + \delta)\mu] \leq \exp(-\delta^2\mu/3),$$

and for $a > 0$,

$$\Pr[X > \mu + a] \leq \exp(-2a^2/n)$$

Proof of Theorem 3.2. We prove the theorem for $\alpha < \delta/6$, in which case we may assume $\gamma > \exp(-\delta|Y|/6)$ since otherwise the right-hand side is non-positive.

Let \mathcal{F}_B denote the set of functions that depend only on Bob's inputs, i.e., $f \in \mathcal{F}_B$ if there exists $f': Y \rightarrow \{0, 1\}$ such that $f(x, y) = f'(y)$ for all x, y . Our family \mathcal{F} will be a subset of $\mathcal{F}_B \times \mathcal{F}_B$, the subset that contains functions that are at most $\delta|Y|$ apart.

$$\mathcal{F} \triangleq \{(f, g) \in \mathcal{F}_B \times \mathcal{F}_B \mid \delta(f, g) \leq \delta\}.$$

It is clear that communication complexity of every function in the support of \mathcal{F} is zero, with zero error (Bob can compute it on his own) and so $\mathcal{F} \subseteq \mathcal{F}_{0,0,\delta}$. So it remains to prove a lower bound on $\text{CC}(\text{AGREE}_{\delta_2, \gamma}(\mathcal{F}))$.

We prove our lower bound by picking a distribution \mathcal{D}_ρ supported mostly on \mathcal{F} and by giving a lower bound on $\text{CC}(\text{AGREE}_{\delta_2, \gamma}(\mathcal{D}_\rho))$. Let $\rho = \delta/2$. The distribution \mathcal{D}_ρ is a simple one. It samples (f, g) as follows. The function f is drawn uniformly at random from \mathcal{F}_B . Then, g is chosen to be a “ ρ -perturbation” of f , namely for every $y \in Y$, $g'(y)$ is chosen to be equal to $f(x, y)$ with probability $1 - \rho$ and $1 - f(x, y)$ with probability ρ . For every $x \in X$, we now set $g(x, y) = g'(x, y)$.

By the Chernoff bound (see Proposition 3.5), we have that $\Pr_{(f,g) \sim \mathcal{D}_\rho}[\delta(f, g) > \delta] = \exp(-\rho|Y|/3) \leq \gamma$. So with overwhelmingly high probability, \mathcal{D}_ρ draws elements from \mathcal{F} . In particular, if some protocol solves $\text{AGREE}_{\delta_2, \gamma}(\mathcal{F})$, then it would also solve $\text{AGREE}_{\delta_2, 2\gamma}(\mathcal{D}_\rho)$.

We thus need to show a lower bound on the communication complexity of $\text{AGREE}_{\delta_2, 2\gamma}(\mathcal{D}_\rho)$. We now note that since this is a distributional problem, by Yao's min-max principle, if there is randomized protocol to solve $\text{AGREE}_{\delta_2, 2\gamma}(\mathcal{D}_\rho)$ with C bits of communication, then there is also a deterministic protocol for the same problem and with the same complexity. Thus, it suffices to lower bound the deterministic communication complexity of $\text{AGREE}_{\delta_2, 2\gamma}(\mathcal{D}_\rho)$. Claim 3.6 shows that any such protocol gives a deterministic protocol for AGREEMENT-DISTILLATION with $k = \Omega_{\delta_2}(|Y|)$. Combining this with Lemma 3.4 gives us the desired lower bound on $\text{CC}(\text{AGREE}_{\delta_2, 2\gamma}(\mathcal{D}_\rho))$ and hence on $\text{CC}(\text{AGREE}_{\delta_2, \gamma}(\mathcal{F}))$. \square

Claim 3.6. *Every protocol for $\text{AGREE}_{\delta_2, \gamma}(\mathcal{D}_\rho)$ is also a protocol for $\text{AGREEMENT-DISTILLATION}_{\gamma, \rho}^k$ for $k = (1 - h(\delta_2))|Y|$, where $h(\cdot)$ is the binary entropy function given by $h(x) = -x \log x - (1 - x) \log(1 - x)$.*

Proof. Suppose Alice and Bob are trying to solve $\text{AGREEMENT-DISTILLATION}_{\gamma, \rho}^k$. They can sample ρ -perturbed strings $(r, s) \in \{0, 1\}^{|Y|}$ and interpret them as functions $f', g': Y \rightarrow \{0, 1\}$ or equivalently as functions $(f, g) \sim \mathcal{D}_\rho$. They can now simulate the protocol for $\text{AGREE}_{\delta, \gamma}(f, g)$ and output q_A and q_B . By definition of AGREE , we have $q_A = q_B$ with probability at least γ . So it suffices to show that $H_\infty(q_A), H_\infty(q_B) \geq k$. But this is obvious since any function q_A is output only if $\delta(f, q_A) \leq \delta_2$ and we have that $|\{f \mid \delta(f, q_A) \leq \delta\}| \leq 2^{h(\delta_2)|Y|}$. Since the probability of sampling f for any f is at most $2^{-|Y|}$, we have that the probability of outputting q_A for any q_A is at most $2^{-(1-h(\delta_2))|Y|}$. In other words, $H_\infty(q_A) \geq (1 - h(\delta_2))|Y|$. Similarly, we can lower bound $H_\infty(q_B)$ and thus we have that the outputs of the protocol for AGREE solve $\text{AGREEMENT-DISTILLATION}$ with $k = (1 - h(\delta_2))|Y|$. \square

4 One-way Communication with Contextual Uncertainty

In this section, we prove Theorem 1.1. We start with a high-level description of the protocol.

4.1 Overview of Protocol

Let μ be a distribution over an input space $X \times Y$. For any function $s: X \times Y \rightarrow \{0, 1\}$ and any $x \in X$, we define the *restriction* of s to x to be the function $s_x: Y \rightarrow \{0, 1\}$ given by $s_x(y) = s(x, y)$ for any $y \in Y$.

We now give a high-level overview of the protocol. First, we consider the particular case of Theorem 1.1 where μ is a product distribution, i.e., $\mu = \mu_X \times \mu_Y$. Note that in this case, $I(X; Y) = 0$ in the right-hand side of Equation (1). We will handle the case of general (not necessarily product) distributions later on.

The general idea is that given inputs (f, x) , Alice can determine the restriction f_x , and she will try to describe it to Bob. For most values of x , f_x will be close (in δ_{μ_Y} -distance) to the function g_x . Bob will try to use the (yet unspecified) description given by Alice in order to determine some function B that is close to g_x . If he succeeds in doing so, he can output $B(y)$ which would equal $g_x(y)$ with high probability over y .

We next explain how Alice will describe f_x , and how Bob will determine some function B that is close to g_x based on Alice's description. For the first part, we let Alice and Bob use shared randomness in order to sample y_1, \dots, y_m , where the y_i 's are drawn independently with $y_i \sim \mu_Y$, and m is a parameter to be chosen later. Alice's description of f_x will then be $(f_x(y_1), \dots, f_x(y_m)) \in \{0, 1\}^m$. Thus, the length of the communication is m bits and we need to show that setting m to be roughly $O(k)$ suffices. Before we explain this, we first need to specify what Bob does with Alice's message.

As a first cut, let us consider the following natural strategy: Bob picks an $\tilde{x} \in X$ such that $g_{\tilde{x}}$ is close to f_x on y_1, \dots, y_m , and sets $B = g_{\tilde{x}}$. It is clear that if $\tilde{x} = x$, then $B = g_{\tilde{x}} = g_x$, and for every $y \in \mu_Y$, we would have $B(y) = g_x(y)$. Moreover, if \tilde{x} is such that $g_{\tilde{x}}$ is close to g_x (which is itself close to f_x), then $B(y)$ would now equal $g_x(y)$ with high probability. It remains to deal with \tilde{x} such that $g_{\tilde{x}}$ is far from g_x . Note that if we first fix any such \tilde{x} and then sample y_1, \dots, y_m , then with high probability, we would reveal that $g_{\tilde{x}}$ is far from g_x . This is because g_x is close to f_x , so $g_{\tilde{x}}$ should also be far from f_x . However, this idea alone cannot deal with all possible \tilde{x} — using a naive union bound over all possible $\tilde{x} \in X$ would require a failure probability of $1/|X|$, which would itself require setting m to be roughly $\log |X|$. Indeed, smaller values of m should not suffice since we have not yet used the fact that $\text{CC}_\epsilon^\mu(g) \leq k$ — but we do so next.

Suppose that Π is a one-way protocol with k bits of communication. Then, note that Alice's message partitions X into 2^k sets, one corresponding to each message. Our modified strategy for Bob is to let him pick a representative x from each set in this partition, and then set $B = g_{\tilde{x}}$ for an \tilde{x} among the representatives for which $g_{\tilde{x}}$ and f are the closest on the samples y_1, \dots, y_m . A simple analysis shows that the g_x 's that lie inside the same set in this partition are close, and thus, if we pick \tilde{x} to be the representative of the set containing x , then $g_{\tilde{x}}$ and f_x will be close on the sampled points. For an other representative, once again if $g_{\tilde{x}}$ is close to g_x , then $g_{\tilde{x}}(y)$ will equal $g_x(y)$ with high probability. For a representative \tilde{x} such that $g_{\tilde{x}}$ is far from g_x (which is itself close to f_x), we can proceed as in the previous paragraph, and now the union bound works out since the total number of representatives is only 2^k .³

We now turn to the case of general (not necessarily product) distributions. In this case, we would like to run the above protocol with y_1, y_2, \dots, y_m sampled independently from $\mu_{Y|x}$ (instead

³We note that a similar idea was used in a somewhat different context by [BJKS02] (following on [KNR99]) in order to characterize one-way communication complexity of any function under product distributions in terms of its VC-dimension.

of μ_Y). Note that Alice knows x and hence knows the distribution $\mu_{Y|x}$. Unfortunately, Bob does not know $\mu_{Y|x}$; he only knows μ_Y as a “proxy” for $\mu_{Y|x}$. While Alice and Bob cannot jointly sample such y_i ’s without communicating (as in the product case), they can still run the correlated sampling protocol of [BR11] in order to agree on such samples while communicating at most $O(m \cdot I(X; Y))$ bits. The original correlated sampling procedure of [BR11] inherently used multiple rounds of communication, but we are able in our case to turn it into a one-way protocol by leveraging the fact that our setup is distributional (see Subsection 4.2 for more details).

The outline of the rest of this section is the following. In Subsection 4.2, we describe the properties of the correlated sampling procedure that we will use. In Subsection 4.3, we give the formal proof of Theorem 1.1.

4.2 Correlated Sampling

We start by recalling two standard notions from information theory. Given two distributions P and Q , the *KL divergence* between P and Q is defined as $D(P||Q) \triangleq \mathbb{E}_{u \sim P}[\log(P(u)/Q(u))]$. Given a joint distribution μ of a pair (X, Y) of random variables with μ_X and μ_Y being the marginals of μ over X and Y respectively, the *mutual information* of X and Y is defined as $I(X; Y) \triangleq D(\mu||\mu_X\mu_Y)$.

The following lemma summarizes the properties of the *correlated sampling* protocol of [BR11].

Lemma 4.1 ([BR11]). *Let Alice be given a distribution P and Bob be given a distribution Q over a common universe \mathcal{U} . There is an interactive public-coin protocol that uses an expected*

$$D(P||Q) + 2 \log(1/\epsilon) + O(\sqrt{D(P||Q)} + 1)$$

bits of communication such that at the end of the protocol:

- *Alice outputs an element a distributed according to P .*
- *Bob outputs an element b such that for each $u \in \mathcal{U}$, $\Pr[b = u \mid a = u] > 1 - \epsilon$.*

Moreover, the message that Bob sends to Alice in any given round consists of a single bit indicating if the protocol should terminate or if Alice should send the next message.

We point out that in general, the correlated sampling procedure in Lemma 4.1 can take more than one round of communication. This is because initially, neither Alice nor Bob knows $D(P||Q)$ and they will need to interactively “discover” it. In our case, we will be using correlated sampling in a “distributional setup”. It turns out that this allows us to use a *one-way version* of correlated sampling which is described in Lemma 4.2 below.

Lemma 4.2. *Let μ be a distribution over (x, y) with marginal μ_X over x , and assume that μ is known to both Alice and Bob. Fix $\epsilon > 0$ and let Alice be given $x \sim \mu_X$. There is a one-way public-coin protocol that uses at most*

$$O(m \cdot I(X; Y)/\epsilon + \log(1/\epsilon)/\epsilon)$$

bits of communication such that with probability at least $1 - \epsilon$ over the public coins of the protocol and the randomness of x , Alice and Bob agree on m samples y_1, y_2, \dots, y_m i.i.d. $\sim \mu(Y|x)$ at the end of the protocol.

Proof. When x is Alice's input, we can consider running the protocol in Lemma 4.1 on the distributions $P \triangleq \prod_{i=1}^m \mu(Y_i|x)$ and $Q \triangleq \prod_{i=1}^m \mu(Y_i)$ and with error parameter $\epsilon/2$. Let Π be the resulting protocol transcript. The expected communication cost of Π is at most

$$\begin{aligned} \mathbb{E}_{x \sim \mu_X} [O(D(P||Q)) + O(\log(1/\epsilon))] &= O(\mathbb{E}_{x \sim \mu_X} [D(P||Q)]) + O(\log(1/\epsilon)) \\ &= O(m \cdot I(X; Y)) + O(\log(1/\epsilon)), \end{aligned} \quad (2)$$

where the last equality follows from the fact that

$$\begin{aligned} \mathbb{E}_{x \sim \mu_X} [D(P||Q)] &= \mathbb{E}_{x \sim \mu_X} \left[\mathbb{E}_{y_1|x, \dots, y_m|x} \left[\log \left(\frac{\prod_{i=1}^m \mu(y_i|x)}{\prod_{i=1}^m \mu(y_i)} \right) \right] \right] \\ &= \sum_{i=1}^m \mathbb{E}_{x \sim \mu_X} \left[\mathbb{E}_{y_1|x, \dots, y_m|x} \left[\log \left(\frac{\mu(y_i|x)}{\mu(y_i)} \right) \right] \right] \\ &= \sum_{i=1}^m \mathbb{E}_{x \sim \mu_X} \left[\mathbb{E}_{y_i|x} \left[\log \left(\frac{\mu(y_i|x)}{\mu(y_i)} \right) \right] \right] \\ &= \sum_{i=1}^m \mathbb{E}_{(x,y) \sim \mu} \left[\log \left(\frac{\mu(y|x)}{\mu(y)} \right) \right] \\ &= m \cdot I(X; Y). \end{aligned}$$

By Markov's inequality applied to (2), we get that with probability at least $1 - \epsilon/2$, the length of the transcript Π is at most

$$\ell \triangleq O(m \cdot I(X; Y)/\epsilon) + O(\log(1/\epsilon)/\epsilon).$$

Conditioned on the event E that the length of Π is at most ℓ bits, the total number of bits sent by Alice to Bob is also at most ℓ .

Note that Lemma 4.1 guarantees that each message of Bob in Π consists of a single bit indicating if the protocol should terminate or if Alice should send the next message. Hence, Bob's messages do not influence the actual bits sent by Alice; they only determine how many bits are sent by her.

In the new one-way protocol Π' , Alice sends to Bob, in a single shot, the first ℓ bits that she would have sent him in protocol Π if he kept refusing to terminate. Upon receiving this message, Bob completes the simulation of protocol Π . The error probability of the new protocol Π' is the probability that either Alice did not send enough bits or that the protocol Π makes an error, which by a union bound is at most

$$\Pr[\bar{E}] + \epsilon/2 \leq \epsilon/2 + \epsilon/2 = \epsilon$$

where \bar{E} denotes the complement of event E . □

4.3 Proof of Theorem 1.1

Recall that in the contextual setting, Alice's input is (f, x) and Bob's input is (g, y) , where $(f, g) \in \text{ow}\mathcal{F}_{k, \epsilon, \delta}^\mu$ and $(x, y) \sim \mu$. Let Π be the one-way protocol for g in the standard setting that shows that $\text{owCC}_\epsilon^\mu(g) \leq k$. Note that Π can be described by an integer $L \leq 2^k$ and functions $\pi: X \rightarrow [L]$ and $\{B_i: Y \rightarrow \{0, 1\}\}_{i \in [L]}$, such that Alice's message on input x is $\pi(x)$, and Bob's output on message i from Alice and on input y is $B_i(y)$. We use this notation below. We also set the parameter $m = \Theta((k + \log(1/\theta))/\theta^2)$, which is chosen such that $2^k \cdot \exp(-\theta^2 m/75) \leq 2\theta/5$.

The protocol. Algorithm 1 describes the protocol we employ in the contextual setting. Roughly speaking, the protocol works as follows. First, Alice and Bob run the one-way correlated sampling procedure given by Lemma 4.2 in order to sample y_1, y_2, \dots, y_m i.i.d. $\sim \mu_{Y|x}$. Then, Alice sends the sequence $(f_x(y_1), \dots, f_x(y_m))$ to Bob. Bob enumerates over $i \in [L]$ and counts the fraction of $z \in \{y_1, \dots, y_m\}$ for which $B_i(z) \neq f_x(z)$. For the index i which minimizes this fraction, Bob outputs $B_i(y)$ and halts.

Algorithm 1: The protocol that handles contextual uncertainty

The setting: Let μ be a probability distribution over a message space $X \times Y$. Alice and Bob are given functions f and g , and inputs x and y , respectively, where $(f, g) \in \text{ow}\mathcal{F}_{k, \epsilon, \delta}^\mu$ and $(x, y) \sim \mu$.

The protocol:

1. Alice and Bob run one-way correlated sampling with error parameter set to $(\theta/10)^2$ in order to sample m values $Z = \{y_1, y_2, \dots, y_m\} \subseteq Y$ each sampled independently according to $\mu_{Y|x}$.
 2. Alice sends $\{f_x(y_i)\}_{i \in [m]}$ to Bob.
 3. For every $i \in [L]$, Bob computes $\text{err}_i \triangleq \frac{1}{m} \sum_{j=1}^m \mathbb{1}(B_i(y_j) \neq f_x(y_j))$. Let $i_{\min} \triangleq \text{argmin}_{i \in [L]} \{\text{err}_i\}$. Bob outputs $B_{i_{\min}}(y)$ and halts.
-

Analysis. Observe that by Lemma 4.2, the correlated sampling procedure requires $O(m \cdot I(X; Y)/\theta^2 + \log(1/\theta)/\theta^2)$ bits of communication. Thus, the total communication of our protocol is at most

$$O(m \cdot I(X; Y)/\theta^2 + \log(1/\theta)/\theta^2) + m = \frac{c(k + \log(\frac{1}{\theta}))}{\theta^2} \cdot \left(1 + \frac{I(X; Y)}{\theta^2}\right) \text{ bits}$$

for some absolute constant c , as promised. The next lemma establishes the correctness of the protocol.

Lemma 4.3. $\Pr_{\Pi, (x, y) \sim \mu} [B_{i_{\min}}(y) \neq g(x, y)] \leq \epsilon + 2\delta + \theta$.

Proof. We start with some notation. For $x \in X$, let $\delta_x \triangleq \delta_{\mu_{Y|x}}(f_x, g_x)$ and let $\epsilon_x \triangleq \delta_{\mu_{Y|x}}(g_x, B_{\pi(x)})$. Note that by definition, $\delta = \mathbb{E}_{x \sim \mu_X}[\delta_x]$ and $\epsilon = \mathbb{E}_{x \sim \mu_X}[\epsilon_x]$. For $i \in [L]$, let $\gamma_{i, x} \triangleq \delta_{\mu_{Y|x}}(f_x, B_i)$. Note that by the triangle inequality,

$$\gamma_{\pi(x), x} = \delta_{\mu_{Y|x}}(f_x, B_{\pi(x)}) \leq \delta_x + \epsilon_x. \quad (3)$$

In what follows, we will analyze the probability that $B_{i_{\min}}(y) \neq g(x, y)$ by analyzing the estimate err_i and the index i_{\min} computed in the above protocol. Note that $\text{err}_i = \text{err}_i(x)$ computed above attempts to estimate $\gamma_{i, x}$, and that both err_i and i_{\min} are functions of x .

Note that Lemma 4.2 guarantees that correlated sampling succeeds with probability at least $1 - \theta^2/100$. Henceforth, we condition on the event that correlated sampling succeeds (we will account for the event where this does happen at the end). By the Chernoff bound, we have for every x and $i \in [L]$

$$\Pr_{y_1, \dots, y_m \sim \mu_{Y|x}} \left[|\gamma_{i, x} - \text{err}_i| > \frac{\theta}{5} \right] \leq \exp\left(-\frac{\theta^2 \cdot m}{75}\right).$$

By a union bound, we have for every $x \in X$,

$$\Pr_{y_1, \dots, y_m \sim \mu_{Y|x}} \left[\exists i \in [L] \text{ s.t. } |\gamma_{i,x} - \text{err}_i| > \frac{\theta}{5} \right] \leq L \cdot \exp\left(-\frac{\theta^2 \cdot m}{75}\right) \leq \frac{2\theta}{5},$$

where the last inequality follows from our choice of $m = \Theta(c \cdot (k + \log(1/\theta))/\theta^2)$.

Now assume that for all $i \in [L]$, we have that $|\gamma_{i,x} - \text{err}_i| \leq \theta/5$, which we refer to below as the ‘‘Good Event’’. Then, for i_{\min} , we have

$$\begin{aligned} \gamma_{i_{\min},x} &\leq \text{err}_{i_{\min}} + \theta/5 && \text{(since we assumed the Good Event)} \\ &\leq \text{err}_{\pi(x)} + \theta/5 && \text{(by definition of } i_{\min}\text{)} \\ &\leq \gamma_{\pi(x),x} + 2\theta/5 && \text{(since we assumed the Good Event)} \\ &\leq \delta_x + \epsilon_x + 2\theta/5. && \text{(By Equation (3))} \end{aligned}$$

Let $W \subseteq \{0, 1\}^n$ be the set of all x for which correlated sampling succeeds with probability at least $1 - \theta/10$ (over the internal randomness of the protocol). By Lemma 4.2 and an averaging argument, $\Pr_{x \sim \mu_X}[x \notin W] \leq \theta/10$. Thus,

$$\begin{aligned} \Pr_{\Pi, (x,y) \sim \mu} [B_{i_{\min}}(y) \neq f(x,y)] &\leq \mathbb{E}_{x \sim \mu_X | x \in W} \left[\Pr_{\Pi, y \sim \mu_{Y|x}} [B_{i_{\min}}(y) \neq f(x,y)] \right] + \theta/10 \\ &\leq \mathbb{E}_{x \sim \mu_X | x \in W} \left[\Pr_{y_1, \dots, y_m, y \sim \mu_{Y|x}} [B_{i_{\min}}(y) \neq f(x,y)] \right] + \theta/5 \\ &= \mathbb{E}_{x \sim \mu_X | x \in W} [\gamma_{i_{\min},x}] + \theta/5 \\ &\leq \mathbb{E}_{x \sim \mu_X | x \in W} [\delta_x + \epsilon_x] + 3\theta/5 \\ &\leq \mathbb{E}_{x \sim \mu_X} [\delta_x + \epsilon_x] + \theta \\ &= \delta + \epsilon + \theta \end{aligned}$$

where the third inequality follows from the fact that the Good Event occurs with probability at least $1 - 2\theta/5$, and from the corresponding upper bound on $\gamma_{i_{\min},x}$. The other inequalities above follow from the definition of the set W and the fact that $\Pr_{x \sim \mu_X}[x \notin W] \leq \theta/10$. Finally, since $\delta(f, g) \leq \delta$, we have that Bob’s output does not equal $g(x, y)$ (which is the desired output) with probability at most $\epsilon + 2\delta + \theta$. \square

5 Lower Bound for Non-Product Distributions

In this section, we prove Theorem 1.3. We start by defining the class of function pairs and distributions that will be used. Consider the parity functions on subsets of bits of the string $x \oplus y \in \{0, 1\}^n$. Specifically, for every $S \subseteq [n]$, let $f_S: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ be defined as $f_S(x, y) \triangleq \bigoplus_{i \in S} (x_i \oplus y_i)$. Let $q = q(n) > 0$ and define

$$\mathcal{F}_q \triangleq \{(f_S, f_T) : |S \Delta T| \leq q \cdot n\}. \quad (4)$$

Next, we define a probability distribution μ_p on $\{0, 1\}^n \times \{0, 1\}^n$ where $p = p(n)$. We do so by giving a procedure to sample according to μ_p . To sample a pair $(x, y) \sim \mu_p$, we draw $x \in_R \{0, 1\}^n$ and let y be a p -noisy copy of x , i.e., $y \sim N_p(x)$. Here, $N_p(x)$ is the distribution on $\{0, 1\}^n$ that

outputs $y \in \{0, 1\}^n$ such that, independently, for each $i \in [n]$, $y_i = 1 - x_i$ with probability p , and $y_i = x_i$ with probability $1 - p$. In other words, $\mu_p(x, y) = 2^{-n} \cdot p^{|x \oplus y|} \cdot (1 - p)^{n - |x \oplus y|}$ for every $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$. Here, $|z|$ denotes the Hamming weight of z , for any $z \in \{0, 1\}^n$.

We will prove Lemmas 5.1 and 5.2 below about the function class \mathcal{F}_q and the distribution μ_p .

Lemma 5.1. *For any $p = p(n)$ and $q = q(n)$, it holds that $\mathcal{F}_q \subseteq \text{ow}\mathcal{F}_{1,0,pqn}^{\mu_p}$.*

In words, Lemma 5.1 says that any pair of functions in \mathcal{F}_q are (pqn) -close in δ_{μ_p} -distance, and any function in \mathcal{F}_q has a one-way zero-error protocol with a single bit of communication. Lemma 5.2 lower bounds the contextually uncertain communication complexity of \mathcal{F}_q under distribution μ_p .

Lemma 5.2. *For any $p = p(n)$, $q = q(n)$ and $\epsilon > 0$, it holds that:*

$$\text{CCU}_{\frac{1}{2}-\epsilon}^{\mu_p}(\mathcal{F}_q) \geq \gamma \cdot \min\{p \cdot n, (q/2) \cdot n\} - \log(1/\epsilon) + \eta,$$

where $\eta = 2^{-\Theta(q \cdot n)}/\epsilon$, and $\gamma > 0$ is an absolute constant.

Note that applying Lemmas 5.1 and 5.2 with $\mathcal{F} = \mathcal{F}_q$, $\mu = \mu_p$ and $p = q = \sqrt{\delta/n}$ (where $\delta > 0$ is any constant) implies Theorem 1.3.

In Subsection 5.1 below, we prove Lemma 5.1 which follows from two simple propositions. The main part of the rest of this section is dedicated to the proof of Lemma 5.2. The idea behind the proof of Lemma 5.2 is to reduce the problem of computing \mathcal{F}_q under μ_p with contextual uncertainty, into the problem of computing a related function in the standard distributional communication complexity model (i.e., *without* uncertainty) under a related distribution. We then use the *discrepancy method* to prove a lower bound on the communication complexity of the new problem. This task itself reduces to upper bounding the spectral norm of a certain communication matrix. The choice of our underlying distribution then implies a tensor structure for this matrix, which reduces the spectral norm computation to bounding the largest singular value of an explicit family of 4×4 matrices.

We point out that our lower bound in Lemma 5.2 is essentially tight up to a logarithmic factor. Namely, one can show using a simple one-way hashing protocol that for any constant $\epsilon > 0$, $\text{owCCU}_\epsilon^{\mu_p}(\mathcal{F}_q) \leq O(r \log r)$ with $r \triangleq \min\{2p \cdot n, q \cdot n\}$.

5.1 Proof of Lemma 5.1

Lemma 5.1 follows from Propositions 5.3 and 5.4 below. We first show that every pair of functions in \mathcal{F}_q are close under the distribution μ_p .

Proposition 5.3. *For every $(f, g) \in \mathcal{F}_q$, it holds that $\delta_{\mu_p}(f, g) \leq pqn$.*

Proof. Any pair of functions $(f, g) \in \mathcal{F}_q$ is of the form $f = f_S$ and $g = f_T$ with $|S \Delta T| \leq q$. Hence,

$$\Pr_{(x,y) \sim \mu} [f(x, y) \neq g(x, y)] = \Pr_{(x,y) \sim \mu} [f_{S \Delta T}(x \oplus y) = 1] \leq 1 - (1 - p)^{|S \Delta T|} \leq 1 - (1 - p)^{qn} \leq pqn.$$

□

Next, we show that there is a simple one-way communication protocol that allows Alice and Bob to compute f_S (for any $S \subseteq [n]$) with just a single bit of communication.

Proposition 5.4. $\text{owCC}(f_S) = 1$.

Proof. Recall that $f_S(x, y) = \bigoplus_{i \in S} (x_i \oplus y_i)$. We write this as $f_S(x, y) = (\bigoplus_{i \in S} x_i) \oplus (\bigoplus_{i \in S} y_i)$. This leads to the simple one-way protocol where Alice computes $b = \bigoplus_{i \in S} x_i$ and sends the single bit result of the computation to Bob. Bob can now compute $b \oplus (\bigoplus_{i \in S} y_i) = f_S(x, y)$ to obtain the value of f_S (with zero error). \square

5.2 Proof of Lemma 5.2

In order to lower bound $\text{CCU}_{\frac{1}{2}-\epsilon}^{\mu_p}(\mathcal{F}_q)$, we define a communication problem in the standard distributional complexity setting that can be reduced to the problem of computing \mathcal{F}_q under contextual uncertainty. The lower bound in Lemma 5.2 is then obtained by proving a lower bound on the communication complexity of the new problem which is defined as follows:

- **Inputs:** Alice's input is a pair (S, x) where $S \subseteq [n]$ and $x \in \{0, 1\}^n$. Bob's input is a pair (T, y) such that $T \subseteq [n]$ and $y \in \{0, 1\}^n$.
- **Distribution:** Let \mathcal{D}_q be a distribution on pairs of Boolean functions (f, g) on $\{0, 1\}^n \times \{0, 1\}^n$ defined by the following sampling procedure. To sample $(f, g) \sim \mathcal{D}_q$, we pick a set $S \subseteq [n]$ uniformly at random and set $f = f_S$. We then pick T to be a $(q/2)$ -noisy copy of S and set $g = f_T$. The distribution on the inputs of Alice and Bob is then described by $\nu_{p,q} = \mathcal{D}_q \times \mu_p$: we sample $(x, y) \sim \mu_p$ and sample $(f, g) \sim \mathcal{D}_q$.
- **Function:** The goal is to compute the function F given by $F((S, x), (T, y)) \triangleq f_T(x \oplus y)$.

Proposition 5.5 below – which follows from a simple Chernoff bound – shows that a protocol computing \mathcal{F}_q under μ_p can also be used to compute the function F in the standard distributional model with $((S, x), (T, y)) \sim \nu_{p,q}$, and with the same amount of communication.

Proposition 5.5. *For every $\epsilon > 0$, it holds that $\text{CCU}_{\frac{1}{2}-\epsilon}^{\mu_p}(\mathcal{F}_q) \geq \text{CC}_{\frac{1}{2}-\epsilon+\epsilon'}^{\nu_{p,q}}(F)$ with $\epsilon' = 2^{-\Theta(q \cdot n)}$.*

In the rest of this section, we will prove the following lower bound on $\text{CC}_{\frac{1}{2}-\epsilon}^{\nu_{p,q}}(F)$, which along with Proposition 5.5, implies Lemma 5.2:

Lemma 5.6. *For every $\epsilon > 0$, it holds that*

$$\text{CC}_{\frac{1}{2}-\epsilon}^{\nu_{p,q}}(F) \geq \gamma \cdot \min\{p \cdot n, (q/2) \cdot n\} - \log(1/\epsilon),$$

where $\gamma > 0$ is an absolute constant.

To prove Lemma 5.6, without loss of generality, we will set $q = 2p$ and prove a lower bound of $\gamma \cdot p \cdot n$ on the communication complexity⁴. So henceforth, we denote $\nu_p \triangleq \nu_{p,2p}$. The proof will use the *discrepancy bound* which is a well-known method for proving lower bounds on distributional communication complexity in the standard model.

⁴We can do so because $\text{CC}_{\frac{1}{2}-\epsilon}^{\nu_{p,q}}(F) \geq \text{CC}_{\frac{1}{2}-\epsilon}^{\nu_{r,2r}}(F)$ with $r \triangleq \min(p, q/2)$, which follows from the fact that Alice can always use her private randomness to reduce the correlation between either (x, y) or (S, T) .

Definition 5.7 (Discrepancy; [KN97]). Let F and ν_p be as above and let R be any rectangle (i.e., any set of the form $R = C \times D$ where $C, D \subseteq \{0, 1\}^{2n}$). Denote

$$\text{DISC}_{\nu_p}(R, F) \triangleq \left| \Pr_{\nu_p} [F((S, x), (T, y)) = 0, ((S, x), (T, y)) \in R] - \Pr_{\nu_p} [F((S, x), (T, y)) = 1, ((S, x), (T, y)) \in R] \right|.$$

The discrepancy of F according to ν_p is $\text{DISC}_{\nu_p}(F) \triangleq \max_R \text{DISC}_{\nu_p}(R, F)$ where the maximum is over all rectangles R .

The next known proposition relates distributional communication complexity to discrepancy.

Proposition 5.8 ([KN97]). For any $\epsilon > 0$, it holds that $\text{CC}_{\frac{1}{2}-\epsilon}^{\nu_p}(F) \geq \log(2\epsilon/\text{DISC}_{\nu_p}(F))$.

We will prove the following lemma.

Lemma 5.9. $\text{DISC}_{\nu_p}(F) \leq 2^{-\gamma \cdot p \cdot n}$ for some absolute constant $\gamma > 0$.

Note that Lemma 5.9 and Proposition 5.8 put together immediately imply Lemma 5.6. The proof of Lemma 5.9 uses some standard facts about the spectral properties of matrices and their tensor powers that we next recall. Let $A \in \mathbb{R}^{d \times d}$ be a real square matrix. Then, $v \in \mathbb{R}^d$ is said to be an eigenvector of A with eigenvalue $\lambda \in \mathbb{R}$ if $Av = \lambda v$. If A is furthermore (symmetric) positive semi-definite, then all its eigenvalues are real and non-negative. We can now define the spectral norm of a (not necessarily symmetric) matrix.

Definition 5.10. The spectral norm of a matrix $A \in \mathbb{R}^{d \times d}$ is given by $\|A\| \triangleq \sqrt{\lambda_{\max}(A^T A)}$, where $\lambda_{\max}(A^T A)$ is the largest eigenvalue of $A^T A$.

Also, recall that given a matrix $A \in \mathbb{R}^{d \times d}$ and a positive integer t , the tensor power matrix $A^{\otimes t} \in \mathbb{R}^{d^t \times d^t}$ is defined by $(A^{\otimes t})_{(i_1, \dots, i_t)} = \prod_{j=1}^t A_{i_j}$ for every $(i_1, \dots, i_t) \in \{1, \dots, d\}^t$. We will use the following standard fact which in particular says that the spectral norm is multiplicative with respect to tensoring.

Fact 5.11. For any matrix $A \in \mathbb{R}^{d \times d}$, vector $u \in \mathbb{R}^d$, scalar $c \in \mathbb{R}$ and positive integer t , we have

1. $\|cA\| = |c| \cdot \|A\|$.
2. $\|A^{\otimes t}\| = \|A\|^t$.
3. $\|Au\|_2 \leq \|A\| \cdot \|u\|_2$, where for any vector $w \in \mathbb{R}^d$, $\|w\|_2$ denotes the Euclidean norm of w , i.e., $\|w\|_2 \triangleq \sqrt{\sum_{i=1}^d w_i^2}$.

The next lemma upper bounds the spectral norm of an explicit family of 4×4 matrices that will be used in the proof of Lemma 5.9. Looking ahead, it is crucial for our purposes that the coefficient of a in the right-hand side of Equation (5) is a constant strictly smaller than 2.

Lemma 5.12. Let $a \in (0, 1)$ be a real number and $N \triangleq N(a) \triangleq \begin{bmatrix} 1 & a & a & -a^2 \\ a & 1 & -a^2 & a \\ a & a^2 & 1 & -a \\ a^2 & a & -a & 1 \end{bmatrix}$. Then,

$$\|N\|_2 \leq 1 + \sqrt{2} \cdot a + a^2 + \frac{a^4}{2} + \frac{a^5}{\sqrt{2}}. \quad (5)$$

Proof. One can verify that

$$N^T N = \begin{bmatrix} (a^2 + 1)^2 & 2a(a^2 + 1) & 2a(1 - a^2) & 0 \\ 2a(a^2 + 1) & (a^2 + 1)^2 & 0 & 2a(1 - a^2) \\ 2a(1 - a^2) & 0 & (a^2 + 1)^2 & -2a(a^2 + 1) \\ 0 & 2a(1 - a^2) & -2a(a^2 + 1) & (a^2 + 1)^2 \end{bmatrix}.$$

Assuming that $a \in (0, 1)$, one can also verify that $N^T N$ has as eigenvectors

$$v_1 \triangleq \begin{bmatrix} \frac{\sqrt{2(a^4+1)}}{1-a^2} \\ \frac{a^2+1}{1-a^2} \\ 1 \\ 0 \end{bmatrix}, \quad v_2 \triangleq \begin{bmatrix} \frac{a^2+1}{1-a^2} \\ \frac{\sqrt{2(a^4+1)}}{1-a^2} \\ 0 \\ 1 \end{bmatrix} \quad \text{with eigenvalue } \lambda_1(a) \triangleq 2a^2 + a^4 + 2a\sqrt{2(a^4 + 1)} + 1,$$

$$\text{and } v_3 \triangleq \begin{bmatrix} \frac{\sqrt{2(a^4+1)}}{a^2-1} \\ \frac{a^2+1}{1-a^2} \\ 1 \\ 0 \end{bmatrix}, \quad v_4 \triangleq \begin{bmatrix} \frac{a^2+1}{1-a^2} \\ \frac{\sqrt{2(a^4+1)}}{a^2-1} \\ 0 \\ 1 \end{bmatrix} \quad \text{with eigenvalue } \lambda_2(a) := 2a^2 + a^4 - 2a\sqrt{2(a^4 + 1)} + 1.$$

Note that for any value of $a \in (0, 1)$, the vectors v_1, v_2, v_3 and v_4 are linearly independent and each of the eigenvalues $\lambda_1(a)$ and $\lambda_2(a)$ has multiplicity 2. Moreover, we have that $\lambda_1(a) \geq \lambda_2(a)$. Hence,

$$\|N\| = \sqrt{\lambda_1(a)} = \sqrt{2a^2 + a^4 + 2a\sqrt{2(a^4 + 1)} + 1}.$$

Applying twice the fact that $\sqrt{1+x} \leq 1 + x/2$ for any $x \geq -1$, we get that

$$\begin{aligned} \|N\| &= \sqrt{1 + 2a^2 + a^4 + 2a\sqrt{2}\sqrt{1+a^4}} \\ &\leq 1 + a^2 + \frac{a^4}{2} + a\sqrt{2}\sqrt{1+a^4} \\ &\leq 1 + a^2 + \frac{a^4}{2} + a\sqrt{2}\left(1 + \frac{a^4}{2}\right) \\ &= 1 + a\sqrt{2} + a^2 + \frac{a^4}{2} + \frac{a^5}{\sqrt{2}}. \end{aligned}$$

□

We are now ready to prove Lemma 5.9.

Proof of Lemma 5.9: Fix any rectangle $R = C \times D$ where $C, D \subseteq \{0, 1\}^{2n}$. We wish to show that $\text{DISC}_{\nu_p}(R, F) \leq 2^{-\gamma \cdot p \cdot n}$. First, note that $\text{DISC}_{\nu_p}(R, F) = |1_C M 1_D|$ where 1_C and 1_D are the 0/1 indicator vectors of C and D (respectively) and M is the $2^{2n} \times 2^{2n}$ real matrix defined by⁵

$$\begin{aligned} M_{((S,x),(T,y))} &\triangleq \nu_p((S, T), (x, y)) \cdot (-1)^{f_T(x \oplus y)} \\ &= \frac{1}{2^{2n}} (1-p)^{2n} (-1)^{\langle T, x \oplus y \rangle} \left(\frac{p}{1-p}\right)^{|S \oplus T| + |x \oplus y|} \end{aligned}$$

⁵We here use the symbols S and T to denote both subsets of $[n]$ and the corresponding 0/1 indicator vectors.

for every $S, x, T, y \in \{0, 1\}^n$. Letting $a \triangleq p/(1-p)$, we can write

$$M_{((S,x),(T,y))} = \frac{1}{2^{2n}} (1-p)^{2n} (N^{\otimes n})_{((S,x),(T,y))}$$

with $N = N(a)$ being the 4×4 real matrix defined by⁶

$$N_{((S_1,x_1),(T_1,y_1))} \triangleq (-1)^{T_1(x_1 \oplus y_1)} a^{|S_1 \oplus T_1| + |x_1 \oplus y_1|} \quad (6)$$

for all $S_1, x_1, T_1, y_1 \in \{0, 1\}$. Using the third property listed in Fact 5.11, we get

$$\text{Disc}_{\nu_p}(R, F) = |1_C M 1_D| \leq \|1_C\|_2 \cdot \|M\| \cdot \|1_D\|_2 \leq \sqrt{2^{2n}} \cdot \|M\| \cdot \sqrt{2^{2n}} = 2^{2n} \cdot \|M\| \quad (7)$$

We now use the first two properties listed in Fact 5.11 to relate $\|M\|$ to $\|N\|$ as follows:

$$\|M\| = \left\| \frac{1}{2^{2n}} (1-p)^{2n} N^{\otimes n} \right\| = \frac{1}{2^{2n}} (1-p)^{2n} \|N\|^n. \quad (8)$$

Using Equation (6), we can check that

$$N = N(a) = \begin{bmatrix} 1 & a & a & -a^2 \\ a & 1 & -a^2 & a \\ a & a^2 & 1 & -a \\ a^2 & a & -a & 1 \end{bmatrix}.$$

Applying Lemma 5.12 with $a = p/(1-p)$ and p sufficiently small (e.g., less than 1/10), we get

$$\|N\| \leq 1 + \sqrt{2} \cdot \left(\frac{p}{1-p}\right) + O(p^2). \quad (9)$$

Combining Equations (7) to (9) above, we conclude that

$$\begin{aligned} \text{Disc}_{\nu_p}(R, F) &\leq (1-p)^{2n} \cdot \left(1 + \sqrt{2} \cdot \left(\frac{p}{1-p}\right) + O(p^2)\right)^n \\ &= \left[(1-p) \cdot \left(1 + p \cdot (\sqrt{2} - 1) + O(p^2)\right) \right]^n \\ &= \left[1 - p \cdot (2 - \sqrt{2}) + O(p^2) \right]^n \\ &\leq 2^{-\gamma \cdot p \cdot n} \end{aligned}$$

for some absolute constant $\gamma > 0$. □

⁶In Equation (6), $T_1(x_1 \oplus y_1)$ denotes the product of the bit T_1 and the bit $(x_1 \oplus y_1)$. Moreover, since $(S_1 \oplus T_1)$ is a single bit, its Hamming weight $|S_1 \oplus T_1|$ is the same as its bit-value, and similarly for $(x_1 \oplus y_1)$.

6 Conclusion and Future Directions

In this work, we introduced and studied a simple model illustrating the role of context in communication and the challenge posed by uncertainty of knowledge of context.

On the technical side, it would be interesting to determine the correct exponent of $I(X; Y)$ in Theorem 1.1. Theorems 1.1 and 1.3 imply that this exponent is between $1/2$ and 1 .

It would also be interesting to prove an analogue of Theorem 1.1 for two-way protocols. Our proof of Theorem 1.1 uses in particular the fact that any low-communication one-way protocol in the standard distributional communication model should have a certain canonical form: to compute $g(x, y)$, Alice tries to describe the entire function $g(x, \cdot)$ to Bob, and this does not create a huge overhead in communication. Coming up with a canonical form of two-way protocols that somehow changes gradually as we morph from g to f seems to be the essence of the challenge in extending Theorem 1.1 to the two-way setting.

On the more conceptual side, arguably, the model considered in this work is a fairly realistic one: communication has some goals in mind which we model by letting Bob be interested in a specific function of the joint information that Alice and Bob possess. Moreover, it is a fairly natural model to posit that the two are not in perfect synchronization about the function that Bob is interested in, but Alice can estimate the function in some sense. One aspect of our model that can be further refined is the specific notion of distance that quantifies the gap between Bob’s function and Alice’s estimate. In this work, we chose the Hamming distance which forms a good first starting point. We believe that it is interesting to propose and study other models of distance between functions that more accurately capture natural forms of uncertainty.

Finally, we wish to emphasize the mix of adversarial and probabilistic elements in our uncertainty model — the adversary picks (f, g) whereas the inputs (x, y) are picked from a distribution. We believe that richer mixtures of adversarial and probabilistic elements could lead to broader settings of modeling and coping with uncertainty — the probabilistic elements offer efficient possibilities that are often immediately ruled out by adversarial choices, whereas the adversarial elements prevent the probabilistic assumptions from being too precise.

References

- [BJKS02] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. Information theory methods in communication complexity. In *17th Annual IEEE Conference on Computational Complexity*, pages 93–102, 2002. 9
- [BM11] Andrej Bogdanov and Elchanan Mossel. On extracting common random bits from correlated sources. *IEEE Transactions on Information Theory*, 57(10):6351–6355, 2011. 7
- [BR11] Mark Braverman and Anup Rao. Information equals amortized communication. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 748–757. IEEE, 2011. 2, 10
- [CGMS15] Clément Louis Canonne, Venkatesan Guruswami, Raghu Meka, and Madhu Sudan. Communication with imperfectly shared randomness. In *Innovations in Theoretical Computer Science, ITCS*, pages 257–262, 2015. 1, 3, 7

- [GJS12] Oded Goldreich, Brendan Juba, and Madhu Sudan. A theory of goal-oriented communication. *J. ACM*, 59(2):8, 2012. 1, 3
- [HS14] Elad Haramaty and Madhu Sudan. Deterministic compression with uncertain priors. In *Innovations in Theoretical Computer Science, ITCS*, pages 377–386, 2014. 1, 3
- [JKKS11] Brendan Juba, Adam Tauman Kalai, Sanjeev Khanna, and Madhu Sudan. Compression without a common prior: an information-theoretic justification for ambiguity in language. In *Innovations in Computer Science, ICS*, pages 79–86, 2011. 1, 3
- [JS08] Brendan Juba and Madhu Sudan. Universal semantic communication I. In *40th Annual ACM Symposium on Theory of Computing*, pages 123–132, 2008. 1, 3
- [JS11] Brendan Juba and Madhu Sudan. Efficient semantic communication via compatible beliefs. In *Innovations in Computer Science, ICS*, pages 22–31, 2011. 1
- [JW13] Brendan Juba and Ryan Williams. Massive online teaching to bounded learners. In *Innovations in Theoretical Computer Science, ITCS*, pages 1–10, 2013. 1
- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997. 4, 16
- [KNR99] Ilan Kremer, Noam Nisan, and Dana Ron. On randomized one-round communication complexity. *Computational Complexity*, 8(1):21–49, 1999. 9
- [Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *11th Annual ACM Symposium on Theory of Computing*, pages 209–213, 1979. 1, 4