

New Limits to Classical and Quantum Instance Compression

(Preliminary version)

Andrew Drucker*

April 13, 2012

Last updated: August 21, 2012

Abstract

Given an instance of a hard decision problem, a limited goal is to *compress* that instance into a smaller, equivalent instance of a second problem. As one example, consider the problem where, given Boolean formulas ψ^1, \dots, ψ^t , we must determine if at least one ψ^j is satisfiable. An *OR-compression scheme* for SAT is a polynomial-time reduction R that maps (ψ^1, \dots, ψ^t) to a string z , such that z lies in some “target” language L' if and only if $\bigvee_j [\psi^j \in \text{SAT}]$ holds. (Here, L' can be arbitrarily complex.) AND-compression schemes are defined similarly. A compression scheme is *strong* if $|z|$ is polynomially bounded in $n = \max_j |\psi^j|$, independent of t .

Strong compression for SAT seems unlikely. Work of Harnik and Naor (FOCS '06/SICOMP '10) and Bodlaender, Downey, Fellows, and Hermelin (ICALP '08/JCSS '09) showed that the infeasibility of strong OR-compression for SAT would show limits to instance compression for a large number of natural problems. Bodlaender et al. also showed that the infeasibility of strong AND-compression for SAT would have consequences for a different list of problems. Motivated by this, Fortnow and Santhanam (STOC '08/JCSS '11) showed that if SAT is strongly OR-compressible, then $\text{NP} \subseteq \text{coNP}/\text{poly}$. Finding similar evidence against AND-compression was left as an open question.

We provide such evidence: we show that strong AND- *or* OR-compression for SAT would imply *non-uniform, statistical zero-knowledge proofs* for SAT—an even stronger and more unlikely consequence than $\text{NP} \subseteq \text{coNP}/\text{poly}$. Our method applies against *probabilistic* compression schemes of sufficient “quality” with respect to the reliability and compression amount (allowing for tradeoff). This greatly strengthens the evidence given by Fortnow and Santhanam against probabilistic OR-compression for SAT. We also give variants of these results for the analogous task of *quantum instance compression*, in which a polynomial-time quantum reduction must output a quantum state that, in an appropriate sense, “preserves the answer” to the input instance.

The central idea in our proofs is to exploit the information bottleneck in an AND-compression scheme for a language L in order to fool a cheating prover in a proof system for \bar{L} . Our key technical tool is a new method to “disguise” information being fed into a compressive mapping; we believe this method may find other applications.

*MIT CSAIL. Email: andy.drucker@gmail.com.

Contents

1	Introduction	3
1.1	Instance compression and parametrized problems	3
1.2	Previous work: results and motivation	4
1.3	Our results	5
1.3.1	Results on classical compression	5
1.3.2	Results on quantum compression	8
1.4	Our techniques	9
1.4.1	The overall approach	9
1.4.2	The Disguising-Distribution Lemma	11
1.4.3	Extension to the quantum case	13
1.5	Organization of the paper	13
2	Preliminaries I	13
2.1	Statistical distance and distinguishability	14
3	Proof of Theorem 1.1	15
4	Preliminaries II	19
4.1	Information theory background	19
4.2	Basic complexity classes and promise problems	21
4.3	Arthur-Merlin protocols	21
4.4	Statistical zero-knowledge and the SD problem	22
4.5	f -compression reductions	24
5	Parametrized problems and parametrized compression	24
5.1	Parametrized problems	25
5.2	OR-expressive and AND-expressive parametrized problems	26
5.3	Parametrized compression	27
5.4	Connecting parametrized compression and f -compression	28
6	Technical lemmas	29
6.1	Distributional stability	29
6.2	Sparsified distributional stability	31
6.3	Building disguising distributions	33
7	Limits to efficient (classical) compression	35
7.1	Complexity upper bounds from OR-compression schemes	35
7.2	Application to AND- and OR-compression of NP-complete languages	38
7.3	f -compression of NP-complete languages for general f	40
7.4	Limits to strong compression for parametrized problems	40
7.5	Application to problems with polynomial kernelizations	41

8	Extension to quantum compression	43
8.1	Trace distance and distinguishability of quantum states	44
8.2	Quantum f -compression	45
8.3	Quantum complexity classes	45
8.4	Quantum distributional stability	47
8.5	Building quantum disguising distributions	47
8.6	Complexity upper bounds from quantum compression schemes	48
9	Questions for further study	49
A	Alternative proofs of distributional stability	53
A.1	A proof based on Raz’s lemma	53
A.2	A proof based on the Average Encoding Theorem	54
B	Our original distributional stability lemma	55
B.1	Entropy and the unreliability of compressive encodings	55
B.2	Bounds on the inverse entropy function	57
B.3	The lemma	58
C	Proof of quantum distributional stability	59

1 Introduction

1.1 Instance compression and parametrized problems

Given an instance of a hard decision problem, we may hope to *compress* that instance into a smaller, equivalent instance, either of the same or of a different decision problem. Here we do not ask to be able to recover the original instance from the smaller instance; we only require that the new instance have the same (yes/no) *answer* as the original. Such *instance compression* may be the first step towards obtaining a solution; this has been a central technique in the theory of *fixed-parameter-tractable* algorithms [DF99, GN07]. Strong compression schemes for certain problems would also have important implications for cryptography [HN10]. Finally, compressing an instance of a difficult problem may also be a worthwhile goal in its own right, since it can make the instance easier to store and communicate [HN10].

It is unknown whether one can efficiently, significantly compress an arbitrary instance of a natural NP-complete language like SAT, the set of satisfiable Boolean formulas.¹ A more limited goal is to design an efficient reduction that achieves compression on instances that are particularly “simple” in some respect. To explore this idea, one needs a formal model defining “simple” instances; the versatile framework of *parametrized problems* [DF99] is one such model, and has been extensively used to study instance compression. A parametrized problem is a decision problem in which every instance has an associated *parameter value* k , giving some measure of the complexity of a problem instance.² As an example, one can parametrize a Boolean formula ψ by the number

¹If we could efficiently reduce instances of some NP-complete problem to shorter instances of the *same* problem, then we could iterate the reduction to solve our problem in polynomial time, implying $P = NP$. However, even if $P \neq NP$, it is still conceivable that SAT might have an efficient compressive reduction to a different target problem—to the Halting problem, say.

²See Section 5.1 for details. The parameter k is explicitly given as part of the input to the algorithm.

of distinct variables appearing in ψ .

An ambitious goal for a parametrized problem P is to compress an arbitrary instance x of the decision problem for P into an equivalent instance x' of a second, “target” decision problem, where the output length $|x'|$ is bounded by a *polynomial* in $k = k(x)$. If P has such a reduction running in time $\text{poly}(|x| + k)$, we say P is *strongly compressible*; we say P is *strongly self-compressible* if the target problem of the reduction is P itself. (In the literature of parametrized problems, a strong self-compression reduction is usually referred to as a *polynomial kernelization*. More generally, a *kernelization* is a polynomial-time self-compression reduction whose output size is bounded by *some* function of the parameter k alone.)

1.2 Previous work: results and motivation

Let VAR-SAT denote the Satisfiability problem for Boolean formulas, parametrized by the number of distinct variables in the formula. In their study of instance compression for NP-hard problems, Harnik and Naor [HN10] asked whether VAR-SAT is strongly compressible.³ They showed that a positive answer would have several significant consequences for cryptography. Notably, they proved that a *deterministic* strong compression reduction for VAR-SAT (with any target problem) would yield a construction of collision-resistant hash functions based on any one-way function—a long-sought goal.

In fact, Harnik and Naor showed that for their applications, it would suffice to achieve strong compression for a simpler parametrized problem, the “OR(SAT) problem:” this is the Satisfiability problem for Boolean formulas expressed as disjunctions $\psi = \bigvee_{j=1}^t \psi_j$, where the parameter is now defined as the maximum bit-length of any sub-formula ψ_j . Strong compression for VAR-SAT easily implies strong compression for OR(SAT). Harnik and Naor defined a hierarchy of decision problems called the “VC hierarchy,” which can be modeled as a class of parametrized problems (see [FS11]). They showed that a strong compression reduction for any of the problems “above” OR(SAT) in this hierarchy would also imply strong compression for OR(SAT); this includes parametrized versions of natural problems like the Clique and Dominating Set problems. While Harnik and Naor’s primary motivation was to *find* a strong compression scheme for OR(SAT) to use in their cryptographic applications, their work also provides a basis for showing *negative* results: in view of the reductions in [HN10], any evidence against strong compression for OR(SAT) is also evidence against strong compression for a variety of other parametrized problems.

In subsequent, independent work, Bodlaender, Downey, Fellows, and Hermelin [BDFH09] also studied the compressibility of OR(SAT) and of related problems; these authors’ motivations came from the theory of *fixed-parameter tractable (FPT)* algorithms [DF99]. An FPT algorithm for a parametrized problem P is an algorithm that solves an arbitrary instance x , with parameter $k = k(x)$, in time $g(k) \cdot \text{poly}(|x| + k)$, for some function $g(\cdot)$. The idea is that even if P is hard in general, an FPT algorithm for P may be practical on instances where the parameter k is small. Now as long as P is decidable, a kernelization reduction for P provides the basis for an FPT algorithm for P : on input x , first compress x , then solve the equivalent, compressed instance. The kernelization approach is one of the most widely-used schemas for developing FPT algorithms.⁴ Of course, one hopes to compress by as large an amount as possible, to maximize the efficiency of the resulting FPT algorithm; this motivates the search for *strong* self-compression reductions.

³Strictly speaking, they asked a slightly different question whose equivalence to this one was pointed out in [FS11].

⁴In fact, *every* problem with an FPT algorithm is kernelizable [CCDF97]. This does not mean, however, that the most efficient FPT algorithms always arise from the kernelization approach.

Strong self-compression reductions are known for parametrized versions of many natural NP-complete problems, such as the Vertex Cover problem; see, e.g., the survey [GN07]. However, for many other such parametrized problems, including numerous problems known to admit FPT algorithms (such as OR(SAT)), no strong compression reduction is known, to any target problem. Bodlaender et al. [BDFH09] conjectured that no strong self-compression reduction exists for OR(SAT). They made a similar conjecture for the closely-related “AND(SAT) problem,” in which one is given Boolean formulas ψ_1, \dots, ψ_t and asked to decide whether $\bigwedge_{j=1}^t [\psi_j \in \text{SAT}]$ holds—that is, whether every ψ_j is individually satisfiable. As with OR(SAT), we parametrize AND(SAT) by the maximum bit-length of any ψ_j .

Bodlaender et al. showed that these conjectures (sometimes referred to as the “OR-” and “AND-conjectures”) would have considerable explanatory power. First, they showed [BDFH09, Theorem 1] that the nonexistence of strong self-compression reductions for OR(SAT) would rule out strong self-compression for a large number of other natural parametrized problems; these belong to a class we call “OR-expressive problems.”⁵ Under the assumption that AND(SAT) does not have strong self-compression, Bodlaender et al. ruled out strong self-compression reductions for a second substantial list of problems [BDFH09, Theorem 2], belonging to a class we will call “AND-expressive.” Despite the apparent similarity of OR(SAT) and AND(SAT), no equivalence between the compression tasks for these two problems is known.

In light of their results, Bodlaender et al. asked for complexity-theoretic evidence against strong self-compression for OR(SAT) and AND(SAT). Fortnow and Santhanam [FS11] provided the first such evidence: they showed that if OR(SAT) has a strong compression reduction (to any target problem), then $\text{NP} \subseteq \text{coNP/poly}$ and the Polynomial Hierarchy collapses to its third level.

The techniques of [BDFH09, FS11] were refined and extended by many researchers to give further evidence against efficient compression for parametrized problems, e.g., in [DLS09, DvM10, BTY11, BJK11a, BJK11b, BJK11c, CFM11, HW12, DM12, Kra12]. (See [DM12] for further discussion and references.) As one notable development that is relevant to our work, Dell and Van Melkebeek [DvM10] combined the techniques of [BDFH09, FS11] with new ideas to provide tight compression-size lower bounds for certain problems that *do* admit polynomial kernelizations. Researchers also used ideas from [BDFH09, FS11] in other areas of complexity, giving new evidence of lower bounds for the length of PCPs [FS11, DvM10] and for the density of NP-hard sets [BH08].

Finding evidence against strong compression for AND(SAT) was left as an open question by these works, however. The limits of *probabilistic* compression schemes for OR(SAT) and for OR-expressive problems (including VAR-SAT) also remained unclear. The results and techniques of [FS11] give evidence only against some restrictive sub-classes of probabilistic compression schemes for OR(SAT): schemes with one-sided error, avoiding false negatives; schemes whose error probability is exponentially small in the length of the entire input; and schemes using $O(\log n)$ random bits, where $n = \max_j |\psi_j|$.

1.3 Our results

1.3.1 Results on classical compression

We complement the results of [FS11] by providing evidence against strong compression for AND(SAT): we prove that such a compression scheme, to any target problem, would also imply $\text{NP} \subseteq \text{coNP/poly}$.

⁵See Section 5.2. The class of OR-expressive problems is not identical to the class described in [BDFH09], but it is closely related and contains their class, as well as other classes of problems identified in [HN10, BJK11a].

In fact, we show that reductions compressing even by a much more modest amount would imply the same conclusion. For concreteness, we state our most “basic” result on compression of AND(SAT) in a self-contained way below.

Theorem 1.1. *Let L be any NP-complete language. Suppose there is a deterministic polynomial-time reduction R that takes an arbitrarily long list of input strings (x^1, \dots, x^t) and outputs a string z , with*

$$z \in L \iff \bigwedge_{j \in [t]} [x^j \in L] .$$

Suppose further that R obeys the output-size bound $|z| \leq (\max_{j \leq t} |x^j|)^{O(1)}$, with the polynomial bound independent of t . Then, $\text{NP} \subseteq \text{coNP/poly}$.

More strongly, we show the following. Suppose there is any second, “target” language L' , a pair of polynomially-bounded functions $t(n), t'(n) : \mathbb{N} \rightarrow \mathbb{N}$ with $t(n) = \omega(1)$ and $t'(n) + 1 \leq t(n)/2$, and a deterministic polynomial-time reduction $R : \{0, 1\}^{t(n) \times n} \rightarrow \{0, 1\}^{\leq t'(n)}$, such that

$$R(x^1, \dots, x^{t(n)}) \in L' \iff \bigwedge_{j \in [t(n)]} [x^j \in L] .$$

Then $\text{NP} \subseteq \text{coNP/poly}$.

We prove Theorem 1.1 in Section 3. In later sections, we will strengthen and generalize Theorem 1.1 using related but more powerful proof techniques. However, we feel it is worthwhile to present a proof of this basic result with a minimum of tools and preliminaries.⁶

The techniques we use to generalize Theorem 1.1 will extend naturally (and in a strong fashion) to the *probabilistic* setting with two-sided error, in which we expect the compression reduction to obey some success-probability guarantee on every input. We show (in Theorem 7.4, item 1) that any sufficiently “high-quality” compression scheme for AND(SAT) would imply $\text{NP} \subseteq \text{coNP/poly}$. Here, “quality” is defined by a certain relationship between the reliability and the compression amount of the reduction, and allows for tradeoff.

We also show (in Theorem 7.4, item 2, and Theorem 7.5) that beyond a second, somewhat more demanding quality threshold, probabilistic compression reductions either for AND(SAT) or for OR(SAT) would imply the existence of *non-uniform, statistical zero-knowledge proofs* for NP languages—a stronger (and even more unlikely) consequence than $\text{NP} \subseteq \text{coNP/poly}$. The more-demanding quality threshold in this second set of results is still rather modest, and allows us to prove the following result as a special case:

Theorem 1.2 (Informal). *Suppose that either of AND(SAT) or OR(SAT) is strongly compressible, with success probability $\geq .5 + 1/\text{poly}(n)$ for an AND or OR of length- n formulas. Then there are non-uniform, statistical zero-knowledge proofs for all languages in NP.*

At the other extreme, where we consider compression schemes with more modest compression amounts, but with greater reliability, our techniques yield the following result:

⁶To be precise, in our elementary proof we avoid any overt use of information-theoretic results and concepts; we also avoid the use of the minimax theorem. These tools are central to our stronger and more general approach (which, in particular, is much better suited for analyzing bounded-error reductions), but familiarity with these tools is not necessary to understand Theorem 1.1. We mention that the decision to use or avoid information theory in the proof is essentially independent of the choice to use or avoid the minimax theorem.

Theorem 1.3 (Informal). *Let $t(n) : \mathbb{N}^+ \rightarrow \mathbb{N}^+$ be any polynomially bounded function. Suppose there is a compression scheme compressing an AND of $t(n)$ length- n SAT instances into an instance z of a second decision problem L' , where $|z| \leq C \cdot t(n) \log t(n)$ for some $C > 0$. If the scheme’s error probability on such inputs is bounded by a sufficiently small inverse-polynomial in n (depending on $t(n)$ and C), then there are non-uniform, statistical zero-knowledge proofs for all languages in NP. The corresponding result also holds for OR-compression.*⁷

Our results give the first strong evidence of hardness for compression of AND(SAT). They also greatly strengthen the evidence given by Fortnow and Santhanam against *probabilistic* compression for OR(SAT), and provide the first strong evidence against probabilistic compression for the potentially-harder problem VAR-SAT. For *deterministic* (or error-free) compression of OR(SAT), the limits established by our techniques also follow from the techniques of [FS11], which apply given an OR-compression scheme with compression bound of form $|z| \leq O(t(n) \log t(n))$.⁸ On the other hand, we provide somewhat stronger complexity-theoretic evidence for these limits to compression.

Using our results on the infeasibility of compression for AND(SAT) and OR(SAT), and building on [HN10, BDFH09, FS11], we give new complexity-theoretic evidence against strong compressibility for a list of interesting parametrized problems with FPT algorithms. (See Theorem 7.7.) This is the first strong evidence against strong compressibility for any of the ten “AND-expressive” problems identified in [BDFH09] (and listed in Section 5.2). For the numerous “OR-expressive” problems identified in [HN10, BDFH09] and other works, this strengthens the negative evidence given by [FS11].

Our methods also extend the known results on limits to compression for parametrized problems that do possess polynomial kernelizations: we can partially extend the results of Dell and Van Melkebeek [DvM10] to the case of probabilistic algorithms with two-sided error. For example, for $d > 1$ and any $\varepsilon > 0$, Dell and Van Melkebeek proved that if the Satisfiability problem for N -variable d -CNFs has a polynomial-time compression reduction with output-size bound $O(N^{d-\varepsilon})$, then $\text{NP} \subseteq \text{coNP/poly}$. Their result applies to co-nondeterministic reductions, and to probabilistic reductions without false negatives; we prove (in Theorem 7.11) that the result also holds for probabilistic reductions with two-sided error, as long as the success probability of the reduction is at least $.5 + N^{-\beta}$ for some $\beta = \beta(d, \varepsilon) > 0$. Using reductions described in [DvM10], we also obtain quantitatively-sharp limits to probabilistic compression for several other natural NP-complete problems, including the Vertex Cover and Clique problems on graphs and hypergraphs. (However, the limits we establish do not give lower bounds on the cost of *oracle communication protocols*; these protocols are a generalization of compression reductions, studied in [DvM10], to which that work’s results do apply. Trying to extend our results to this model seems like an interesting challenge for further study.)

Our results about AND(SAT) and OR(SAT) follow from more general results about arbitrary languages. For any language L , we follow previous authors and consider the “OR(L) problem,” in which one is given a collection x^1, \dots, x^t of strings, and is asked to determine whether at least one of them is a member of L . We show (in Theorem 7.1, item 1) that if a sufficiently “high-quality” probabilistic polynomial-time compression reduction exists for the OR(L) problem, then

⁷In fact, *error-free* OR-compression of this sort for SAT would give non-uniform *perfect* zero-knowledge proofs for NP, and error-free AND-compression for SAT would give non-uniform perfect zero-knowledge proofs for coNP; see Theorem 7.3.

⁸This is not explicitly shown in [FS11], but follows from the technique of [FS11, Theorem 3.1]; see also [DvM10, Lemma 3] for a more general result that makes the achievable bounds clear.

$L \in \text{NP/poly}$. (As before, “high-quality” is defined by a relation between the reliability of the reduction and the compression amount.) We also show (in Theorem 7.1, item 2) that a polynomial-time compression scheme for $\text{OR}(L)$ meeting a more demanding standard of quality implies that L possesses non-uniform statistical zero-knowledge proof systems, and lies in $\text{NP/poly} \cap \text{coNP/poly}$. (For deterministic compression, the conclusion $L \in \text{coNP/poly}$ was established earlier in [FS11].) Applying these results to $L := \overline{\text{SAT}}$ gives our hardness-of-compression results for $\text{AND}(\text{SAT})$; applying the second set of results to $L := \text{SAT}$ gives our improved negative results for $\text{OR}(\text{SAT})$.

In unpublished work, Buhrman [Buh] constructed an oracle A such that, for every NP^A -complete language L , the decision problem $\text{AND}(L)$ does not have a P^A -computable strong compression reduction. This gave earlier, indirect evidence against efficient strong compression for the $\text{AND}(\text{SAT})$ problem—or at least, it indicated that exhibiting such a compression reduction would require novel techniques. Now, inspection of the proofs reveals that our new results on compression for $\text{OR}(L)$ are all perfectly relativizing. This allows us to identify many more oracles obeying the property of Buhrman’s oracle: namely, we may take any A for which $\text{NP}^A \not\subseteq \text{coNP}^A/\text{poly}$. For example, this holds with probability 1 for a *random oracle* [BG81].⁹ Such an oracle can also be obtained through a simple diagonalization argument.

For any Boolean function $f : \{0, 1\}^* \rightarrow \{0, 1\}$, we may generalize the $\text{OR}(L)$ decision problem to the problem $f \circ L$, in which one is given a collection of strings x^1, \dots, x^t and must output $f(L(x^1), \dots, L(x^t))$. So far it would seem that our negative results are fairly specific to the case where the outer “combining function” f is either AND or OR . However, by an idea suggested in [FS11, Section 7], our negative result on compression for $\text{AND}(\text{SAT})$, combined with Fortnow and Santhanam’s negative results on compression for $\text{OR}(\text{SAT})$, actually implies the following: for *any* Boolean function f that depends on all of its input bits for each input length, no strong compression schemes exist from $f \circ \text{SAT}$ to a target language $L' \in \text{NP}$, unless $\text{NP} \subseteq \text{coNP/poly}$. (See Theorem 7.6.) Note the new requirement in this result that L' be in NP ; if the combining function f is monotone, this requirement may be dropped. The quantitative bounds we obtain on the achievable compression amount are somewhat weaker for general f than for $f \in \{\text{AND}, \text{OR}\}$, however; developing a better understanding of the situation for other combining functions could be another interesting goal for future work.

1.3.2 Results on quantum compression

Up to this point, we have discussed compression reductions in which the input and output are both “classical” bit-strings. However, from the perspective of quantum computing and quantum information [NC00], it is natural to ask about the power of compression reductions that output a *quantum state*. An “ n -qubit state” is a quantum superposition over classical n -bit strings; a vast body of research has explored the extent to which information can be succinctly encoded within and retrieved from such quantum states. If quantum computers become a practical reality, quantum instance compression schemes could help to store and transmit hard computational problems; compressing an instance might also be a first step towards its solution by a quantum algorithm.

We propose the following quantum generalization of classical instance compression: a *quantum compression reduction* for a language L is a quantum algorithm that, on input x , outputs a quantum state ρ on some number q of qubits—hopefully with $q \ll |x|$, to achieve significant compression. Our

⁹In [BG81] it is shown that $\text{NP}^A \not\subseteq \text{coNP}^A$ for random A ; the technique readily extends to give the stronger claim above.

correctness requirement is that there should exist *some* quantum measurement \mathcal{M}_q , depending only on q , such that for every x compressing to q qubits, $\mathcal{M}_q(\rho) = L(x)$ holds with high probability over the inherent randomness in the measurement $\mathcal{M}_q(\rho)$. We do not require that \mathcal{M}_q be an efficiently-performable measurement; this is by analogy to the general version of the classical compression task, in which the target language of the reduction may be arbitrarily complex.

Our results for quantum compression are closely analogous to our results in the classical case. First, we show that for any language L , if a sufficiently “high-quality” quantum polynomial-time compression reduction exists for the $\text{OR}(L)$ problem, then L possesses a *non-uniform, 2-message quantum interactive proof system* (with a single prover). Second, we show that a sufficiently higher-quality quantum polynomial-time compression reduction for $\text{OR}(L)$ implies that L possesses a *non-uniform quantum statistical zero-knowledge proof system*. Remarkably, the two “quality thresholds” in our quantum results are essentially *the same* as in the corresponding results for the classical case.¹⁰ It follows that, unless there exist surprisingly powerful quantum proofs of unsatisfiability for Boolean formulas, the limits we establish for probabilistic compression of $\text{AND}(\text{SAT})$ and $\text{OR}(\text{SAT})$ hold just as strongly for quantum compression.¹¹

1.4 Our techniques

In this section we will focus on describing our strongest and most general techniques. As mentioned earlier, we also present a similar, but more “elementary” approach to prove Theorem 1.1. We will give some self-contained intuition about that approach in Section 3. That strategy bears some similarities to work of Fortnow and Santhanam [FS11] on the hardness of compression for $\text{OR}(\text{SAT})$. In particular, it shares an *incremental* approach to defining non-uniform advice for a proof system; in each case, the stage-based construction makes progress in correctly classifying more and more strings of a given input length.

1.4.1 The overall approach

We first describe our techniques for the classical case; these form the basis for the quantum case as well. Our first two general results, giving complexity upper bounds on any language L for which $\text{OR}(L)$ has a sufficiently high-quality compression reduction (Theorem 7.1, items 1 and 2), are both based on a single reduction that we describe next. This reduction applies to compression reductions mapping some number $t(n) \leq \text{poly}(n)$ of inputs of length n to an output string z of length $|z| = O(t(n) \log t(n))$.

Fix any language L such that $\text{OR}(L)$ has a possibly-probabilistic compression reduction

$$R(x^1, \dots, x^t) : \{0, 1\}^{t \times n} \longrightarrow \{0, 1\}^{\leq t'}$$

with some target language L' , along with parameters t', t satisfying $t' \leq O(t \log t) \leq \text{poly}(n)$.¹² We will use R to derive upper bounds on the complexity of L . (The reader may keep in mind the main

¹⁰We do place a minor additional restriction on quantum compression reductions for $\text{OR}(L)$: we require that the reduction, on input (x^1, \dots, x^t) , outputs a quantum state of size determined by $(\max_j |x^j|)$ and t .

¹¹We remark that *3-message* quantum interactive proofs are known to be fully as powerful as quantum interactive proofs in which polynomially many messages are exchanged [Wat03], and that these proof systems are equal in power to PSPACE in the uniform setting [JJUW11]. However, *2-message* quantum proof systems seem much weaker, and are not known to contain coNP .

¹²Here we pay exclusive attention to R 's behavior on tuples of strings of some equal length n .

intended setting $L = \overline{\text{SAT}}$, which we will use to derive our hardness results for the compression of AND(SAT). No special properties of this language will be used in the argument, however.)

A simple, motivating observation is that if we take a string $y \in L$ and “insert” it into a tuple $\bar{x} = (x^1, \dots, x^t)$ of elements of \bar{L} , replacing some x^j to yield a modified tuple \bar{x}' , then the values

$$R(\bar{x}), R(\bar{x}')$$

are *different* with high probability—for, by the “OR-respecting” property of R , we will with high probability have $R(\bar{x}) \in \bar{L}'$, $R(\bar{x}') \in L'$. More generally, for any *distribution* \mathcal{D} over t -tuples of inputs from \bar{L} , let $\mathcal{D}[y, j]$ denote the distribution obtained by sampling $\bar{x} \sim \mathcal{D}$ and replacing x^j with y ; then the two output distributions

$$R(\mathcal{D}), R(\mathcal{D}[y, j])$$

are *far apart* in statistical distance. (Of course, the strength of the statistical-distance lower bound we get will depend on the reliability of our compression scheme.)

We want this property to serve as the basis for an interactive proof system by which a computationally powerful Prover can convince a skeptical polynomial-time (but non-uniform) Verifier that a string y lies in L . The idea for our initial, randomized protocol (which we will later derandomize) is that Prover will make his case by demonstrating his ability to *distinguish* between the two R -output distributions described above, when Verifier privately chooses one of the two distributions, samples from it, and sends the sample to Prover.¹³ But then to make our proof system meaningful, Verifier also needs to *fool* a cheating Prover in the case $y \notin L$. To do this, we want to choose \mathcal{D}, j in such a way that the distributions $R(\mathcal{D}), R(\mathcal{D}[y, j])$ are as close as possible whenever $y \notin L$.

We may not be able to achieve this for an index j that is poorly-chosen. For instance, $R(\bar{x})$ may always copy the first component x^1 as part of the output string z , so taking $j = 1$ would fail badly. To get around this, we choose our replacement index j *uniformly at random*, aiming in this way to make R “insensitive” to the insertion of y .¹⁴ As R is a compression scheme, it doesn’t have room in its output string to replicate its entire input, so there is reason for hope.

This invites us to search for a distribution \mathcal{D}^* over $(\bar{L}_n)^t$ with the following properties:

- (i) For every $y \in \bar{L}_n$, if we select $j \in [t]$ uniformly then the expected statistical distance $\mathbb{E}_j [||R(\mathcal{D}^*) - R(\mathcal{D}^*[y, j])||_{\text{stat}}]$ is “not too large;”¹⁵
- (ii) \mathcal{D}^* is efficiently sampleable, given non-uniform advice of length $\text{poly}(n)$.

Condition (i) is quite demanding: we need a single distribution \mathcal{D}^* rendering R insensitive to the insertion of *any* string $y \in \bar{L}_n$ —a set which may be of exponential size. Condition (ii) is also a strong restriction: \bar{L}_n may be a complicated set, and in general we can only hope to sample from

¹³Interactive proofs based on distinguishing tasks have seen many uses in theoretical computer science, and indeed we will rely upon known protocols of this kind in our work; see Section 4.4.

¹⁴We emphasize that the “insensitivity” we are looking for is *statistical*; we are not asking that y have small effect on the output of R for most *particular* outcomes to $\bar{x} \sim \mathcal{D}$. This latter goal may not be achievable, e.g., if R outputs the sum of all its input strings x^i taken as vectors over \mathbb{F}_2^n .

¹⁵For our purposes, it actually suffices to bound $||R(\mathcal{D}^*) - R(\mathcal{D}^*[y, \mathbf{j}])||_{\text{stat}}$, where \mathbf{j} is a uniform value sampled “internally” as part of the distribution. In our streamlined proof of Theorem 1.1, we will use this idea. However, our techniques will yield the stronger property in condition (i) above, and this is the course we will follow in proving our general results.

distributions over $(\bar{L}_n)^t$ in which t -tuples are formed out of a fixed “stockpile” of $\text{poly}(n)$ elements of \bar{L}_n , hard-coded into the non-uniform advice.

Remarkably, it turns out that such a distribution \mathcal{D}^* can always be found. In fact, in item (i), we can force the two distributions to be non-negligibly close (with expected statistical distance $\leq 1 - \frac{1}{\text{poly}(n)}$) whenever the output-size bound t' obeyed by R is $O(t \log t)$; the distributions will be much closer when $t' \ll t$. We call our key technical result (Lemma 6.6), guaranteeing the existence of such a \mathcal{D}^* , the “*Disguising-Distribution Lemma*.”

Assuming this lemma for the moment, we use \mathcal{D}^* as above to reduce any membership claim for L to a distinguishing task for a Prover-Verifier protocol. Given any input y , we’ve constructed two distributions $\mathfrak{R} = R(\mathcal{D}^*)$ and $\mathfrak{R}' = R(\mathcal{D}^*[y, \mathbf{j}])$ (with \mathbf{j} uniform), where each distribution is sampleable in non-uniform polynomial time. Our analysis guarantees some lower bound $D = D(n)$ on $\|\mathfrak{R} - \mathfrak{R}'\|_{\text{stat}}$ in the case $y \in L$, and some upper bound $d = d(N)$ on this distance when $y \notin L$. (These parameters depend on the reliability and compression guarantees of R .) If $D(n) - d(n) \geq \frac{1}{\text{poly}(n)}$, we can give non-uniform distinguishing protocols for L , which can be converted to public-coin protocols and then non-uniformly derandomized to show that $L \in \text{NP}/\text{poly}$. Also, if $D(n)^2 - d(n) \geq \frac{1}{\text{poly}(n)}$ then, using a powerful result due to Sahai and Vadhan [SV03], we can derive a non-uniform, statistical zero-knowledge proof system for L . This also implies $L \in \text{NP}/\text{poly} \cap \text{coNP}/\text{poly}$.

1.4.2 The Disguising-Distribution Lemma

The Disguising-Distribution Lemma, informally described in Section 1.4.1, is a statement about the behavior of $R(x^1, \dots, x^t)$ on a specified product subset S^t of inputs ($S = \bar{L}_n$ in our application). This lemma is a “generic” result about the behavior of compressive mappings; it uses no properties of R other than R ’s output-size bound.¹⁶ In view of its generality and interest, we are hopeful that the lemma will find other applications.

Our proof of this lemma uses two central ideas. First, we interpret the search for the “disguising distribution” \mathcal{D}^* as a two-player game between a “disguising player” (choosing \mathcal{D}^*) and an opponent who chooses y ; we can then apply simple yet powerful principles of game theory. Second, to build a winning strategy for the disguising player, we will exploit an information bottleneck in R stemming from its compressive property.¹⁷

To describe the proof, it is helpful to first understand how one may obtain the distribution \mathcal{D}^* if we drop the efficient-sampleability requirement on \mathcal{D}^* , and focus on the “disguising” requirement (condition (i)). To build \mathcal{D}^* in this relaxed setting, we will appeal to the *minimax theorem* for two-player, zero-sum games; applied here, it tells us that to guarantee the existence of a \mathcal{D}^* that succeeds in disguising all strings $y \in \bar{L}_n$, it is enough to show how to build a \mathcal{D}_Y^* that succeeds in expectation, when y is sampled from some fixed (but arbitrary) distribution Y over \bar{L}_n .

Here, a natural idea springs to mind: let \mathcal{D}_Y^* just be a product distribution over t copies of Y ! In this case, inserting $y \sim Y$ into \mathcal{D}_Y^* at a random location is equivalent to conditioning on the outcome of a randomly-chosen coordinate of a sample from \mathcal{D}_Y^* . The intuition here is that, due to the output-size bound on R , the distribution $R(\mathcal{D}_Y^*)$ shouldn’t have enough “degrees of freedom” to be affected much by this conditioning.

¹⁶Indeed, in our application we have essentially no control on R ’s behavior when we consider its restriction to inputs from S^t , so a generic result is needed.

¹⁷This is hardly the first paper in which such a bottleneck plays a crucial and somewhat unexpected role. For example, an interesting and slightly similar application of information-theoretic tools to the study of *metric embeddings* was found recently by Regev [Reg11].

We show (in Lemma B.4, with similar, alternative results presented in Appendices A and B) that for any product distribution $\bar{x} \sim (\mathcal{D}_1, \dots, \mathcal{D}_t)$ over t -tuple inputs to R , conditioning on the value of $x^j \sim \mathcal{D}_j$ for a uniformly-chosen index $j \in [t]$ has bounded expected effect on the output distribution $R(\bar{x})$. That is, the *expected statistical distance* between the pre- and post-conditioned distributions is bounded non-negligibly away from 1 (provided that $t' \leq O(t \log t)$). We refer to this important property of R as “*distributional stability*.”

In our original proof that our compressive mapping R is distributionally stable, we gave a simple (non-constructive) way to use R as a one-shot *encoding method* for independent, unbiased bits b_1, \dots, b_t . The encoding Enc has a desirable property: for each component $j \in [t]$ whose expected “influence” on the output distribution of R is noticeable (when we fix a single value $x^j \sim \mathcal{D}_j$), our encoding transmits b_j with noticeable advantage over a random guess. We can then deduce strong upper bounds on the influence of a typical component j , using the output-size bound on R and elementary information-theoretic bounds on the reliability of compressive encodings. This analysis succeeds when $t' \leq t - 2$. In our original draft, we used more elaborate techniques (which involved modifying the mapping R itself) to analyze the case when $t \leq t' \leq O(t \log t)$.

Several researchers pointed out to us that the distributional stability property can be established in a different way, using Kullback-Leibler divergence and an inequality due to Pinsker (see Theorem 4.7). This approach allows us to analyze the case when $t \leq t' \leq (1 + \varepsilon)t$, for a modest $\varepsilon > 0$. As this author noted later, the divergence-based approach can be combined with an alternative to Pinsker’s inequality—a bound due to Vajda (Theorem 4.8; see [FHT03, RW09] for more information on both of these inequalities)—to show that the mapping R has a non-negligible amount of distributional stability as long as $t' \leq O(t \log t)$. Thus we feel that the divergence-based approach is ultimately the most convenient one to work with in general; this is the approach we now use in the main body of the paper.

Colleagues additionally helped us to understand that the distributional stability property for mappings with $t' \leq (1 + \varepsilon)t$ can also be established using other similar, known results that follow from the same divergence/Pinsker-based techniques: a lemma of Raz [Raz98], and the “Average Encoding Theorem” of Klauck et al. [KNTSZ07]. The latter was used in [KNTSZ07] to identify a stability property for trace and Hellinger distance metrics, for the inputs to a problem in quantum communication complexity; this was used for a round-elimination argument. Their proof is for inputs drawn from the uniform distribution, but extends readily to general distributions and can be used to derive the kind of lemma we need. We describe these alternative proofs of distributional stability in Appendix A,¹⁸ and we describe our own original, encoding/decoding-based approach in Appendix B. We feel that all of these approaches to proving distributional stability are interesting and worth understanding.

Using the distributional-stability property of compressive mappings under product input-distributions, we then establish a certain “sparsified variant” of this property (Lemma 6.3), which allows us to replace each \mathcal{D}_j with a small set sampled from \mathcal{D}_j ;¹⁹ this is an important tool in addressing the efficient-sampleability requirement on our desired \mathcal{D}^* . Using this variant, we use the minimax theorem to show (in Lemma 6.4) that there exists a *distribution* \mathfrak{D} over product input-distributions to R —with each product distribution defined over small subsets of S —such that, in expectation,

¹⁸Russell Impagliazzo suggested the use of Raz’s lemma; Salil Vadhan also helped me to understand the connection. Ashwin Nayak and S. Vadhan suggested direct proofs of distributional stability based on divergence and Pinsker’s inequality, which we now use as our main approach. Dieter van Melkebeek also suggested the relevance of Pinsker’s inequality. I thank all of these researchers.

¹⁹For convenience in the proof, we assume $\mathcal{D}_j = \mathcal{D}_{j'}$ for all j, j' .

\mathfrak{D} disguises the random insertion of any string $y \in S$ at a uniformly-chosen position j . Finally, in Lemma 6.6 we obtain our desired “disguising distribution” \mathcal{D}^* as a sparsified version of \mathfrak{D} , using a result due to Lipton and Young [LY94] and, independently, to Althöfer [Alt94], that guarantees the existence of sparsely-supported, nearly-optimal strategies in 2-player, zero-sum games.

1.4.3 Extension to the quantum case

Our techniques for studying quantum compression are closely analogous to the classical case. The main technical difference is that the output $R(\mathcal{D})$ of our compression reduction, on any input distribution \mathcal{D} , is now a (mixed) *quantum state*. In this setting, to carry out an analogue of the argument sketched in Sections 1.4.1 and 1.4.2, we need a “disguising distribution” for R that meets a modified version of condition (i) from Section 1.4.1:

- (i') For every $y \in \bar{L}_n$, if we select $j \in [t]$ uniformly then, for any quantum measurement \mathcal{M} , the expected statistical distance $\mathbb{E}_j [|\mathcal{M}(R(\mathcal{D}^*)) - \mathcal{M}(R(\mathcal{D}^*[y, j]))|_{\text{stat}}]$ is not too large.

A basic measure of distance between quantum states, the *trace distance*, is relevant here: if two states ρ, ρ' are at trace distance $\|\rho - \rho'\|_{\text{tr}} \leq \delta$, then for any measurement \mathcal{M} , the statistical distance $|\mathcal{M}(\rho) - \mathcal{M}(\rho')|_{\text{stat}}$ is at most δ . (In fact, this property *characterizes* the trace distance.) Thus to satisfy condition (i'), it will be enough to construct \mathcal{D}^* so as to upper-bound $\mathbb{E}_j [|\mathcal{M}(R(\mathcal{D}^*)) - \mathcal{M}(R(\mathcal{D}^*[y, j]))|_{\text{tr}}]$, for uniformly-chosen j . We do this by essentially the same techniques as in the classical case. The one significant difference is that here, we need to establish a “stability property” for trace distance, analogous to the stability property for statistical distance described in Section 1.4.2. This can be obtained using the same basic divergence-based techniques as in the classical case, with the help of suitable tools from quantum information theory.²⁰

1.5 Organization of the paper

In Section 2, we present the “bare minimum” of preliminaries needed to understand our proof of Theorem 1.1. We present this proof in Section 3.

The rest of the paper is devoted to proving stronger and more general results. In Section 4, we give the additional needed preliminary material for our work, including our definitions of compression reductions. In Section 5, we formally introduce parametrized problems and AND- and OR-expressive problems. In Section 6, we prove the main technical lemmas we use to obtain our results on limits of efficient instance compression (with alternative proofs of the first such lemma appearing in Appendices A and B). Our results for the classical setting are proved in Section 7, and the quantum results are proved in Section 8. Finally, in Section 9 we present questions for future study.

2 Preliminaries I

Definition 2.1. The binary entropy function $H(\alpha) : [0, 1] \rightarrow [0, 1]$ is defined by

$$H(\alpha) := -\alpha \log_2 \alpha - (1 - \alpha) \log_2(1 - \alpha)$$

²⁰Our original approach to proving distributional stability also admits a quantum version, although we no longer present it here.

on $(0, 1)$, with $H(0) = H(1) := 0$.

$\binom{n}{k}$ denotes the binomial coefficient $n!k!/(n-k)!$. We will use the following standard, simple bound (see, e.g., [vL99, Chapter 1]) on the number of binary vectors of low Hamming weight:

Fact 2.2. *For $t \in \mathbb{N}$ and $\alpha \in (0, .5)$, we have*

$$\sum_{0 \leq \ell \leq \alpha t} \binom{t}{\ell} \leq 2^{H(\alpha)t}.$$

2.1 Statistical distance and distinguishability

All distributions in this paper will take finitely many values; let $\text{supp}(\mathcal{D})$ be the set of values assumed by \mathcal{D} and let $\mathcal{D}(u) := \Pr[\mathcal{D} = u]$.

For a probability distribution \mathcal{D} and $t \geq 1$, we let $\mathcal{D}^{\otimes t}$ denote a t -tuple of outputs sampled independently from \mathcal{D} . We let \mathcal{U}_K denote the uniform distribution over a multiset K .

The *statistical distance* of two distributions $\mathcal{D}, \mathcal{D}'$ over a shared universe of outcomes is defined as

$$\|\mathcal{D} - \mathcal{D}'\|_{\text{stat}} := \frac{1}{2} \sum_{u \in \text{supp}(\mathcal{D}) \cup \text{supp}(\mathcal{D}')} |\mathcal{D}(u) - \mathcal{D}'(u)|.$$

We will use the following familiar “distinguishability interpretation” of the statistical distance. Suppose a value $b \in \{0, 1\}$ is selected uniformly, unknown to us, and a sample $u \in U$ is drawn from \mathcal{D} if $b = 0$, or from \mathcal{D}' if $b = 1$. We observe u , and our goal is to correctly guess b . It is a basic fact that, for *any* $\mathcal{D}, \mathcal{D}'$, our maximum achievable success probability in this “distinguishing” experiment is precisely $\frac{1}{2}(1 + \|\mathcal{D} - \mathcal{D}'\|_{\text{stat}})$. Furthermore, the optimal distinguishing algorithm may without loss of generality be a deterministic “maximum-likelihood” rule $ML(b|u)$: guess “ $b = 1$ ” if and only if $\Pr[b = 1|u] \geq 1/2$. Similarly, we define a maximum-likelihood rule $ML(X|Y)$ for guessing any random variable X based on the observed value of any other random variable Y : simply guess the likeliest value of X conditioned on the observation (breaking ties arbitrarily).

The following fact follows from the distinguishability characterization of $\|\cdot\|_{\text{stat}}$; it is a convenient weakening of that principle.

Fact 2.3. *If X, Y are random variables over some shared domain S , and $\Delta := \|X - Y\|_{\text{stat}}$, then there exists a subset $T \subseteq S$ such that*

$$\Pr_{x \sim X}[x \in T] \geq \Delta \quad \text{and} \quad \Pr_{y \sim Y}[y \notin T] \geq \Delta.$$

We will also use the following facts:

Fact 2.4. *If X, Y are random variables over some shared domain S , and $R(X)$ is any (possibly randomized) function taking inputs from S , then*

$$\|R(X) - R(Y)\|_{\text{stat}} \leq \|X - Y\|_{\text{stat}}.$$

Fact 2.5 ([SV03], Fact 2.3). *Suppose (X_1, X_2, Y_1, Y_2) are distributions on a shared probability space Ω , that X_1 is independent of X_2 , and that Y_1 is independent of Y_2 . Then,*

$$\|(X_1, X_2) - (Y_1, Y_2)\|_{\text{stat}} \leq \|X_1 - Y_1\|_{\text{stat}} + \|X_2 - Y_2\|_{\text{stat}}.$$

3 Proof of Theorem 1.1

This section presents a proof that the “AND-conjecture” of Bodlaender, Downey, Fellows, and Hermelin [BDFH09] holds true unless $\text{NP} \subseteq \text{coNP}/\text{poly}$. As discussed earlier, in this section we aim for a proof that avoids information theory and the minimax theorem. In later sections we will obtain stronger and more general results with these tools.

It will be convenient to consider mappings $R : \{0, 1\}^{t \times n} \rightarrow \{0, 1\}^{\leq t'}$, for fixed n, t, t' . For $A \subseteq \{0, 1\}^n$, let \mathbf{R}_A denote the distribution $\mathbf{R}_A := R(\mathcal{U}_A^{\otimes t})$, and for each $a \in \{0, 1\}^n$, define the distribution

$$\mathbf{R}_A[a] := R(\mathcal{U}_A^{\otimes(j-1)}, a, \mathcal{U}_A^{\otimes(t-j)}) ,$$

where $\mathbf{j} \sim \mathcal{U}_{[t]}$.

Define the *standout factor*

$$\beta(a, A) := \|\mathbf{R}_A[a] - \mathbf{R}_A\|_{\text{stat}} . \tag{1}$$

The basic idea of our proof of Theorem 1.1 is as follows: we will show that for each $n > 0$, there exists a $\text{poly}(n)$ -size collection of $\text{poly}(n)$ -size sets $A_i \subseteq L_n$,²¹ such that every other element $x \in L_n$ will have standout factor $\beta(x, A_i) < 1 - \Omega(1)$ for at least one A_i . On the other hand, each $x \notin L_n$ will have standout factor 1 against each A_i .²² Thus, if a polynomial-time Verifier “quizzes” a Prover by randomly sampling, either from \mathbf{R}_{A_i} or from $\mathbf{R}_{A_i}[x]$ on each i , then Prover will be able to reliably guess which distribution was sampled from if and only if $x \notin L$. By known results, this leads to the conclusion $L \in \text{coNP}/\text{poly}$.

Toward this end, the next lemma is our main technical tool:

Lemma 3.1. *Let $R : \{0, 1\}^{t \times n} \rightarrow \{0, 1\}^{\leq t'}$ be given. Let $A \subseteq \{0, 1\}^n$ be a set of size $M \geq 100t$, and suppose that we select $a^* \sim \mathcal{U}_A$. Then if t is sufficiently large and $t' < 2(t - 1)$, we have*

$$\mathbb{E} [\beta(a^*, A \setminus a^*)] \leq 1 - 10^{-4} . \tag{2}$$

Lemma 3.1 establishes that certain distributions are (at least slightly) “stable” under modification. Related facts, with information-theoretic proofs, appear in [Raz98, KNTSZ07] (see Appendix A), and these can be readily used to obtain our lemma. A distinctive aspect of Lemma 3.1, however, is that it establishes the closeness of the output distribution of R induced by an input to R containing a string a^* , to one from an input distribution to R that does not support a^* . This “apples-to-oranges” comparison is key to our application of Lemma 3.1: we will use it to build small ($\text{poly}(n)$ -size) subsets of L_n that serve as helpful non-uniform advice to prevent *exponential*-size chunks of L_n from being accepted by Verifier. In the “minimax-free” proof being presented here, we will do so in an iterative fashion until all of L_n is “covered” by our advice. This is reminiscent of the incremental approach of Fortnow and Santhanam [FS11] to defining their advice, in their proof that the OR-conjecture holds unless $\text{NP} \subseteq \text{coNP}/\text{poly}$.

In the more general proofs we give in later sections, Lemmas 6.2 and 6.3 will play a role analogous (but not identical) to that of Lemma 3.1 in the current proof.

²¹(here, $L_n = L \cap \{0, 1\}^n$)

²²We note that this amounts to a weakened version of the Disguising-Distribution Lemma of Section 6.

Proof of Lemma 3.1. Suppose to the contrary that $\mathbb{E}[\beta(a^*, A \setminus a^*)] > 1 - 10^{-4}$. Call $a \in A$ “distinctive” if $\beta(a, A \setminus a) \geq .99$; the measure β is bounded by 1, so more than a .99 fraction of $a \in A$ are distinctive.

For each $a \in A$, let $T = T_a$ be the set given by Fact 2.3, with $X := \mathbf{R}_{A \setminus a}[a], Y := \mathbf{R}_{A \setminus a}$; then for all distinctive $a \in A$, we have

$$\Pr_{z \sim \mathbf{R}_{A \setminus a}[a]}[z \in T_a] \geq .99, \quad \Pr_{z \sim \mathbf{R}_{A \setminus a}}[z \notin T_a] \geq .99. \quad (3)$$

Let us index A as $A = \{a^1, \dots, a^M\}$. Define a random R -input $\mathbf{x} = (x^1, \dots, x^t) \sim \mathcal{U}_A^{\otimes t}$, and for $i \in [M]$ let $\text{Incl}_i(\mathbf{x})$ be the indicator variable for the event that at least one of the elements x^j is equal to a^i . We also define the indicator variable

$$\text{Corr}_i(\mathbf{x}) := [\text{Incl}_i(\mathbf{x}) \Leftrightarrow (R(\mathbf{x}) \in T_{a^i})] = \neg[\text{Incl}_i(\mathbf{x}) \oplus (R(\mathbf{x}) \in T_{a^i})].$$

The idea is that $R(\mathbf{x}) \in T_{a^i}$ “suggests” that a^i was included among the inputs to R , while $R(\mathbf{x}) \notin T_{a^i}$ suggests the opposite; $\text{Corr}_i(\mathbf{x})$ checks whether the suggestion given is correct.

It is easy to see that, if we condition on $[\text{Incl}_i(\mathbf{x}) = 0]$, then $R(\mathbf{x})$ is distributed as $\mathbf{R}_{A \setminus a^i}$. In this case, the conditional probability that $[\text{Corr}_i(\mathbf{x}) = 1]$ holds is at least .99, provided i is distinctive.

On the other hand, suppose we condition on $[\text{Incl}_i(\mathbf{x}) = 1]$. Then the conditional probability that a^i appears *twice* among the coordinates of \mathbf{x} is, by basic counting, at most $t/M \leq .01$. (After conditioning on any leftmost occurrence of a^i , there are only $t - 1$ indices which could contain the next occurrence of a^i ; and each plays this role with probability at most $1/M$.) Thus under this conditioning, $R(\mathbf{x})$ is .01-close to the distribution $\mathbf{R}_{A \setminus a^i}[a^i]$, so that $[\text{Corr}_i(\mathbf{x}) = 1]$ holds with probability at least $.99 - .01 = .98$ if a^i is distinctive.

It is also the case that $\sum_{i \in [M]} \text{Incl}_i(\mathbf{x}) \geq .95t$ with probability at least .99 (for sufficiently large t), since $t/M \leq .01$. Combining all of our work, we find that for large enough t , with probability at least .5 the following conditions hold:

1. $\sum_{i \in [M]} \text{Incl}_i(\mathbf{x}) \geq .95t$;
2. $\sum_{i \in [M]} [\text{Incl}_i(\mathbf{x}) \wedge \text{Corr}_i(\mathbf{x})] \geq .9t$;
3. $\sum_{i \in [M]} \text{Corr}_i(\mathbf{x}) \geq .9M$.

Say that \mathbf{x} is *good* if all of these conditions hold.

Now fix any R -output $z \in \{0, 1\}^{\leq t'}$; we are going to derive an upper bound U on the number of good inputs \mathbf{x} for which $R(\mathbf{x}) = z$. Since every \mathbf{x} maps to a string of length $\leq t'$ under R , it will follow that

$$2^{t'+1} \geq \frac{.5|A^{\times t}|}{U} = \frac{.5M^t}{U}, \quad (4)$$

which will yield a contradiction to our settings.

First, suppose $z \in T_{a^i}$ for more than $t + .1M$ indices $i \in [M]$. Then for any \mathbf{x} such that $R(\mathbf{x}) = z$, there are more than $.1M$ indices for which $\text{Incl}_i(\mathbf{x}) = 0$ yet $z \in T_{a^i}$. For such i , $\text{Corr}_i(\mathbf{x}) = 0$. Thus \mathbf{x} is not good. So to have *any* good inputs \mathbf{x} map to it under R , z must satisfy

$$|\{i : z \in T_{a^i}\}| \leq t + .1M. \quad (5)$$

Next, suppose $R(\mathbf{x}) = z$ and that $\mathbf{x} = (x^1, \dots, x^t)$ contains more than $.15t$ components x^j whose value is any element $x^j = a^i \in A$ for which $z \notin T_{a^i}$. If \mathbf{x} is good, then by property 1 of good strings, among these components we can find a subcollection of more than $.1t$ components x^j whose values are pairwise distinct. For each $a^i = x^j$ in this subcollection, we have $\text{Incl}_i(\mathbf{x}) = 1$ yet $\text{Corr}_i(\mathbf{x}) = 0$. Thus $\sum_{i \in [M]} [\text{Incl}_i(\mathbf{x}) \wedge \text{Corr}_i(\mathbf{x})] < .9t$, so \mathbf{x} is not good—a contradiction. Thus any good \mathbf{x} for which $R(\mathbf{x}) = z$ can contain at most $.15t$ components x^j whose value $x^j = a^i$ satisfies $z \notin T_{a^i}$.

Combining this observation with Eq. (5), there is a set $A' \subseteq A$ of size at most $t + .1M \leq .11M$, such that for any good \mathbf{x} mapping to z under R , at least $.85t$ components x^j satisfy $x^j \in A'$. We can now bound the number of good inputs \mathbf{x} mapping to z under R ; any such \mathbf{x} is specifiable by:

- a set of at most $.15t$ “exceptional” indices $j \in [t]$;
- the values of x^j on these exceptional indices;
- the values of x^j on all other indices, which must lie in A' .

The number of such \mathbf{x} is at most

$$\begin{aligned} \sum_{0 \leq t' \leq .15t} \binom{t}{t'} M^{t'} (.11M)^{t-t'} &\leq (.11)^{.85t} M^t \cdot \sum_{0 \leq t' \leq .15t} \binom{t}{t'} \\ &\leq (.11)^{.85t} M^t \cdot 2^{H(.15)t} \\ &< 4^{-t} M^t, \end{aligned}$$

using Fact 2.2 and a calculation. Thus we may take as our bound $U := 4^{-t} M^t$, so that by Eq. (4),

$$2^{t'+1} \geq .5 \cdot 4^t = 2^{2t-1},$$

which contradicts our assumption that $t' < 2(t-1)$. This proves Lemma 3.1. \square

Proof of Theorem 1.1. We will show that the existence of the reduction R for L implies that there exists a two-message, private-coin, *interactive proof system* between a polynomial-time-bounded Verifier and a computationally unbounded Prover to prove that a given string $x \in \{0, 1\}^n$ lies in \bar{L} . The proof system will be executable using $\text{poly}(n)$ bits of non-uniform advice on length- n inputs; Prover will be able to make Verifier accept with probability 1 if $x \notin L$, and with probability at most $1 - \Omega(1)$ if $x \in L$. It then follows from known results on interactive proof systems and non-uniform derandomization [GS86, Adl78] that $\bar{L} \in \text{NP}/\text{poly}$ (see Theorem 4.11 and the proof of Theorem 4.15 for details), which gives our desired conclusion.

Using the existence of the reduction R and Lemma 3.1, we will prove the following claim:

Claim 3.2. *There exist multisets $A_1, \dots, A_{q(n) \leq \text{poly}(n)} \subseteq L_n$, each of size bounded by some $s(n) \leq \text{poly}(n)$, such that, for all $x \in \{0, 1\}^n \setminus \left(\bigcup_{i \in [q(n)]} A_i\right)$:*

1. *If $x \in \bar{L}_n$, then $\beta(x; A_i) = 1$ for all $i \in [q(n)]$;*
2. *If $x \in L_n$, there is an $i \in [q(n)]$ for which $\beta(x; A_i) \leq 1 - 10^{-5}$.*

Assuming the truth of Claim 3.2 for the moment, we use it to prove Theorem 1.1. For inputs of length n to our interactive proof system, we let the non-uniform advice be a description of the sets $A_1, \dots, A_{q(n)}$ given by Claim 3.2, along with the value $t(n)$. The proof system works as follows. On input $x \in \{0, 1\}^n$, Verifier first checks if x is in one of the sets A_i . If so, Verifier knows that $x \in L$. Otherwise, Verifier and Prover execute the following procedure in parallel for $i = 1, 2, \dots, q(n)$:

- Verifier privately flips an unbiased coin $b_i \sim \mathcal{U}_{\{0,1\}}$;
- Verifier privately samples strings $y^{i,1}, \dots, y^{i,t(n)} \in \{0, 1\}^n$ independently from \mathcal{U}_{A_i} ;
- If $b_i = 0$ then Verifier sets

$$z = z(i) := R(y^{i,1}, \dots, y^{i,t(n)}) ;$$

otherwise ($b_i = 1$), Verifier samples $\mathbf{j} = \mathbf{j}(i) \sim \mathcal{U}_{[t(n)]}$ and sets

$$z := R(y^{i,1}, \dots, y^{i,\mathbf{j}-1}, x, y^{i,\mathbf{j}+1}, \dots, y^{i,t(n)}) .$$

- Verifier sends z to Prover.
- Prover makes a guess \tilde{b}_i for the value of b_i .

Verifier accepts iff $\tilde{b}_i = b_i$ for all i .

This protocol is clearly polynomial-time executable by Arthur given $t(n)$ and the description of $A_1, \dots, A_{q(n)}$, and these sets are of polynomial size and polynomial in number. Now let us analyze the behavior of the protocol (assuming $x \notin \bigcup_i A_i$). First, suppose that $x \in \bar{L}_n$. In this case, we have

$$\|\mathbf{R}_{A_i}[x] - \mathbf{R}_{A_i}\|_{\text{stat}} = 1$$

for each i , by the first property of our sets A_i . Thus, Prover can guess b_i with perfect confidence for each i , and can cause Verifier to accept with probability 1.

Next, suppose that $x \in L_n$. Then by the second property of our sets, there exists an $i^* \in [q(n)]$ such that

$$\|\mathbf{R}_{A_{i^*}}[x] - \mathbf{R}_{A_{i^*}}\|_{\text{stat}} \leq 1 - 10^{-5} .$$

By the distinguishability characterization of statistical distance, and the independence of the trials $i = 1, 2, \dots, q(n)$, this implies that the probability that Prover guesses b_{i^*} correctly is at most $1 - .5 \cdot 10^{-5}$. Thus Verifier rejects with probability $\Omega(1)$. So our interactive proof has the desired properties. As discussed earlier, this implies $\bar{L} \in \text{NP/poly}$. \square

Proof of Claim 3.2. Fixing attention to a single value of n , let $(t, t') = (t(n), t'(n))$. Assume that t is large enough to apply Lemma 3.1. (Note that then t' satisfies the assumptions of that lemma as well.) Let $M := 100t$. We define a sequence of sets $S_1 \supseteq S_2 \supseteq \dots \supseteq S_{q(n)+1} = \emptyset$, each contained in L_n , and a sequence of sets $A_1, A_2, \dots, A_{q(n)}$, with all elements of A_i drawn from S_i .

Let $S_1 := L_n$. Inductively, having defined S_i , we define A_i, S_{i+1} as follows. If $|S_i| < M$, we let $A_i := S_i$ and $S_{i+1} := \emptyset$, and set $q(n) := i$, terminating the construction at this stage. Otherwise ($|S_i| \geq M$), we let A_i be a uniformly random size- $(M - 1)$ subset of S_i . We let

$$S_{i+1} := \{a \in S_i \setminus A_i : \beta(a, A_i) > 1 - 10^{-5}\} .$$

The procedure clearly terminates, since $|S_{i+1}| \leq |S_i| - (M - 1)$ whenever $S_{i+1} \neq \emptyset$. Let us verify that these A_i satisfy conditions 1-2 of the Claim; we will then argue that $q(n) \leq \text{poly}(n)$ (with high probability over the randomness in the construction).

First, suppose $x \in \bar{L}_n \setminus \left(\bigcup_{i \in [q(n)]} A_i\right)$. Then with attention to Eq. (1), note that R always outputs an element of \bar{L}' when x is one of the inputs to R . On the other hand, when all inputs to R are drawn from some $A_i \subseteq S_i \subseteq L_n$, R outputs an element of L' . Thus these two cases are perfectly distinguishable, and $\beta(x, A_i) = 1$ for each i , as needed.

Next suppose $x \in L_n \setminus \left(\bigcup_{i \in [q(n)]} A_i\right)$. Let $i \in [1, q(n)]$ be the unique index such that $x \in S_i \setminus S_{i+1}$. Then by the definitions, we have $\beta(x, A_i) = \|\mathbf{R}_{A_i}[x] - \mathbf{R}_{A_i}\|_{\text{stat}} \leq 1 - 10^{-5}$.

Finally, we argue that $q(n) \leq \text{poly}(n)$ with high probability. Note that when we generate A_i as a uniform set of size $M - 1$, we may equivalently generate A_i by first generating a uniform set $\hat{A}_i \subseteq S_i$ of size M , then selecting a uniform element a^* of \hat{A}_i to discard to form A_i .

By Lemma 3.1, $\mathbb{E}_{a^*}[\beta(a^*, A_i)] \leq 1 - 10^{-4}$. Then with probability at least .9 over our randomness at this stage, a^* satisfies $\beta(a^*, A_i) \leq 1 - 10^{-5}$. But a^* is distributed as a uniform element of $S_i \setminus A_i$. Thus,

$$\mathbb{E}[|S_{i+1}|] \leq .1(|S_i| - |A_i|).$$

Thus $q(n) = O(n)$ with high probability. This completes the proof of Claim 3.2. \square

4 Preliminaries II

Now we collect facts and definitions that will inform our work in the rest of the paper as we prove more general results.

4.1 Information theory background

Recall from Section 2 that $H(\alpha)$ denotes the binary entropy function on $[0, 1]$. For a finitely-supported random variable Z , we let

$$H_{rv}(Z) := \sum_{z \in \text{supp}(Z)} -\Pr[Z = z] \log_2 \Pr[Z = z]$$

denote the Shannon entropy of Z . Then, for two possibly-dependent random variables Y, Z ,

$$H_{rv}(Z|Y) := \mathbb{E}_{y \sim Y}[H_{rv}(Z_{[Y=y]})] = H_{rv}((Y, Z)) - H_{rv}(Y)$$

denotes the *entropy of Z conditional on Y* . ($Z_{[Y=y]}$ denotes Z conditioned on the event $[Y = y]$.)

Fact 4.1 (Subadditivity of entropy). *For all X, Y , $H_{rv}((X, Y)) \leq H_{rv}(X) + H_{rv}(Y)$ and $H_{rv}(X|Y) \leq H_{rv}(X)$, with equality holding in each case iff X, Y are independent. Similarly, $H_{rv}(X|(Y, Z)) \leq H_{rv}(X|Y)$.*

Definition 4.2 (Mutual information). *The mutual information between random variables X, Y is defined as $I(X; Y) := H_{rv}(X) + H_{rv}(Y) - H_{rv}((X, Y))$.*

The next fact follows easily from the definitions.

Fact 4.3. *Mutual information obeys the following properties, for all random variables X, Y, Z :*

1. $I(X; Y) = I(Y; X)$;
2. $I(X; (Y, Z)) = I(X; Y) + I((X, Y); Z) - I(Y; Z)$;
3. $I(X; (Y, Z)) \geq I(X; Y)$;
4. $I(X; Z) = 0$ if X, Z are independent.

Lemma 4.4. *If X^1, \dots, X^t are independent, then*

$$I(Y; (X^1, \dots, X^t)) \geq \sum_{j \in [t]} I(Y; X^j) .$$

Our proof of this standard claim follows steps in [Nay99a, p. 33].

Proof. We have

$$\begin{aligned} I(Y; (X^1, \dots, X^t)) &= I(Y; X^t) + I((Y, X^t); (X^1, \dots, X^{t-1})) - \underbrace{I(X^t; (X^1, \dots, X^{t-1}))}_{=0, \text{ by Fact 4.3, item 4}} \\ &\geq I(Y; X^t) + I(Y; (X^1, \dots, X^{t-1})) , \end{aligned}$$

where we used item 2 of Fact 4.3 in the first step, and items 1 and 3 in the second step. Iterating in this way gives the Lemma. \square

The next definition is a useful, non-symmetric measure of difference between random variables.

Definition 4.5 (KL divergence). *The (binary) Kullback-Leibler divergence, or KL divergence between random variables X, Y , is denoted $D_{\text{KL}}(X||Y)$ and defined as*

$$D_{\text{KL}}(X||Y) := \sum_{x \in \text{supp}(X)} \Pr[X = x] \cdot \log_2 \left(\frac{\Pr[X = x]}{\Pr[Y = x]} \right) .$$

The convention is that for $p \neq 0$, we have $p \log_2(p/0) = +\infty$. So D_{KL} may be infinite. We have the following basic equivalence (see [CT06, Chapter 2]):

Fact 4.6. *Let X, Y be any random variables; let X' be distributed as X and independent of Y . The mutual information and Kullback-Leibler divergence satisfy*

$$I(X; Y) = D_{\text{KL}}((X, Y)|| (X', Y)) .$$

A proof of the following important result can be found in [CT06] (see Lemma 11.6.1, p. 370).

Theorem 4.7 (Pinsker's inequality, stated for binary KL divergence). *For any random variables Z, Z' ,*

$$D(Z||Z') \geq \frac{2}{\ln 2} \cdot \|Z - Z'\|_{\text{stat}}^2 .$$

When $\|Z - Z'\|_{\text{stat}} \approx 1$, the following bound, known as Vajda's inequality (see [FHT03, RW09]), gives better information on the divergence:

Theorem 4.8 (Vajda's inequality, stated for binary KL divergence). *For any random variables Z, Z' , let $\Delta := \|Z - Z'\|_{\text{stat}}$. Then,*

$$D(Z||Z') \geq \frac{1}{\ln 2} \left(\ln \left(\frac{1 + \Delta}{1 - \Delta} \right) - \frac{2\Delta}{1 + \Delta} \right) \geq \frac{1}{\ln 2} \left(\ln \left(\frac{1}{1 - \Delta} \right) - 1 \right) .$$

4.2 Basic complexity classes and promise problems

We assume familiarity with the basic complexity classes NP and coNP and the higher levels Σ_k^p, Π_k^p of the Polynomial Hierarchy PH. (For the needed background in complexity theory, see [AB09].) In this paper we define NP, coNP, etc. as classes of *languages* (not promise problems).

We also assume familiarity with the general model of polynomial-size, non-uniform advice, and with the non-uniform classes NP/poly and coNP/poly. It is considered unlikely that $\text{NP} \subseteq \text{coNP/poly}$. In particular, this would imply a collapse of the Polynomial Hierarchy:

Theorem 4.9 ([Yap83]). *If $\text{NP} \subseteq \text{coNP/poly}$, then $\text{PH} = \Sigma_3^p = \Pi_3^p$.*

We use pr-NP, pr-coNP, etc. to denote the analogous complexity classes for *promise problems*. Recall that, for a class C of promise problems, $\text{coC} = \{(\Pi_Y, \Pi_N) : (\Pi_N, \Pi_Y) \in \text{C}\}$. A *many-to-one* reduction B from the promise problem $\Pi = (\Pi_Y, \Pi_N)$ to $\Pi' = (\Pi'_Y, \Pi'_N)$ is a mapping satisfying $B(\Pi_Y) \subseteq \Pi'_Y, B(\Pi_N) \subseteq \Pi'_N$. (This definition applies as well to the special case where one or both of the promise problems are languages.) When we refer to NP-complete problems in this paper, we mean problems complete under deterministic, polynomial-time many-to-one reducibility.

All of the results we prove in this paper about limits of compression for *languages* L and language complexity classes readily extend to the setting of compression for promise problems (under the analogous definitions). However, for notational simplicity we will state our main results for languages, and will only use promise problems and promise classes where doing so helps to streamline our proofs and our result statements.

4.3 Arthur-Merlin protocols

We will make use of the model of (public-coin, two-round) *Arthur-Merlin protocols*. To be precise, these are protocols \mathcal{P} , defined by a deterministic polynomial-time predicate $A(x, r, w)$, which operate as follows. On an input x , visible to both a polynomial-time bounded verifier (Arthur) and to a computationally-unbounded prover (Merlin):

1. Arthur generates a uniformly random string r and sends it to Merlin;
2. Merlin sends a response string w to Arthur;
3. Arthur accepts if $A(x, r, w) = 1$, otherwise rejects.

We require that $|r|, |w|$ each be pre-specified lengths $\leq \text{poly}(n)$, where $n = |x|$, and that these lengths be computable in $\text{poly}(n)$ time given 1^n .

We will need to work with promise problems having Arthur-Merlin protocols. Say that such a protocol \mathcal{P} *defines* a promise problem $\Pi = (\Pi_Y, \Pi_N)$ with *completeness* $c(n)$ and *soundness* $s(n)$ if

1. For all $x \in \Pi_Y$, some Merlin strategy causes Arthur to accept with probability $\geq c(n)$;
2. For all $x \in \Pi_N$, all Merlin strategies cause Arthur to accept with probability $\leq s(n)$.

Let $\text{pr-AM}_{c(n),s(n)}$ denote the class of promise problems definable by an Arthur-Merlin protocol with completeness $c(n)$ and soundness $s(n)$; let $\text{pr-AM} := \text{pr-AM}_{1,1/3}$. Then, $\text{pr-coAM} = \{(\Pi_Y, \Pi_N) : (\Pi_N, \Pi_Y) \in \text{pr-AM}\}$.

Theorem 4.10 ([FGM+89]). *For any parameters $s(n), c(n) \in (0, 1]$ that are polynomial-time computable²³ and satisfy $\frac{1}{\text{poly}(n)} < s(n) < c(n) - \frac{1}{\text{poly}(n)}$, we have $\text{pr-AM}_{c(n),s(n)} = \text{pr-AM}$. If we drop the requirement $s(n) > \frac{1}{\text{poly}(n)}$, but keep the gap requirement, we still have $\text{pr-AM}_{c(n),s(n)} \subseteq \text{pr-AM}$.*

The next, well-known result follows from the non-uniform derandomization technique of Adleman [Adl78]:

Theorem 4.11. $\text{pr-AM} \subseteq \text{pr-NP/poly}$. *Similarly, $\text{pr-coAM} \subseteq \text{pr-coNP/poly}$.*

4.4 Statistical zero-knowledge and the SD problem

Next we will define the *statistical zero-knowledge* class SZK. Actually, we will only work with its promise-problem analogue pr-SZK .²⁴ Informally, these are the promise problems (Π_Y, Π_N) for which a (private-coin) interactive proof of membership in Π_Y can be given, in which the verifier *learns (almost) nothing*—*except* to become convinced that the input y indeed lies in Π_Y ! The “learns nothing” requirement is cashed out by requiring that the verifier be able to *simulate* interactions with the intended prover strategy on any input y , such that if $y \in \Pi_Y$, the resulting distribution is negligibly close in statistical distance to the true distribution generated by their interaction.

Making this definition formal is somewhat delicate. (For details, and for more information on these and related classes, see [SV03].) Fortunately, there is a simple (but non-trivial) characterization of pr-SZK . First, given a Boolean circuit $C = C(r)$ with k output gates, and an ordering on these gates, let \mathcal{D}_C denote the output distribution of C on a uniformly random input r . (This is a random variable over $\{0, 1\}^k$.) We use the following problem:

Definition 4.12. *For parameters $0 \leq d \leq D \leq 1$, define the promise problem $\text{SD}_{\leq d}^{\geq D} = (\Pi_Y, \Pi_N)$ as follows:*

$$\begin{aligned} \Pi_Y &:= \{ \langle C, C' \rangle : \|\mathcal{D}_C - \mathcal{D}_{C'}\|_{\text{stat}} \geq D \} , \\ \Pi_N &:= \{ \langle C, C' \rangle : \|\mathcal{D}_C - \mathcal{D}_{C'}\|_{\text{stat}} \leq d \} . \end{aligned}$$

Define $\text{SD}_{> D}^{\leq d}$ analogously, switching the “yes” and “no” cases. In this definition, both $d = d(n)$ and $D = D(n)$ may be parameters depending on the input length $n = |\langle C, C' \rangle|$.

It is shown in [SV03] that the standard, complicated definition of pr-SZK is equivalent to the following simpler one, which we take as our definition:

Definition 4.13. *Let pr-SZK be defined as the class of promise problems for which there is a many-to-one,²⁵ deterministic polynomial-time reduction from Π to $\text{SD}_{\leq 1/3}^{\geq 2/3}$.*

The constants $2/3, 1/3$ in the above definition are not arbitrary; it is unknown whether we get the same class if we replace them by $.51, .49$. However, we have the following result:

Theorem 4.14 (Follows from [SV03]; described as Theorem 1 in [GV11]). *Suppose $0 \leq d = d(n) < D = D(n) \leq 1$ are polynomial-time computable, and satisfy $D^2 > d + \frac{1}{\text{poly}(n)}$. Then, $\text{SD}_{\leq d}^{\geq D} \in \text{pr-SZK}$.*

²³(say, as rational values represented by their numerator and denominator)

²⁴Often the promise class is denoted SZK.

²⁵Recall the definition in Section 4.2.

When we merely have $D - d \geq \frac{1}{\text{poly}(n)}$, the following weaker, standard result holds:

Theorem 4.15. *Suppose $0 \leq d = d(n) < D = D(n) \leq 1$ are polynomial-time computable and satisfy $D > d + \frac{1}{\text{poly}(n)}$. Then, $\text{SD}_{\leq d}^{\geq D} \in \text{pr-AM}$.*

Proof sketch. We describe a *private-coin* two-message protocol, in which the verifier has a source of random bits not viewable by the prover; any such protocol can be efficiently converted into a public-coin one [GS86].

Let $m = m(n) \leq \text{poly}(n)$ be a large value. On input $\langle C, C' \rangle$, Verifier chooses b_1, \dots, b_m uniformly at random and, for $i \in [m]$, samples

$$z^i \sim \mathcal{D}_C \quad \text{if } b_i = 0, \quad z^i \sim \mathcal{D}_{C'} \quad \text{if } b_i = 1,$$

independently for each i . Prover is asked to try to guess the values b_1, \dots, b_m .

If m is chosen appropriately large then, using the distinguishability interpretation of statistical distance (see Section 2.1),

1. If $\|\mathcal{D}_C - \mathcal{D}_{C'}\|_{\text{stat}} > D$ then Prover can, with high probability, guess at least a $\frac{1}{2} \left(1 + \frac{D-d}{2}\right)$ fraction of the bits b_i correctly;
2. If $\|\mathcal{D}_C - \mathcal{D}_{C'}\|_{\text{stat}} < d$ then Prover cannot, except with low probability, guess this fraction of the b_i s correctly.

Thus, Verifier can use this threshold as an acceptance criterion, so that the protocol has the desired completeness-soundness gap. After converting to a public-coin protocol, we find that $\text{SD}_{\leq d}^{\geq D} \in \text{pr-AM}_{2/3, 1/3} = \text{pr-AM}$ (using Theorem 4.10). \square

We will also use the following important results about pr-SZK:

Theorem 4.16 ([Oka00]). *pr-SZK is closed under complement.*

Theorem 4.17. $\text{pr-SZK} \subseteq \text{pr-AM} \cap \text{pr-coAM} \subseteq \text{pr-NP/poly} \cap \text{pr-coNP/poly}$.

The containment in pr-coAM is due to Fortnow [For87]; containment in pr-AM was first shown by Aiello and Håstad [AH91].²⁶ The second containment in Theorem 4.17 uses Theorem 4.11.

Finally, one of our results (Theorem 7.3) will make use of the class pr-PZK of problems having (honest-verifier) *perfect* zero-knowledge proofs. This is a subclass of pr-SZK. We will not define pr-PZK (see, e.g., [SV03]); unfortunately it has no known simple characterization analogous to Definition 4.13 for pr-SZK. We will, however, use the following result:

Theorem 4.18 ([SV03], Proposition 5.7). $\text{SD}_{\leq \frac{1}{5}}^{\geq 1} \in \text{pr-PZK}$.

Next we combine tools described in Sections 4.3 and 4.4, reformulating them slightly.

Theorem 4.19. *Let $0 \leq d = d(n) < D = D(n) \leq 1$ be (not necessarily computable) parameters.*

1. *If $D > d + \frac{1}{\text{poly}(n)}$, then $\text{SD}_{\leq d}^{\geq D} \in \text{pr-NP/poly}$.*

²⁶These works treat language classes, but the proofs extend without change to the promise-problem setting. Also, these works analyze a so-called “honest-verifier” model of statistical zero-knowledge proofs; these were shown to have the same expressive power as “cheating-verifier” statistical zero-knowledge proofs in [GSV98].

2. If we have the stronger gap $D^2 > d + \frac{1}{\text{poly}(n)}$, then $\text{SD}_{\leq d}^{\geq D}$ is many-to-one reducible to $\text{SD}_{\leq 1/3}^{\geq 2/3} \in \text{pr-SZK}$, in non-uniform polynomial time. Also, $\text{SD}_{\leq d}^{\geq D} \in \text{pr-coNP/poly}$.

Proof sketch. For item 1, we essentially combine Theorem 4.15 with Theorem 4.11. The only extra ingredient needed is to encode sufficiently accurate approximations of $d(n), D(n)$ into the non-uniform advice for length n , and to use these in defining the private-coin protocol as in the proof of Theorem 4.15. We then convert this non-uniform protocol into an NP/poly one by the same techniques from [GS86] (which shows how to convert private-coin to public-coin protocols), Theorem 4.10 (to get perfect completeness), and Theorem 4.11 (to derandomize).

Similarly, for item 2, we essentially combine Theorems 4.14 and 4.17, except that at each step we need to incorporate approximations of $d(n), D(n)$ as (additional) non-uniform advice. \square

4.5 f -compression reductions

Here we define a class of compression reductions for the problems $f \circ L$ introduced in Section 1.3.1, in which one is given (x^1, \dots, x^m) and must compute $f(L(x^1), \dots, L(x^m))$. Our main focus will be the case where f is the OR or AND function of its input bits. The problem $f \circ L$ will be formally defined as a parametrized problem in Section 5.1, but it will be useful to have a specialized definition for this problem as well; here we won't explicitly rely on the parametrized-problem framework.

Our next definition is modeled on definitions in [BDFH09, FS11], with some differences. Notably, we will consider reductions where a quantitative compression guarantee is only made when all the input strings x^j are of some equal length n , and the number of input strings x^j is equal to some value $t_1(n)$ determined by n . The error bound will also be a function of n . This specialization is mostly to reduce clutter in our work, and will not lead to loss of generality: we will be *ruling out* the existence of compression reductions (under complexity-theoretic assumptions, and for all $t_1(n)$ that are sufficiently large compared to other parameters), so ruling out even compression algorithms that work only in narrow input-regimes will lead to stronger results.

Definition 4.20 (Probabilistic f -compression reductions). *Let L, L' be two languages, and let $f : \{0, 1\}^* \rightarrow \{0, 1\}$ be a Boolean function. Let $t_1(n), t_2(n) : \mathbb{N}^+ \rightarrow \mathbb{N}^+$ and $\xi(n) : \mathbb{N}^+ \rightarrow [0, 1]$ be given.*

A probabilistic f -compression reduction for L , with parameters $(t_1(n), t_2(n), \xi(n))$ and target language L' , is a randomized mapping $R(x^1, \dots, x^m)$ outputting a string z , such that for all $(x^1, \dots, x^{t_1(n)}) \in \{0, 1\}^{t_1(n) \times n}$,

1. $\Pr_R[L'(z) = f(L(x^1), \dots, L(x^{t_1(n)}))] \geq 1 - \xi(n)$;
2. $|z| \leq t_2(n)$.

If some reduction R as above is computable in probabilistic polynomial time, we say that L is PPT- f -compressible with parameters $(t_1(n), t_2(n), \xi(n))$. (This does not require that $(t_1(n), t_2(n), \xi(n))$ themselves be computable.)

5 Parametrized problems and parametrized compression

A central aim of our work is to better understand the limitations of efficient compressive reductions for a variety of parametrized problems. For this we need to formally define parametrized problems

and an appropriate model of probabilistic compression for these problems. However, some readers may be satisfied to understand our work on the limits of efficient AND- and OR-compression (as defined in Section 4.5) for SAT and other NP-complete languages. To prove these results, including Theorem 1.3 in the Introduction, we will not need the definitions of this section, and readers may choose to skip ahead to Section 6. (We will find it convenient to prove Theorem 1.2 using the definitions below; however, this result can also be derived directly from our Theorem 7.1, item 2 with little trouble.)

5.1 Parametrized problems

We will use the following definition:

Definition 5.1 ([DF99]). *A parametrized problem is a subset of binary strings of the form $\langle x, 1^k \rangle$, for $x \in \{0, 1\}^*$ and $k > 0$ (under some natural binary encoding of such tuples).*

Thus, our convention is that a parametrized problem is just a particular type of decision problem, i.e., a language.²⁷ However, we will use P to denote a generic parametrized problem, as opposed to an “ordinary” language, denoted L . Sometimes, as in the Introduction, we speak of “parametrized versions” of an ordinary decision problem L . There is no single, canonical way to go from a decision problem to a parametrized problem; often, however, a parametrized problem can be formed from a decision problem L in a natural way. For example, we formally define VAR-SAT, OR(SAT), and AND(SAT) from the Introduction as follows:

Definition 5.2. *Fix some natural encoding of tuples of bit-strings, and some encoding of Boolean formulas as bit-strings. Define*

1. VAR-SAT := $\{\langle \psi, 1^k \rangle \mid \psi \text{ is satisfiable and contains } \leq k \text{ distinct variables}\}$;
2. OR(SAT) := $\{\langle \psi_1, \dots, \psi_t, 1^k \rangle \mid \text{at least one } \psi_j \text{ is satisfiable, and each } \psi_j \text{ is of bit-length } \leq k\}$;
3. AND(SAT) := $\{\langle \psi_1, \dots, \psi_t, 1^k \rangle \mid \text{every } \psi_j \text{ is satisfiable, and each } \psi_j \text{ is of bit-length } \leq k\}$.

We also generalize items 2 and 3 above:

Definition 5.3. *Let $L \subseteq \{0, 1\}^*$, and $f : \{0, 1\}^* \rightarrow \{0, 1\}$. Define*

1. OR(L) := $\{\langle (x^1, \dots, x^t), 1^k \rangle \mid \bigvee_{j=1}^t L(x^j) = 1 \text{ and } |x^j| \leq k \text{ for each } j\}$;
2. AND(L) := $\{\langle (x^1, \dots, x^t), 1^k \rangle \mid \bigwedge_{j=1}^t L(x^j) = 1 \text{ and } |x^j| \leq k \text{ for each } j\}$;
3. $f \circ L$:= $\{\langle (x^1, \dots, x^t), 1^k \rangle \mid f(L(x^1), \dots, L(x^t)) = 1 \text{ and } |x^j| \leq k \text{ for each } j\}$.

²⁷In this definition we are following [FS11]. In [BDFH09] and many other works, parametrized problems are defined as a subset of $\{0, 1\}^* \times \mathbb{N}^+$ (the parameter is still presented as part of the input); they refer to the corresponding subset of strings of form $\langle x, 1^k \rangle$ as the “unparametrized version” or “classical version” of the problem.

5.2 OR-expressive and AND-expressive parametrized problems

Our compression lower bounds will apply to two classes of parametrized problems. As we will explain, these classes are closely related to classes identified earlier in [HN10, BDFH09, BJK11a, BTY11]; the classes we introduce will help to apply our techniques uniformly to these various earlier classes.

Definition 5.4 (OR- and AND-expressive problems). *A parametrized problem P is OR-expressive, with parameter $S(n) \leq \text{poly}(n)$, if there exists an NP-complete language L and a deterministic polynomial-time reduction B . Whenever B receives an input of form $\langle (x^1, \dots, x^t), 1^n \rangle$, for any $t, n \in \mathbb{N}^+$, B outputs a tuple*

$$\langle \langle y^1, 1^{k_1} \rangle, \dots, \langle y^s, 1^{k_s} \rangle \rangle.$$

We have the following properties:

1. $\langle (x^1, \dots, x^t), 1^n \rangle \in \text{OR}(L) \iff \exists i \in [s] : \langle y^i, 1^{k_i} \rangle \in P$;
2. $s \leq S(n)$ (in particular, the bound is independent of t);
3. For each $i \in [s]$, $|y^i| \leq (t+n)^{O(1)}$ and $k_i \leq n^{O(1)}$.

Define AND-expressive problems identically, except we replace condition 1 above by

$$1'. \langle (x^1, \dots, x^t), 1^n \rangle \in \text{AND}(L) \iff \forall i \in [s] : \langle y^i, 1^{k_i} \rangle \in P.$$

The results of [BDFH09] imply that a variety of natural parametrized problems are OR- or AND-expressive:

Theorem 5.5 (Follows from [BDFH09]). 1. OR(SAT) is OR-expressive with $S(n) = 1$. Also, each of the following parametrized problems are OR-expressive with $S(n) \leq \text{poly}(n)$:

- k -Path, k -Cycle, k -Exact Cycle and k -Short Cheap Tour,
- k -Graph Minor Order Test and k -Bounded Treewidth Subgraph Test,
- k -Planar Graph Subgraph Test and k -Planar Graph Induced Subgraph Test,
- (k, σ) -Short Nondeterministic Turing Machine Computation,
- w -Independent Set, w -Clique and w -Dominating Set,

defined in [BDFH09].

2. AND(SAT) is AND-expressive with $S(n) = 1$. Also, each of the following parametrized problems are AND-expressive with $S(n) \leq \text{poly}(n)$:

- k -Cutwidth, k -Modified Cutwidth, and k -Search Number,
- k -Pathwidth, k -Treewidth, and k -Branchwidth,
- k -Gate Matrix Layout and k -Front Size,
- w -3-Coloring and w -3-Domatic Number,

also defined in [BDFH09].

In [BDFH09], the authors define a notion of *compositionality* for parametrized problems. If a parametrized problem P is compositional and NP-complete, then it is OR-expressive, with respect to the NP-complete language $L = P$. Also, if P is NP-complete and \overline{P} is compositional, then P is AND-expressive. These facts follow almost immediately from the definitions. Theorem 5.5 then follows from the compositionality results proved in [BDFH09]. In a number of the problems above we can actually take $S(n) = 1$.

Bodlaender, Jansen, and Kratsch [BJK11a] introduced a notion of *cross-compositionality* of parametrized problems, generalizing compositionality. They showed that the evidence against efficient compression against compositional problems given by [BDFH09, FS11] can be extended to cross-compositional problems. Cross-compositional problems are also OR-expressive, as follows from the definitions [BJK11a, Section 3].²⁸ As shown in [BJK11a], this class includes interesting parametrized versions of the Clique, Chromatic Number, and Feedback Vertex Set problems.

AND-expressiveness results are fewer in number, although this is partly due to the fact that, after the results of [FS11] appeared, OR-expressiveness results were preferentially sought. Another example of an AND-expressive problem (not known to be OR-expressive) is presented in [BJK11c].

We also have the following result, derived from the earlier work of [HN10]:

Theorem 5.6 (Follows from [HN10, FS11]). *Each of the problems Clique, Dominating Set,²⁹ Integer Programming, described in [HN10] and modeled as parametrized problems in [FS11] (with slightly distinctive, but natural, parametrizations), are OR-expressive, with $S(n) = 1$.*

A class of reductions between parametrized problems, called *W-reductions*, is used in these works (see [FS11, Definition 2.10]); OR(SAT) is shown to *W-reduce* to each of the problems listed in Theorem 5.6. This immediately implies that these problems are OR-expressive with $S(n) = 1$. Also, if an OR-expressive parametrized problem P *W-reduces* to a second problem Q , then Q is also OR-expressive. This technique was used in [BTY11] to derive additional hardness-of-compression results for problems not easily captured by the compositionality framework; our new results apply to these problems as well.

We remark that the polynomial bounds involved in the reductions of Theorems 5.5 and 5.6 are fairly modest.

5.3 Parametrized compression

We define compression reductions for parametrized problems as follows, following [FS11] (but with some added flexibility in our definitions):

Definition 5.7 (Probabilistic parametrized compression reductions). *Let P be a parametrized problem and L' be a language, and say we are given two functions*

$$c(m, k, w) : (\mathbb{N}^+)^3 \rightarrow \mathbb{N}^+, \quad \xi(m, k, w) : (\mathbb{N}^+)^3 \rightarrow [0, 1] .$$

Say that a randomized mapping $R : \{0, 1\}^ \rightarrow \{0, 1\}^*$ is a (c, ξ) -parametrized compression reduction for P , with target language L' , if for all inputs of form $\langle y, 1^k, 1^w \rangle$, $R(\langle y, 1^k, 1^w \rangle)$ outputs a string z such that:*

²⁸Strictly speaking, according to their definition, cross-compositional problems are OR-expressive under the minor restriction on the reduction in Definition 5.4 that the input $\langle (x^1, \dots, x^t), 1^n \rangle$ satisfy $t \leq 2^{n^a}$, for some $a > 0$. This is of no importance to us, since we will always work with the case $t \leq \text{poly}(n)$; we could have required this in Definition 5.4, and could prove the same variety of hardness results.

²⁹(these are different parametrized problems than w -Clique and w -Dominating Set in Theorem 5.5 above)

1. $\Pr_R[L'(z) = P(\langle y, 1^k \rangle)] \geq 1 - \xi(|y|, k, w)$;
2. $|z| \leq c(|y|, k, w)$.

We call c the compression bound and ξ the error bound of the reduction; we call w the confidence parameter.

For a parametrized problem P , if some reduction R as above is computable in probabilistic polynomial time, we say that P is PPT-compressible with parameters (c, ξ) .

We will not be exploring the full range of possible parameter values in the above definition, but we believe it provides a reasonable framework for future work. (Only a few interesting examples of randomized parametrized compression reductions seem to be known; see [HN10, KW12].) The idea of a confidence parameter w , that one can use to increase the reliability of the compression at the expense of a potentially larger output size, is natural for probabilistic compression and will be useful in our work. (The same basic notion was used earlier in [FS11].)

Next, we define a notion of “strong” compressibility as in the Introduction, preserving flexibility in the error bound:

Definition 5.8. *Say that P is strongly PPT-compressible with error bound $\xi(m, k, w)$, if P is PPT-compressible (to some target language L') with error bound ξ and some compression bound c satisfying $c(m, k, 1) \leq k^{O(1)}$, with the polynomial bound independent of m .*

Using the majority-vote technique of [FS11, Proposition 5.1], we have the following easy result:

Lemma 5.9. *Let $a > 0$. Suppose that P is strongly PPT-compressible with error bound satisfying $\xi(m, k, 1) \leq .5 - k^{-O(1)}$. Then, P is also PPT-compressible with compression bound $c'(m, k, w) \leq k^{O(1)} \cdot w$ and error bound $\xi'(m, k, w) \leq 2^{-w}$.*

5.4 Connecting parametrized compression and f -compression

The next lemma shows that to give evidence against efficient compression for “expressive” parametrized problems, it suffices to give evidence against efficient AND- and OR-compression for NP-complete languages. This lemma is modeled on [BDFH09, Lemma 2], but with some slight complications due to the probabilistic setting. For simplicity we only treat strong compression in the result below; our techniques also extend to give evidence against more modest compression amounts for expressive problems. (For more modest compression amounts, the obtainable results are weaker when the parameter $S(n)$ in the definition of expressiveness is fast-growing.)

Lemma 5.10. *Let L be an NP-complete language.*

1. *Suppose that the parametrized problem P is OR-expressive with respect to L , with parameter $S(n) \leq \text{poly}(n)$. If P is strongly PPT-compressible with error bound $\xi(m, k, 1) \leq .5 - k^{-O(1)}$, then for any polynomially-bounded function $T(n) : \mathbb{N}^+ \rightarrow \mathbb{N}^+$, L is PPT-OR-compressible with parameters*

$$t_1(n) = T(n), \quad t_2(n) \leq S(n) \cdot n^{O(1)}, \quad \xi'(n) \leq 2^{-n}.$$

2. *Suppose P is AND-expressive with respect to L . If P is strongly PPT-compressible with error bound $\xi(m, k, 1) \leq .5 - k^{-O(1)}$, then L is PPT-AND-compressible with parameters $(t_1(n), t_2(n), \xi'(n))$ as in item 1.*

Proof of Lemma 5.10. We will prove item 1 above; item 2 is proved similarly. Let R be the PPT compression reduction R for P given by Lemma 5.9. Let L'_0 be the target language of R . Let B be the reduction for P and L as in Definition 5.4.

We define an OR-compression reduction R' for L , with target language $L' := \text{OR}(L'_0)$, as follows. In defining R' , we let $t_1(n) := T(n)$. On inputs $x^1, \dots, x^{T(n)} \in \{0, 1\}^{T(n) \times n}$, the reduction first applies B to $\langle (x^1, \dots, x^{T(n)}), 1^n \rangle$, yielding a tuple $(\langle y^1, 1^{k_1} \rangle, \dots, \langle y^s, 1^{k_s} \rangle)$. Next, for each $i \in [s]$, R' applies R to the string $\langle y^i, 1^{k_i}, 1^{2n} \rangle$ (here we are selecting the confidence parameter $w := 2n$ for R), yielding an output z^i . Then R' outputs $\langle (z^1, \dots, z^s), 1^M \rangle$, where $M := \max_i |z^i|$.

R' is clearly polynomial-time computable. Now let us analyze its compression and reliability properties. First, each y^i is of bit-length $|y^i| \leq (T(n) + n)^{O(1)}$, and $k_i \leq n^{O(1)}$, by item 3 of Definition 5.4. Then by the compression guarantee for R , each z^i is of bit-length $\leq n^{O(1)} \cdot w = n^{O(1)}$. Thus for the output-size bound of R' we may take $t_2(n) \leq S(n) \cdot n^{O(1)}$, as needed.

Now we bound the error of R' . Using the correctness property of B (Definition 5.4, item 1), the equivalence

$$\langle (z^1, \dots, z^s), 1^M \rangle \in \text{OR}(L'_0) \iff \bigvee_{j=1}^{T(n)} [x^j \in L]$$

holds as long as each application of R , namely $R(\langle y^i, 1^{k_i} \rangle)$ for $i \in [s]$, is successful. By a union bound, this occurs with probability $\geq 1 - S(n) \cdot 2^{-2n}$, which is larger than $1 - 2^{-n}$ for sufficiently large n . (For smaller n , R' may solve its input problem directly by brute force.) Thus for the error bound $\xi'(n)$ for R' , we may take $\xi'(n) \leq 2^{-n}$. \square

6 Technical lemmas

In this section we present our main technical lemmas. Our final goal in this section will be the “Disguising-Distribution Lemma,” our key technical tool for our main results.

6.1 Distributional stability

Here we define the notion of “distributional stability” described in Section 1.4.2.

Definition 6.1. *Let U be some finite universe, and let $T, n \geq 1$ be integers. Given a possibly-randomized mapping $F(x^1, \dots, x^T) : \{0, 1\}^{T \times n} \rightarrow U$, and a collection $\mathcal{D}_1, \dots, \mathcal{D}_T$ of mutually independent distributions over $\{0, 1\}^n$, for $j \in [T]$ let*

$$\gamma_j := \mathbb{E}_{y \sim \mathcal{D}_j} [\|F(\mathcal{D}_1, \dots, \mathcal{D}_{j-1}, y, \mathcal{D}_{j+1}, \dots, \mathcal{D}_T) - F(\mathcal{D}_1, \dots, \mathcal{D}_T)\|_{\text{stat}}] .$$

For $\delta \in [0, 1]$, say that F is δ -distributionally stable (or δ -DS) with respect to $\mathcal{D}_1, \dots, \mathcal{D}_T$ if

$$\frac{1}{T} \sum_{j=1}^T \gamma_j \leq \delta .$$

Lemma 6.2. *Let $R(x^1, \dots, x^t) : \{0, 1\}^{t \times n} \rightarrow \{0, 1\}^{\leq t'}$ be any possibly-randomized mapping, for any $n, t, t' \in \mathbb{N}^+$. R is δ -distributionally stable with respect to any independent input distributions $\mathcal{D}_1, \dots, \mathcal{D}_t$, where we may take either of the following two bounds:*

1. $\delta := \sqrt{\frac{\ln 2}{2} \cdot \frac{t'+1}{t}}$;
2. $\delta := 1 - 2^{-\frac{t'}{t}-3}$.

Our proof of Lemma 6.2, item 1 follows suggestions by Ashwin Nayak and Salil Vadhan; item 2 is a small modification using Vajda's inequality. When $t'/t = 1 - \Omega(1)$, the bound given in item 1 above is within constant factors of the bound from our original distributional stability lemma, Lemma B.4. On the other hand, when $t'/t = 1 - \alpha \approx 1$, the bound in Lemma 6.2, item 1 is better (i.e., smaller) by a $\Theta(\log \frac{1}{\alpha})$ factor. We don't know how to prove a version of item 2 above with the methods of Lemma B.4; this alternative bound is important for our work. In an earlier draft we used a more complicated workaround to prove the results obtainable from item 2.

Proof of Lemma 6.2. Define independent random variables $X^j \sim \mathcal{D}_j$ over $\{0, 1\}^n$, for $j \in [t]$. Let $\mathbf{R} := R(X^1, \dots, X^t)$.

The entropy of \mathbf{R} is at most $\log_2 \left(\left| \{0, 1\}^{\leq t'} \right| \right) < t'+1$. Thus, the mutual information $I((X^1, \dots, X^t); \mathbf{R})$ is less than $t' + 1$. By the independence of the X^j s, Lemma 4.4 gives

$$\sum_{j \in [t]} I(X^j; \mathbf{R}) < t' + 1. \quad (6)$$

By Fact 4.6,

$$I(X^j; \mathbf{R}) = D_{\text{KL}}((X^j, \mathbf{R}) \parallel (Y^j, \mathbf{R})), \quad (7)$$

where $Y^j \sim \mathcal{D}_j$ is independent of \mathbf{R} . By Theorem 4.7,

$$\begin{aligned} D_{\text{KL}}((X^j, \mathbf{R}) \parallel (Y^j, \mathbf{R})) &\geq \frac{2}{\ln 2} \cdot \|(X^j, \mathbf{R}) - (Y^j, \mathbf{R})\|_{\text{stat}}^2 \\ &= \frac{2}{\ln 2} \cdot \mathbb{E}_{x^j \sim \mathcal{D}_j} \left[\left| R(\mathcal{D}_1, \dots, \mathcal{D}_{j-1}, x^j, \mathcal{D}_{j+1}, \dots, \mathcal{D}_t) - R(\mathcal{D}_1, \dots, \mathcal{D}_t) \right|_{\text{stat}} \right]^2, \end{aligned}$$

where the equality follows from the distinguishability interpretation of statistical distance. Using this, we find

$$\begin{aligned} &\left(\frac{1}{t} \sum_{j \in [t]} \mathbb{E}_{x^j \sim \mathcal{D}_j} \left[\left| R(\mathcal{D}_1, \dots, \mathcal{D}_{j-1}, x^j, \mathcal{D}_{j+1}, \dots, \mathcal{D}_t) - R(\mathcal{D}_1, \dots, \mathcal{D}_t) \right|_{\text{stat}} \right] \right)^2 \\ &\leq \frac{1}{t} \sum_{j \in [t]} \mathbb{E}_{x^j \sim \mathcal{D}_j} \left[\left| R(\mathcal{D}_1, \dots, \mathcal{D}_{j-1}, x^j, \mathcal{D}_{j+1}, \dots, \mathcal{D}_t) - R(\mathcal{D}_1, \dots, \mathcal{D}_t) \right|_{\text{stat}} \right]^2 \\ &\text{(by Jensen's inequality)} \\ &< \frac{\ln 2}{2} \cdot \frac{t'+1}{t}. \end{aligned}$$

Thus, R is $\sqrt{\frac{\ln 2}{2} \cdot \frac{t'+1}{t}}$ -distributionally stable with respect to $\mathcal{D}^1, \dots, \mathcal{D}^t$. This proves item 1 of the Lemma.

For item 2, we apply the alternative bound, Vajda's inequality (Theorem 4.8), to each $j \in [t]$, to find

$$\begin{aligned} D_{\text{KL}}((X^j, \mathbf{R}) \parallel (Y^j, \mathbf{R})) &\geq \frac{1}{\ln 2} \left(\ln \left(\frac{1}{1 - \|(X^j, \mathbf{R}) - (Y^j, \mathbf{R})\|_{\text{stat}}} \right) - 1 \right) \\ &= \frac{1}{\ln 2} \left(\ln \left(\frac{1}{\varepsilon_j} \right) - 1 \right), \end{aligned}$$

where we define

$$\varepsilon_j := 1 - \mathbb{E}_{x^j \sim \mathcal{D}_j} \left[\left| R(\mathcal{D}_1, \dots, \mathcal{D}_{j-1}, x^j, \mathcal{D}_{j+1}, \dots, \mathcal{D}_t) - R(\mathcal{D}_1, \dots, \mathcal{D}_t) \right|_{\text{stat}} \right]$$

and note that $\varepsilon_j > 0$ since the support of (Y^j, \mathbf{R}) contains that of (X^j, \mathbf{R}) . Averaging over $j \in [t]$ and applying Eqs. (6) and (7),

$$\frac{t' + 1}{t} \geq \frac{1}{t} \sum_{j \in [t]} \frac{1}{\ln 2} \left(\ln \left(\frac{1}{\varepsilon_j} \right) - 1 \right),$$

i.e.,

$$\frac{1}{t} \sum_{j \in [t]} \ln \left(\frac{1}{\varepsilon_j} \right) \leq \frac{(\ln 2)(t' + 1)}{t} + 1.$$

The function $f(x) = \ln(1/x)$ has second derivative $x^{-2} > 0$ for $x > 0$, and so Jensen's inequality gives

$$\ln \left(\frac{1}{\frac{1}{t} \sum_{j \in [t]} \varepsilon_j} \right) \leq \frac{(\ln 2)(t' + 1)}{t} + 1.$$

This implies

$$\frac{1}{t} \sum_{j \in [t]} \varepsilon_j \geq \left(e^{\frac{(\ln 2)(t' + 1)}{t} + 1} \right)^{-1} \geq 2^{-\frac{t'}{t} - 3},$$

which proves item 2. □

6.2 Sparsified distributional stability

Here we prove a technical lemma showing that if a mapping F is distributionally stable with respect to i.i.d. inputs, then F also obeys a slightly different stability property, in which we replace an input distribution \mathcal{D} with a “sparsified” version of \mathcal{D} .

Lemma 6.3. *Let U be a finite set, and let $F(x^1, \dots, x^T) : \{0, 1\}^{T \times n} \rightarrow U$ be given. Suppose F is δ -distributionally stable with respect to input distribution $\mathcal{D}^{\otimes T}$, for every distribution \mathcal{D} over $\{0, 1\}^n$.*

Fix some distribution \mathcal{D} over $\{0, 1\}^n$, and let x^1, \dots, x^d be independently sampled from \mathcal{D} . Let $k^ \sim \mathcal{U}_{[d]}$.*

Let $\widehat{\mathcal{D}}$ denote the distribution defined by sampling uniformly from the multiset $\{x^k\}_{k \neq k^}$. (This distribution is itself a random variable, determined by x^1, \dots, x^d and by k^* .) Define*

$$\beta_j := \mathbb{E}_{k^*, x^1, \dots, x^d} \left[\left| F(\widehat{\mathcal{D}}^{\otimes(j-1)}, x^{k^*}, \widehat{\mathcal{D}}^{\otimes(T-j)}) - F(\widehat{\mathcal{D}}^{\otimes T}) \right|_{\text{stat}} \right],$$

where all the $\widehat{\mathcal{D}}$ s are to be mutually independent (for fixed values of x^1, \dots, x^d and k^*). Then,

$$\frac{1}{T} \sum_{j=1}^t \beta_j \leq \delta + 2T/d.$$

Proof. Let $\widetilde{\mathcal{D}}$ denote the distribution, determined by x^1, \dots, x^d , that samples uniformly from the multiset $\{x^k\}_{k \in [d]}$. By an easy calculation, for any values of x^1, \dots, x^d and k^* we can bound

$$\left\| \widetilde{\mathcal{D}} - \widehat{\mathcal{D}} \right\|_{\text{stat}} \leq 1/d.$$

It follows that

$$\left\| F\left(\widetilde{\mathcal{D}}^{\otimes T}\right) - F\left(\widehat{\mathcal{D}}^{\otimes T}\right) \right\|_{\text{stat}} \leq \left\| \widetilde{\mathcal{D}}^{\otimes T} - \widehat{\mathcal{D}}^{\otimes T} \right\|_{\text{stat}} \leq T/d,$$

where in the last step we used Fact 2.5 and the fact that for any assignment to x^1, \dots, x^d and to k^* , the T copies of $\widetilde{\mathcal{D}}$ used are mutually independent, as are the copies of $\widehat{\mathcal{D}}$.

By identical reasoning, for any assignment to x^1, \dots, x^d and to k^* , and for any index $j \in [T]$ we have

$$\left\| F\left(\widetilde{\mathcal{D}}^{\otimes(j-1)}, x^{k^*}, \widetilde{\mathcal{D}}^{\otimes(T-j)}\right) - F\left(\widehat{\mathcal{D}}^{\otimes(j-1)}, x^{k^*}, \widehat{\mathcal{D}}^{\otimes(T-j)}\right) \right\|_{\text{stat}} \leq (T-1)/d.$$

Using the triangle inequality for $\|\cdot\|_{\text{stat}}$, for any values x^1, \dots, x^d, k^* and any index $j \in [T]$ we always have

$$\begin{aligned} & \left\| F\left(\widehat{\mathcal{D}}^{\otimes(j-1)}, x^{k^*}, \widehat{\mathcal{D}}^{\otimes(T-j)}\right) - F\left(\widehat{\mathcal{D}}^{\otimes T}\right) \right\|_{\text{stat}} \\ & \leq \left\| F\left(\widehat{\mathcal{D}}^{\otimes(j-1)}, x^{k^*}, \widehat{\mathcal{D}}^{\otimes(T-j)}\right) - F\left(\widetilde{\mathcal{D}}^{\otimes(j-1)}, x^{k^*}, \widetilde{\mathcal{D}}^{\otimes(T-j)}\right) \right\|_{\text{stat}} \\ & \quad + \left\| F\left(\widetilde{\mathcal{D}}^{\otimes(j-1)}, x^{k^*}, \widetilde{\mathcal{D}}^{\otimes(T-j)}\right) - F\left(\widetilde{\mathcal{D}}^{\otimes T}\right) \right\|_{\text{stat}} \\ & \quad + \left\| F\left(\widetilde{\mathcal{D}}^{\otimes T}\right) - F\left(\widehat{\mathcal{D}}^{\otimes T}\right) \right\|_{\text{stat}} \\ & \leq \left\| F\left(\widetilde{\mathcal{D}}^{\otimes(j-1)}, x^{k^*}, \widetilde{\mathcal{D}}^{\otimes(T-j)}\right) - F\left(\widetilde{\mathcal{D}}^{\otimes T}\right) \right\|_{\text{stat}} + 2T/d. \end{aligned} \tag{8}$$

Now suppose we fix any values x^1, \dots, x^d , leaving k^* undetermined. The value k^* is uniform on $[d]$, so that x^{k^*} is distributed exactly according to $\widehat{\mathcal{D}}$. Under our conditioning, let

$$\gamma_j = \gamma_j\left(\{x^k\}_{k \in [d]}\right) := \mathbb{E}_{k^*} \left[\left\| F\left(\widetilde{\mathcal{D}}^{\otimes(j-1)}, x^{k^*}, \widetilde{\mathcal{D}}^{\otimes(T-j)}\right) - F\left(\widetilde{\mathcal{D}}^{\otimes T}\right) \right\|_{\text{stat}} \right].$$

By our original assumption, F is δ -DS with respect to input distribution $\widetilde{\mathcal{D}}^{\otimes T}$. Thus, for any x^1, \dots, x^d we have

$$\frac{1}{t} \sum_{j=1}^t \gamma_j \leq \delta. \tag{9}$$

Now γ_j is itself a random variable, determined by x^1, \dots, x^d , and from Eq. (8) we have

$$\beta_j \leq \mathbb{E}[\gamma_j] + 2T/d.$$

Using linearity of expectation, we find that

$$\frac{1}{t} \sum_{j=1}^t \beta_j \leq \delta + 2T/d .$$

□

6.3 Building disguising distributions

In the next lemmas we show how the distributional stability of a mapping F can be used to obtain a “disguising distribution” for F . In Lemma 6.6 we will apply this to give disguising distributions for any sufficiently compressive mapping R .

Recall that \mathcal{U}_K denotes the uniform distribution over a multiset K .

Lemma 6.4. *Suppose $F(x^1, \dots, x^T) : \{0, 1\}^{T \times n} \rightarrow U$ obeys the assumption of Lemma 6.3: namely, F is δ -distributionally stable with respect to input distribution $\mathcal{D}^{\otimes T}$, for every distribution \mathcal{D} over $\{0, 1\}^n$.*

Let $S \subseteq \{0, 1\}^n$, and fix some value $d > 0$. There exists a distribution \mathcal{K} over size- d multisets $K \subseteq S$, such that for every $y \in S$, the following holds:

$$\mathbb{E}_{K \sim \mathcal{K}, j^* \sim \mathcal{U}_{[T]}} \left[\left\| F \left(\mathcal{U}_K^{\otimes(j^*-1)}, y, \mathcal{U}_K^{\otimes(T-j^*)} \right) - F \left(\mathcal{U}_K^{\otimes T} \right) \right\|_{\text{stat}} \right] \leq \delta + 2T/(d+1) .$$

(Here the copies of \mathcal{U}_K are to be mutually independent for fixed K , although the set $K \sim \mathcal{K}$ used is the same for each copy.)

Proof. Consider the following two-player, simultaneous-move, zero-sum game:

- **Player 1:** chooses a size- d multiset $K \subseteq S$.
- **Player 2:** chooses a string $y \in S$.
- **Payoff:** Player 2 receives a payoff equal to

$$\mathbb{E}_{j^* \sim \mathcal{U}_{[T]}} \left[\left\| F \left(\mathcal{U}_K^{\otimes(j^*-1)}, y, \mathcal{U}_K^{\otimes(T-j^*)} \right) - F \left(\mathcal{U}_K^{\otimes T} \right) \right\|_{\text{stat}} \right] .$$

(Note that this payoff is a determinate value, given (K, y) .)

Consider any randomized strategy by Player 2, specified by a distribution $y \sim Y$ over S . In response, let \mathcal{K}_Y be the randomized Player-1 strategy that chooses a size- d multiset K of elements sampled independently from Y .

To bound the expected payoff under the strategy-pair (\mathcal{K}_Y, Y) , note that we can equivalently generate $(K, y) \sim (\mathcal{K}_Y, Y)$ as follows. First, sample x^1, \dots, x^{d+1} independently from Y . Sample $k^* \sim \mathcal{U}_{[d+1]}$, set $y := x^{k^*}$, and let

$$K := \{x^1, \dots, x^{k^*-1}, x^{k^*+1}, \dots, x^{d+1}\} .$$

It is easily verified that $(K, y) \sim (\mathcal{K}_Y, Y)$ as desired.

Then Lemma 6.3, applied to our initial distributional-stability assumption on F , informs us that

$$\mathbb{E}_{j^* \sim \mathcal{U}_{[t]}, K, y} \left[\left\| F \left(\mathcal{U}_K^{\otimes(j^*-1)}, y, \mathcal{U}_K^{\otimes(T-j^*)} \right) - F \left(\mathcal{U}_K^{\otimes T} \right) \right\|_{\text{stat}} \right] \leq \delta + 2T/(d+1) .$$

Thus Player 2's expected payoff against \mathcal{K}_Y is at most $\delta + 2T/(d+1)$.

As Y was arbitrary, the minimax theorem tells us that there exists a distribution \mathcal{K} over Player-1 moves that forces Player 2's expected payoff under *every* strategy to be at most $\delta + 2T/(d+1)$. The result follows. \square

Lemma 6.5. *Let U be a finite set, and let $F(x^1, \dots, x^T) : \{0, 1\}^{T \times n} \rightarrow U$ be given. Suppose F is δ -distributionally stable with respect to input distribution $\mathcal{D}^{\otimes T}$, for every distribution \mathcal{D} over $\{0, 1\}^n$.*

Let $S \subseteq \{0, 1\}^n$, and fix $d > 0$. Given any $\varepsilon > 0$, let $s := \lceil (.5 \ln 2)n/\varepsilon^2 \rceil$. Then there exists a collection K_1, \dots, K_s of size- d multisets contained in S , such that for every $y \in S$ the following holds:

$$\mathbb{E}_{a \sim \mathcal{U}_{[s]}, j^* \sim \mathcal{U}_{[t]}} \left[\left\| F \left(\mathcal{U}_{K_a}^{\otimes(j^*-1)}, y, \mathcal{U}_{K_a}^{\otimes(T-j^*)} \right) - F \left(\mathcal{U}_{K_a}^{\otimes T} \right) \right\|_{\text{stat}} \right] \leq \delta + 2T/(d+1) + \varepsilon .$$

Proof. This is an immediate application (to the game in Lemma 6.4) of a general result due to Lipton and Young [LY94, Theorem 2], showing that all two-player, zero-sum games have sparsely-supported, nearly-optimal player strategies. (Essentially the same result was proved independently by Althöfer [Alt94], and a more general result for many-player, non-zero-sum games was proved later in [LMM03].) The support size required in the Lipton-Young-Althöfer result depends logarithmically on the number of pure strategies available to the player we are opposing; in our case, Player 2 has a choice of $|S| \leq 2^n$ strings y , so we get $s = O(n/\varepsilon^2)$. In their proof technique applied to our setting, the K_1, \dots, K_s are obtained by sampling independently from the distribution \mathcal{K} given by Lemma 6.4, giving a successful outcome with nonzero probability. \square

Lemma 6.6 (Disguising-Distribution Lemma). *Let $R(x^1, \dots, x^t) : \{0, 1\}^{t \times n} \rightarrow \{0, 1\}^{\leq t'}$ be any possibly-randomized mapping, for $t, t' \in \mathbb{N}^+$. Let $S \subseteq \{0, 1\}^n$, and fix $d > 0$. Given any $\varepsilon > 0$, let $s := \lceil (.5 \ln 2)n/\varepsilon^2 \rceil$. Let*

$$\widehat{\delta} := \min \left\{ \sqrt{\frac{\ln 2}{2} \cdot \frac{t'+1}{t}}, 1 - 2^{-\frac{t'}{t}-3} \right\} .$$

Then there exists a collection K_1, \dots, K_s of size- d multisets contained in S , such that for every $y \in S$, we have

$$\begin{aligned} \mathbb{E}_{a \sim \mathcal{U}_{[s]}, j^* \sim \mathcal{U}_{[t]}} \left[\left\| R \left(\mathcal{U}_{K_a}^{\otimes(j^*-1)}, y, \mathcal{U}_{K_a}^{\otimes(t-j^*)} \right) - R \left(\mathcal{U}_{K_a}^{\otimes t} \right) \right\|_{\text{stat}} \right] \\ \leq \widehat{\delta} + 2t/(d+1) + \varepsilon . \end{aligned}$$

Proof. This follows immediately from the combination of Lemmas 6.2 and 6.5, applied to $F := R$ (and with $T := t$). \square

7 Limits to efficient (classical) compression

In this section, we show that a sufficiently high-quality PPT-OR-compression reduction for any language L implies that $L \in \text{NP}/\text{poly}$. We also show that above a higher threshold of quality, such a compression reduction implies that L has non-uniform, statistical zero-knowledge proofs, which in particular implies $L \in \text{coNP}/\text{poly}$ as well. We will then apply these results to give evidence against efficient probabilistic compression for AND(SAT) and OR(SAT), as described in the Introduction, and for other parametrized problems with either of the two “expressiveness” properties described in Section 5.2. We will also present our result on f -compression reductions for more general combining functions f , and our result extending the work of Dell and Van Melkebeek [DvM10] on problems with polynomial kernelizations.

7.1 Complexity upper bounds from OR-compression schemes

Theorem 7.1. *Let L be any language. Suppose $t_1(n), t_2(n) : \mathbb{N}^+ \rightarrow \mathbb{N}^+$ are (not necessarily computable) functions. Suppose that there exists a PPT-OR-compression reduction $R(x^1, \dots, x^t) : \{0, 1\}^{t_1(n) \times n} \rightarrow \{0, 1\}^{\leq t_2(n)}$ for L with parameters $t_1(n), t_2(n)$, error bound $\xi(n) < .5$, and some target language L' . Let*

$$\widehat{\delta} := \min \left\{ \sqrt{\frac{\ln 2}{2} \cdot \frac{t_2(n) + 1}{t_1(n)}}, \quad 1 - 2^{-\frac{t_2(n)}{t_1(n)} - 3} \right\}.$$

1. *If for some constant $c > 0$ we have*

$$1 - 2\xi(n) - \widehat{\delta} \geq \frac{1}{n^c}, \quad (10)$$

then $L \in \text{NP}/\text{poly}$.

2. *If for some $c > 0$ we have the (stronger) bound*

$$(1 - 2\xi(n))^2 - \widehat{\delta} \geq \frac{1}{n^c}, \quad (11)$$

then there is a many-to-one reduction from L to a promise problem in pr-SZK. The reduction is computable in non-uniform polynomial time; in particular, this implies $L \in \text{NP}/\text{poly} \cap \text{coNP}/\text{poly}$.

We remark that, using the technique of [FS11, Proposition 5.1], one can reduce the error bound $\xi(n)$ of an OR-compression scheme, at the cost of increasing the output-length bound $t_2(n)$. (The idea is to perform multiple, independent applications of R to the fixed input tuple $(x^1, \dots, x^{t_1(n)})$ and to concatenate the results in the output, using a majority-vote rule to define a new target language.) With this amplification, we can in some cases apply Theorem 7.1 where its assumptions do not hold for the original scheme—or, we may obtain the stronger conclusion in item 2 of Theorem 7.1 in cases where only item 1 would apply directly.

Proof of Theorem 7.1. We will use the same basic reduction to prove items 1 and 2. First, with non-uniformity it is easy to handle length- n inputs whenever $L_n = \{0, 1\}^n$, so let us assume from this point on that \overline{L}_n is nonempty.

Using R , we define a deterministic, non-uniform polynomial-time reduction \mathcal{R} that, on input $y \in \{0, 1\}^n$, builds a description of two circuits C, C' . The aim is that $\|\mathcal{D}_C - \mathcal{D}_{C'}\|_{\text{stat}}$ should be large if $y \in L$, and small if $y \notin L$. \mathcal{R} works as follows:

- **Non-uniform advice for length n :** a description of the value $t_1(n)$, and the multisets $K_1, \dots, K_s \subseteq \bar{L}_n$ given by Lemma 6.5 with

$$(t, t') := (t_1(n), t_2(n)), \quad S := \bar{L}_n, \quad d := \lceil 8t_1(n) \cdot n^c \rceil, \quad \varepsilon := \frac{1}{4n^c}.$$

(Here $c > 0$ is as in Eq. (10) or Eq. (11), according to which item of the Theorem we are proving.) Note that d and the value s given by Lemma 6.5 are both $\leq \text{poly}(n)$ under these settings, so our advice is of polynomial length.

- **On input $y \in \{0, 1\}^n$:** let \mathcal{R} output descriptions $\langle C, C' \rangle$ of the following two randomized circuits:

- **Circuit C :** samples $a \sim \mathcal{U}_{[s]}$, then samples

$$\bar{x} = (x^1, \dots, x^{t_1(n)}) \sim \mathcal{U}_{K_a}^{\otimes t_1(n)},$$

and outputs $z := R(\bar{x})$.

- **Circuit C' :** samples values

$$a \sim \mathcal{U}_{[s]}, \quad j^* \sim \mathcal{U}_{[t_1(n)]};$$

then, samples

$$\bar{x} \sim \left(\mathcal{U}_{K_a}^{\otimes (j^*-1)}, y, \mathcal{U}_{K_a}^{\otimes (t_1(n)-j^*)} \right),$$

and outputs $z := R(\bar{x})$.

Claim 7.2. *The following holds:*

1. If $y \in L$, then

$$\|\mathcal{D}_C - \mathcal{D}_{C'}\|_{\text{stat}} \geq D(n) := 1 - 2\xi(n);$$

2. If $y \notin L$, then

$$\|\mathcal{D}_C - \mathcal{D}_{C'}\|_{\text{stat}} \leq d(n) := \hat{\delta} + \frac{1}{2n^c}. \quad (12)$$

We defer the proof of Claim 7.2, and use it to prove the two items of Theorem 7.1.

For item 1 of Theorem 7.1, if Eq. (10) holds (for sufficiently large n), then $D(n) - d(n) \geq \frac{1}{n^c}$.

Now $D(n), d(n)$ were parametrized in terms of $n = |y|$, but the gap $D(n) - d(n)$ is also at least inverse-polynomial in the length $N \leq \text{poly}(n)$ of the output description $\langle C, C' \rangle$. Thus our reduction \mathcal{R} reduces any instance y of the decision problem for L , to an equivalent instance $\mathcal{R}(y) = \langle C, C' \rangle$ of the promise problem $\text{SD}_{\leq d'(N)}^{\geq D'(N)}$, with different parameters $D'(N), d'(N)$ still satisfying the gap condition $D' - d' \geq \frac{1}{\text{poly}(N)}$.

By item 1 of Theorem 4.19, $\text{SD}_{\leq d'}^{\geq D'} \in \text{pr-NP/poly}$. Let $(A, \{a_N\}_{N>0})$ be a nondeterministic, non-uniform polynomial-time algorithm and advice family solving $\text{SD}_{\leq d'}^{\geq D'}$. Then by applying

$(A, \{a_N\})$ to $\mathcal{R}(y)$, we obtain a nondeterministic, non-uniform polynomial-time algorithm for solving L . This shows $L \in \text{NP/poly}$, proving item 1 of the Theorem.

Next, for item 2 of Theorem 7.1, if Eq. (11) holds for sufficiently large n , then $D(n)^2 - d(n) \geq \frac{1}{n^c}$. Arguing as in the previous case, but this time applying item 2 of Theorem 4.19, we exhibit a nonuniform polynomial-time reduction from L to $\text{SD}_{\leq \frac{D'}{d'}}^{\geq \frac{D'}{d'}}$, where this time $D'(N)^2 - d'(N) \geq \frac{1}{\text{poly}(N)}$. This problem can in turn be reduced to $\text{SD}_{\leq 1/3}^{\geq 2/3} \in \text{pr-SZK}$ in non-uniform polynomial time, by the second assertion of Theorem 4.14. This also yields $L \in \text{NP/poly} \cap \text{coNP/poly}$, and completes the proof of Theorem 7.1. \square

Proof of Claim 7.2. (1.) First, suppose $y \in L$. We will use the distinguishing interpretation of statistical distance (see Section 2.1) to argue that $\|\mathcal{D}_C - \mathcal{D}_{C'}\|_{\text{stat}}$ is large. Suppose an unbiased coin $b \sim \mathcal{U}_{\{0,1\}}$ is flipped, unseen by us, and we receive a sample $z \sim \mathcal{D}_C$ if $b = 0$, or $z \sim \mathcal{D}_{C'}$ if $b = 1$. Consider the distinguisher that outputs the guess $\tilde{b} := 0$ if $z \in \overline{L'}$, or $\tilde{b} := 1$ if $z \in L'$.

We lower-bound the success probability $\Pr[\tilde{b} = b]$ as follows. Say we condition on $[b = 0]$, so that $z \sim \mathcal{D}_C$. The distributions \mathcal{U}_{K_a} are supported on $\overline{L_n}$, so in the execution of C we get $\bar{x} \in (\overline{L_n})^{t_1(n)}$. Then it follows from the OR-compression property of R for L that $\Pr[z \in \overline{L'}] \geq 1 - \xi(n)$. On the other hand, suppose we condition on $[b = 1]$, so that $z \sim \mathcal{D}_{C'}$. In an execution of C' the input tuple \bar{x} contains $y \in L_n$; thus, by the OR-compression property of R , we have $\Pr[z \in L'] \geq 1 - \xi(n)$. So regardless of the value of b , our distinguisher succeeds with probability $\geq 1 - \xi(n)$. Thus, $1 - \xi(n) \leq \frac{1}{2}(1 + \|\mathcal{D}_C - \mathcal{D}_{C'}\|_{\text{stat}})$. This proves item 1.

(2.) Now suppose $y \notin L$; we must upper-bound $\|\mathcal{D}_C - \mathcal{D}_{C'}\|_{\text{stat}}$. Consider the distinguishing experiment between C and C' as in item 1. If we regard the random variables a and j^* (the latter used only by C') to be part of the joint probability space of both algorithms (noting that a is identically distributed in the two circuits), then revealing the values a, j^* along with z to the distinguisher cannot decrease the distinguisher's maximum achievable success probability. Now conditioned on revealed values a, j^* , the maximum achievable success probability in the modified distinguishing experiment is

$$\frac{1}{2} \left(1 + \left\| \left\| R \left(\mathcal{U}_{K_a}^{\otimes t_1(n)} \right) - R \left(\mathcal{U}_{K_a}^{\otimes (j^*-1)}, y, \mathcal{U}_{K_a}^{\otimes (t_1(n)-j^*)} \right) \right\|_{\text{stat}} \right) \right),$$

from which we conclude that

$$\|\mathcal{D}_C - \mathcal{D}_{C'}\|_{\text{stat}} \leq \mathbb{E}_{a \sim \mathcal{U}_{[s]}, j^* \sim \mathcal{U}_{[t_1(n)]}} \left[\left\| \left\| R \left(\mathcal{U}_{K_a}^{\otimes t_1(n)} \right) - R \left(\mathcal{U}_{K_a}^{\otimes (j^*-1)}, y, \mathcal{U}_{K_a}^{\otimes (t_1(n)-j^*)} \right) \right\|_{\text{stat}} \right] . \quad (13)$$

By our choice of K_1, \dots, K_s and Lemma 6.6, the right-hand side of Eq. (13) is at most

$$\widehat{\delta} + 2t_1(n)/(d+1) + \varepsilon < \widehat{\delta} + 2 \cdot \frac{1}{4n^c}, \quad (14)$$

by our settings to d, ε . This proves Eq. (12) and completes the proof of Claim 7.2. \square

The next result gives a useful consequence of Theorem 7.1 for the case where the compression bound $t_2(n)$ is on the order of $t_1(n) \cdot \log_2(t_1(n))$, and also points out a strengthening of the result's conclusion in the case of error-free compression.

Theorem 7.3. *Let L be any language. Suppose $t_1(n), t_2(n) : \mathbb{N}^+ \rightarrow \mathbb{N}^+$ satisfy $t_2(n) \leq C \cdot t_1(n) \log t_1(n)$ and $t_1(n) \leq n^{C'}$, for some $C, C' > 0$. Suppose that R is a PPT-OR-compression reduction $R(x^1, \dots, x^{t_1(n)}) : \{0, 1\}^{t_1(n) \times n} \rightarrow \{0, 1\}^{\leq t_2(n)}$ for L with parameters $t_1(n), t_2(n)$, error bound $\xi(n) < .5$, and some target language L' .*

1. *If $\xi(n) < n^{-C \cdot C'} / 32$, then there is a non-uniform polynomial-time many-to-one reduction from L to a promise problem in pr-SZK.*
2. *Suppose further that R is error-free (i.e., $\xi(n) = 0$). Then, there is a non-uniform polynomial-time many-to-one reduction from L to a promise problem in pr-PZK.*

Proof. (1.) We bound the quantity $\widehat{\delta}$ from Theorem 7.1:

$$\begin{aligned} \widehat{\delta} &\leq 1 - 2^{-\frac{t_2(n)}{t_1(n)} - 3} \\ &\leq 1 - 2^{-C \log_2(t_1(n))} / 8 \\ &\leq 1 - t_1(n)^{-C} / 8 \\ &\leq 1 - n^{-C \cdot C'} / 8. \end{aligned}$$

If $\xi(n) < n^{-C \cdot C'} / 32$, then the left-hand quantity in Eq. (11) is $\geq \frac{1}{\text{poly}(n)}$, and the desired conclusion then follows from Theorem 7.1, item 2.

(2.) Looking into the proof of Theorem 7.1, item 2, we see that it gives a non-uniform polynomial-time many-to-one reduction from L to $\text{SD}_{\leq d'(N)}^{\geq D'(N)}$, where in the current case, using Claim 7.2, we have

$$D'(N) = 1, \quad d'(N) \leq 1 - \frac{1}{\text{poly}(N)}.$$

This problem can in turn be uniformly many-to-one reduced to $\text{SD}_{\leq .5}^{\geq 1}$ by mapping a circuit-distribution pair $\langle C, C' \rangle$ to $\langle C^{\otimes T}, (C')^{\otimes T} \rangle$, where $C^{\otimes T}$ is the circuit that outputs T samples drawn independently from C , and where $T \leq \text{poly}(n)$ is chosen suitably large. Finally, $\text{SD}_{\leq .5}^{\geq 1} \in \text{pr-PZK}$ by Theorem 4.18. \square

7.2 Application to AND- and OR-compression of NP-complete languages

Throughout this section, for parameters $t_1(n), t_2(n)$, we will use the shorthand

$$\widehat{\delta} := \min \left\{ \sqrt{\frac{\ln 2}{2} \cdot \frac{t_2(n) + 1}{t_1(n)}}, \quad 1 - 2^{-\frac{t_2(n)}{t_1(n)} - 3} \right\}$$

Here is our first main result giving evidence against efficient AND-compression for NP-complete languages:

Theorem 7.4. *Suppose that for some NP-complete language L , any target language L' , and an error bound $\xi(n) < .5$, L has a PPT-AND-compression reduction R with target language L' , with parameters $t_1(n), t_2(n) : \mathbb{N}^+ \rightarrow \mathbb{N}^+$ and error bound $\xi(n) < .5$.*

1. If

$$1 - 2\xi(n) - \widehat{\delta} \geq \frac{1}{\text{poly}(n)}, \quad (15)$$

then $\text{NP} \subseteq \text{coNP}/\text{poly}$ and $\text{PH} = \Sigma_3^p = \Pi_3^p$.

2. If we have the bound

$$(1 - 2\xi(n))^2 - \widehat{\delta} \geq \frac{1}{\text{poly}(n)}, \quad (16)$$

then L (and every other language in NP) is many-to-one reducible in non-uniform polynomial time to a problem in pr-SZK , and $\text{NP} \subseteq \text{coNP}/\text{poly}$.

3. The conclusion of item 2 holds if $t_2(n) \leq C \cdot t_1(n) \log(t_1(n))$ and if $\xi(n)$ is a sufficiently small inverse-polynomial function of n (determined by t_1 and the constant C).

Item 3 above establishes the assertion of Theorem 1.3 from the Introduction for the case of AND-compression.

Proof of Theorem 7.4. (1.) The reduction R is also a PPT-OR-compression for \bar{L} , with target language \bar{L}' , and with the same parameters.

If Eq. (15) holds in case 1, we apply item 1 of Theorem 7.1 to \bar{L} , concluding that $\bar{L} \in \text{NP}/\text{poly}$, i.e., $L \in \text{coNP}/\text{poly}$. The consequence for PH is from Theorem 4.9.

(2.) Similarly, if Eq. (16) holds in case 2, we apply item 2 of of Theorem 7.1 to \bar{L} , giving a non-uniform many-to-one reduction from \bar{L} to a problem $\Pi = (\Pi_Y, \Pi_N) \in \text{pr-SZK}$. This is also a reduction from L to (Π_N, Π_Y) , which by Theorem 4.14 also lies in pr-SZK . The extension to other languages in NP follows from the NP -completeness of L .

(3.) In this case we apply Theorem 7.3, item 1 to \bar{L} . □

The next theorem gives evidence for the infeasibility of efficient OR-compression for NP -complete languages.

Theorem 7.5. 1. Suppose that for some NP -complete language L , any target language L' , and an error bound $\xi(n) < .5$, L has a PPT-OR-compression reduction R with target language L' , with parameters $t_1(n), t_2(n) : \mathbb{N}^+ \rightarrow \mathbb{N}^+$ and error bound $\xi(n) < .5$. If

$$(1 - 2\xi(n))^2 - \widehat{\delta} \geq \frac{1}{\text{poly}(n)}, \quad (17)$$

then L (and every other language in NP) is reducible in non-uniform polynomial time to a problem in pr-SZK , and $\text{NP} \subseteq \text{coNP}/\text{poly}$.

2. The conclusion of item 1 holds if $t_2(n) \leq C \cdot t_1(n) \log(t_1(n))$ and if $\xi(n)$ is a sufficiently small inverse-polynomial function of n (determined by t_1 and C).

Item 2 completes the proof of Theorem 1.3 from the Introduction.

Proof of Theorem 7.5. (1.) This time, if Eq. (17) holds, we just apply item 2 of Theorem 7.1 to L itself.

(2.) In this case we apply item 1 of Theorem 7.3 to L . □

7.3 f -compression of NP-complete languages for general f

As discussed in Section 1.3.1, our techniques, combined with ideas of [FS11, Section 7], imply some limitations to efficient strong f -compression of SAT or other NP-complete languages, for any “reasonable” combining function f .

Theorem 7.6. *Suppose $f : \{0, 1\}^* \rightarrow \{0, 1\}$ depends on all coordinates, and suppose that R is a strong f -compression reduction for an NP-complete language L with a target language $L' \in \text{NP}$, computable in deterministic³⁰ polynomial time. Then $\text{NP} \subseteq \text{coNP/poly}$.*

Proof sketch. Suppose the strong f -compression reduction maps inputs $(x^1, \dots, x^t) \in \{0, 1\}^{t \times n}$ to an output z of length at most n^c (independent of t); we have

$$z \in L' \iff f(L(x^1), \dots, L(x^t)) = 1 .$$

Fix an input length n , and let $N := n^{2c+1}$. If f is non-monotone on inputs of size N , say with respect to the first coordinate, then we can non-uniformly fix some $N-1$ strings $x^2, \dots, x^N \in \{0, 1\}^n$ such that for $x \in \{0, 1\}^n$,

$$R(x, x^2, \dots, x^N) \in L' \iff L(x) = 0 .$$

This reduces the membership question for \bar{L}_n to a membership question for $L' \in \text{NP}$.

Otherwise, f is monotone for length- N inputs. Then by results of Nisan [Nis91], we can “embed” an AND or OR of size \sqrt{N} into f_N . That is, there exists some subset $S \subseteq [N]$ of size at least $\sqrt{N} = n^{c+.5}$, and an assignment to all inputs to f outside of S , such that the restricted function is an AND or OR of its remaining coordinates. Thus, there exists a collection $(x^\ell)_{\ell \notin S}$ of length- n inputs to R , such that for any setting to the remaining (length- n) inputs with indices in S ,

$$R(x^1, x^2, \dots, x^N) \in L' \iff \text{OP}_{\ell \in S} (L(x^\ell)) = 1 ,$$

where here OP is either AND or OR. Thus for this input length n , we have (non-uniformly) either an AND-compression or OR-compression reduction from L to target L' , compressing by a polynomial amount. In either case, we can apply the techniques of Section 7.2 to get a non-uniform proof system for membership in \bar{L}_n . Combining our work for each input length, we find that $\bar{L} \in \text{NP/poly}$. As L is NP-complete, the Theorem is proved. \square

7.4 Limits to strong compression for parametrized problems

Next, we use Theorem 7.1 to give evidence against strong compressibility for “expressive” parametrized problems. The result we give below is a simple-to-state, representative example; the quantitative settings studied here are not the only interesting ones our techniques can handle.

Theorem 7.7. *Say that P is OR-expressive or AND-expressive, e.g., one of the problems listed in Theorems 5.5 and 5.6. Suppose additionally that P is strongly PPT-compressible³¹ with error bound $\xi(m, k, w)$ satisfying $\xi(m, k, 1) \leq .5 - k^{-O(1)}$ (independent of m), i.e., with success probability $\geq .5 + k^{-O(1)}$. Then, every language in NP is many-to-one reducible in non-uniform polynomial time to a problem in pr-SZK (and $\text{NP} \subseteq \text{coNP/poly}$).*

³⁰(this assumption is for simplicity)

³¹(as in Definition 5.8)

Theorem 1.2 from the Introduction follows, by considering the special cases $P = \text{OR}(\text{SAT})$ and $P = \text{AND}(\text{SAT})$.

Proof of Theorem 7.7. Suppose first that P is OR-expressive, with respect to the NP-complete language L and with some parameter $S(n) \leq \text{poly}(n)$. We apply item 1 of Lemma 5.10 to L and the assumed strong compression reduction for P . Using some function $T(n) \leq \text{poly}(n)$ to be determined, and with $w(n) := 1$, we obtain a PPT-OR-compression for L with parameters

$$t_1(n) = T(n), \quad t_2(n) \leq S(n) \cdot n^{O(1)}, \quad \xi'(n) \leq 2^{-n}.$$

(Here, the bound on t_2 is independent of the choice of $T(n)$.) We evaluate

$$(1 - 2\xi'(n))^2 - \sqrt{\frac{\ln 2}{2} \cdot \frac{t_2(n) + 1}{t_1(n)}} \geq (1 - 4 \cdot 2^{-n}) - \sqrt{\frac{\ln 2}{2} \cdot \frac{S(n) \cdot n^{O(1)} + 1}{T(n)}},$$

for some $a > 0$ (using $S(n) \leq \text{poly}(n)$). The expression above can be made greater than .5 for large n by choosing a sufficiently fast-growing $T(n) \leq \text{poly}(n)$. Under such a setting, Eq (17) holds for $(t_1(n), t_2(n), \xi'(n))$. We can then apply the first assertion of Theorem 7.5, item 1 to our PPT-OR-compression for L , which yields the desired conclusion.

The case where P is AND-expressive is handled analogously; in this case we apply Lemma 5.10, item 2 and the first assertion of Theorem 7.4, item 2. \square

We can also apply Theorem 7.3 to show that, if any NP-complete language L is PPT-OR-compressible by an error-free reduction with $t_2(n) = O(t_1(n) \log(t_1(n)))$, then NP has non-uniform *perfect* zero-knowledge proofs. From a deterministic AND-compression reduction for L of this type, we get non-uniform perfect zero-knowledge proofs for coNP. (Note that unlike pr-SZK, pr-PZK is not known to be closed under complement.)

7.5 Application to problems with polynomial kernelizations

In this section we prove new limits to efficient compression for the Satisfiability problem on d -CNFs, and for some problems on graphs and hypergraphs, partially extending results of Dell and Van Melkebeek [DvM10] to handle two-sided error. First, we need some background.

Definition 7.8 (Hypergraphs, vertex covers, and cliques). *For any integer $d \geq 2$, a d -uniform hypergraph, or d -hypergraph, is a set H of size- d subsets of a vertex set $V = [N]$. A vertex cover in a d -uniform hypergraph H is a subset of vertices that intersects all hyperedges in H . A subset $V' \subseteq V$ is a clique in H if every size- d subset of V' is a member of H .*

Clearly H has a vertex cover of size s exactly if the “complement” hypergraph $\overline{H} := \{e : |e| = d \wedge e \notin H\}$ contains a clique of size $N - s$.

Definition 7.9. *Define the parametrized problems*

d -Vertex Cover $:= \{\langle (H, s), 1^N \rangle : H \text{ is a } d\text{-hypergraph on } [N] \text{ and contains a vertex cover of size } s\}$,

d -Clique $:= \{\langle (H, s), 1^N \rangle : H \text{ is a } d\text{-hypergraph on } [N] \text{ and contains a clique of size } s\}$.³²

Also define the parametrized d -CNF Satisfiability problem

$$d\text{-SAT}_{\text{par}} := \{\langle \psi, 1^N \rangle : \psi \text{ is a satisfiable } d\text{-CNF on } N \text{ variables}\}.$$

³²This is a different parametrized problem than the two clique-based problems mentioned in Section 5.2.

We will prove new limits on efficient compression for these problems with the help of the following powerful, ingenious reduction of Dell and Van Melkebeek.

Theorem 7.10 ([DvM10], Lemma 2). *Fix $d \geq 2$, and let $T(n) : \mathbb{N}^+ \rightarrow \mathbb{N}^+$ be polynomially bounded. There is a deterministic polynomial-time OR-compression reduction³³ R^* for $L = 3\text{-SAT}$,³⁴ with target language $L' = d\text{-Clique}$. For the first parameter we have $t_1(n) = T(n)$. The $d\text{-Clique}$ instance $\langle (H, s), 1^N \rangle$ output by R^* satisfies*

$$N = O\left(n \cdot \max\left(n, T(n)^{1/d+o(1)}\right)\right).$$

By straightforwardly combining Theorem 7.10 with our Theorem 7.4, we will prove the following theorem:

Theorem 7.11. *Let $d \geq 2, \varepsilon > 0$ be given. There is a $\beta = \beta(d, \varepsilon) > 0$ for which the following holds. Suppose that $d\text{-Clique}$ has a polynomial-time compression reduction with output-size bound $O(N^{d-\varepsilon})$ and success probability $.5 + N^{-\beta}$; that is (in the terms of Definition 5.7), suppose that $d\text{-Clique}$ is PPT-compressible with parameters c, ξ satisfying*

$$c(M, N, 1) \leq O(N^{d-\varepsilon}), \quad \xi(M, N, 1) \leq .5 - N^{-\beta},$$

with any target language L' .

Then, every language in NP is many-to-one reducible in non-uniform polynomial time to a problem in pr-SZK (and $\text{NP} \subseteq \text{coNP/poly}$).

The same result holds if we replace $d\text{-Clique}$ with $d\text{-Vertex Cover}$ or $d\text{-SAT}_{\text{par}}$.

Theorem 7.11 gives a version of [DvM10, Theorems 1 and 2] that applies to probabilistic reductions with two-sided error. However, our result does not apply to the more general setting of *oracle communication protocols*, to which those earlier results do apply (for co-nondeterministic protocols, and protocols avoiding false negatives).

Dell and Van Melkebeek use their techniques to show compression lower bounds for several other interesting graph problems (including the Feedback Vertex Set, Bounded-Degree Deletion, and Non-Planar Deletion problems) via reductions from 2-Vertex Cover [DvM10, Section 5.2]. Using our results and the reductions in [DvM10], one can also obtain similarly strong compression lower bounds for these problems for the two-sided error setting.

Proof of Theorem 7.11. We already described a simple reduction (in both directions) between the $d\text{-Vertex Cover}$ and $d\text{-Clique}$ problems that preserves the parameter N . Also, an instance of $d\text{-Vertex Cover}$ on N vertices is efficiently reducible to a $d\text{-SAT}$ instance over $O(N)$ variables [DvM10, Lemma 5]. Thus, it suffices to prove the result for $d\text{-Clique}$.

Let R be the compression reduction assumed to exist for $d\text{-Clique}$, with the value $\beta > 0$ to be determined later. Let $C > d$ be a large integer value, also to be determined.

We will define an OR-compression reduction R' for $L = 3\text{-SAT}$ and target language L' from our assumption; this will allow us to apply Theorem 7.5. R' works as follows. We let $t_1(n) := n^C$. On input formulas $\psi_1, \dots, \psi_{n^C}$, each of bit-length n , the reduction first computes $\langle (H, s), 1^N \rangle := R^*(\psi_1, \dots, \psi_{n^C})$, where R^* is as in Theorem 7.10. Next, R' outputs the value $z := R(\langle (H, s), 1^N \rangle)$.

³³(as in Definition 4.20)

³⁴Here 3-SAT is just the usual language $\{\langle \psi \rangle : \psi \text{ is a satisfiable 3-CNF}\}$.

R' is clearly polynomial-time computable. To analyze R' , fix length- n formulas $\psi_1, \dots, \psi_{n^C}$, and let

$$b := \bigvee_{j=1}^{n^C} [\psi_j \in \text{3-SAT}] .$$

By the OR-compression property of the deterministic mapping R^* , we have

$$[b = 1] \iff \langle (H, s), 1^N \rangle \in d\text{-Clique} .$$

Then by the assumed reliability guarantee of R ,

$$\begin{aligned} \Pr[L'(z) = b] &\geq .5 + N^{-\beta} \\ &\geq .5 + \left(O \left(n \cdot \max \left(n, n^{C/d+o(1)} \right) \right) \right)^{-\beta} \\ &\geq .5 + n^{-\beta(1+C/d)+o(1)} . \end{aligned}$$

Thus the error bound $\xi(n)$ of our reduction R' is at most $.5 - n^{-\beta(1+C/d)+o(1)}$. Also, by the compression guarantee of R , the output z satisfies

$$\begin{aligned} |z| &\leq O(N^{d-\varepsilon}) \\ &\leq O \left(\left(n^{1+C/d+o(1)} \right)^{d-\varepsilon} \right) \\ &\leq O \left(n^{C-1+o(1)} \right) , \end{aligned}$$

with the last step valid provided we take $C > d(d+1)/\varepsilon$. Thus as an output-size bound for R' , we may take $t_2(n) = O(n^{C-1+o(1)})$. We evaluate

$$\begin{aligned} (1 - 2\xi(n))^2 - \sqrt{\frac{\ln 2}{2} \cdot \frac{t_2(n) + 1}{t_1(n)}} &\geq 4n^{-2(\beta(1+C/d)-o(1))} - O(n^{-.5+o(1)}) \\ &\geq n^{-\Omega(1)} , \end{aligned}$$

provided we take $\beta < .25(1+C/d)^{-1}$. Thus under these settings, Eq. (17) holds. Then Theorem 7.5, item 1 gives the desired conclusion, since $L = \text{3-SAT}$ is NP-complete. \square

8 Extension to quantum compression

In this section we will show that our results on OR- and AND-compression have analogues for the model in which the compression scheme is allowed to be a quantum algorithm, outputting a quantum state.

We assume familiarity with the basics of quantum computing and quantum information (for the needed background, consult [NC00]). However, readers without this background should be able to follow the overall structure of the argument if they are willing to regard “qubits,” “quantum operations” “quantum algorithms,” and “quantum measurements” as certain types of black-box objects, and accept some known facts about them. In particular, a “mixed state on m qubits” is a “quantum superposition” over classical m -bit strings. Let

$$\text{MS}_m$$

denote the collection of m -qubit mixed states. (MS_m can be identified with the set of 2^m -by- 2^m , trace-1, positive-semidefinite complex matrices.)

A “quantum operation” is a certain type of mapping $OP : \text{MS}_m \rightarrow \text{MS}_{m'}$, for some $m, m' > 0$. (The operations allowed by quantum physics are the *completely positive, trace-preserving (CPTP) maps*; these are a subset of the linear transformations mapping $\text{MS}_m \subset \mathbb{C}^{m \times m}$ into $\text{MS}_{m'} \subset \mathbb{C}^{m' \times m'}$.) We let

$$\text{OP}_{m,m'}$$

denote the valid quantum operations from m -qubit into m' -qubit states.

“Quantum measurements” are measurements performed on quantum states to yield information about these states; in the quantum setting, measurements are inherently probabilistic, and alter the states being measured. See [NC00, Chapter 2] for a formal definition. Quantum states turn out to inherit some of the information-theoretic limitations of their classical counterparts; this fact will be the basis for our results on quantum compression.

8.1 Trace distance and distinguishability of quantum states

The *trace distance* is a metric on mixed quantum states from a shared state space [NC00]; we denote the trace distance between $\rho, \rho' \in \text{MS}_m$ by $\|\rho - \rho'\|_{\text{tr}} \in [0, 1]$. Formally, treating ρ, ρ' as matrices,

$$\|\rho - \rho'\|_{\text{tr}} := \frac{1}{2} \text{Tr} \left[\sqrt{(\rho - \rho')^2} \right] .$$

This distance is intimately related to the distinguishability of quantum states. Suppose ρ, ρ' are two known states, and we are sent one or the other, each with equal probability (depending on the outcome of an unbiased coin flip $b \in \{0, 1\}$). We want to guess b , by applying some series of quantum operations and measurements. For any ρ, ρ' , it is known [NC00, Theorem 9.1] that our success probability at this task is maximized by using a single binary measurement,³⁵ depending on ρ, ρ' , and that our maximum achievable success probability equals

$$\frac{1}{2} (1 + \|\rho - \rho'\|_{\text{tr}}) .$$

A probability distribution over mixed states is again a mixed state. Thus for a distribution \mathcal{D} over a finite universe U and a mapping $R : U \rightarrow \text{MS}_m$, $R(\mathcal{D})$ defines a quantum state. We use the following standard claim concerning such states, which follows from the distinguishability characterization of $\|\cdot\|_{\text{tr}}$:

Claim 8.1. *For any distributions $\mathcal{D}, \mathcal{D}'$ over a shared finite universe U , and any mapping $R : U \rightarrow \text{MS}_m$, we have*

$$\|R(\mathcal{D}) - R(\mathcal{D}')\|_{\text{tr}} \leq \|\mathcal{D} - \mathcal{D}'\|_{\text{stat}} .$$

Similarly, for any valid quantum operation $OP \in \text{OP}_{m,m'}$ and states $\rho, \rho' \in \text{MS}_m$, we have

$$\|OP(\rho) - OP(\rho')\|_{\text{tr}} \leq \|\rho - \rho'\|_{\text{tr}} .$$

³⁵(i.e., a measurement with two possible outcomes)

8.2 Quantum f -compression

The following notion of quantum compression is modeled on Definition 4.20. The definition is made slightly more complicated by the fact that we no longer have the notion of a “target language” for our reduction; instead, we will require that the answer to our original instance of the decision problem $f \circ L$ be recoverable by some quantum measurement performed on the output state. (This measurement need not be efficiently performable, however.)

Definition 8.2 (Quantum f -compression reductions). *Let L be a language, and let $f : \{0, 1\}^* \rightarrow \{0, 1\}$ be a Boolean function. Let $t_1(n), t_2(n) : \mathbb{N}^+ \rightarrow \mathbb{N}^+$ and $\xi(n) : \mathbb{N}^+ \rightarrow [0, 1]$ be given.*

A quantum f -compression reduction for L , with parameters $t_1(n), t_2(n), \xi(n)$, is a mapping $R(x^1, \dots, x^m)$ outputting a mixed state ρ . There must also exist a family of (not necessarily efficiently-performable) binary quantum measurements $\{\mathcal{M}_n\}_{n>0}$ on $t_2(n)$ -qubit states. We require the following properties: for all $(x^1, \dots, x^{t_1(n)}) \in \{0, 1\}^{t_1(n) \times n}$,

1. *The state $\rho = R(x^1, \dots, x^{t_1(n)})$ is on $t_2(n)$ qubits;*

2. *We have*

$$\Pr \left[\mathcal{M}_n(\rho) = f \left(L(x^1), \dots, L(x^{t_1(n)}) \right) \right] \geq 1 - \xi(n).$$

If some reduction R as above is computable in quantum polynomial time, we say that L is QPT- f -compressible with parameters $(t_1(n), t_2(n), \xi(n))$.

8.3 Quantum complexity classes

We will be using the class QIP[k] of languages definable by k -message, quantum interactive proof systems [Wat03]. Our treatment of these proof systems will be informal, since all the technical properties we need are summarized in theorems from prior work (for details see [Wat03, Wat02]). These are proof systems in which a computationally-unbounded Prover exchanges quantum messages with a quantum polynomial-time Verifier; a total of $k = k(n)$ messages are exchanged. Verifier sends the first message if k is even, or Prover if k is odd, and the parties alternate thereafter. We take $\text{QIP} := \bigcup_{c>0} \text{QIP}[n^c]$.

It was shown in [Wat03, KW00] that for any $3 \leq k(n) \leq \text{poly}(n)$, $\text{PSPACE} \subseteq \text{QIP}[k(n)] = \text{QIP}[3]$; the latter class was recently shown to equal PSPACE [JJUW11]. Importantly for us, however, the class QIP[2] is not known to contain even coNP . The power of 3-message quantum proof systems is in contrast to the classical (private-coin) interactive-proof classes IP[$k(n)$], where for any constant $k \geq 2$, $\text{IP}[k] = \text{IP}[2] = \text{AM}$, and the latter class is believed to be much weaker than $\text{IP}[\text{poly}(n)] = \text{PSPACE}$.

In what follows, we will actually find it more convenient to work with the promise-problem classes $\text{pr-QIP}[k]$.³⁶ The results we’ve summarized carry over to the promise setting as well.

A model of quantum statistical zero-knowledge proofs was proposed by Watrous [Wat02], and used to define the class QSZK of promise problems having polynomial-time proof systems of this type.³⁷ We will use pr-QSZK to denote this class. Watrous showed in [Wat02] that Sahai and Vadhan’s “statistical distance characterization” of pr-SZK , embodied in Definition 4.13, has a

³⁶This is to avoid having to define non-uniform versions of these classes, just as we avoided defining non-uniform versions of AM and SZK.

³⁷Watrous’s original model was of *honest-verifier* quantum statistical zero-knowledge proof systems; he later showed that these proof systems are equivalent in power to “cheating-verifier” ones [Wat09].

quantum analogue. First, we need a promise problem involving trace distance. For a quantum circuit C with an m -qubit output register, let ρ_C denote the output state of C on some fixed input state (say, the all-zeros state). We consider circuits built from a fixed, finite “universal” gate-set (see [NC00, Chapter 4]).

Definition 8.3. For parameters $0 \leq d \leq D \leq 1$, define the promise problem $\text{TD}_{\leq d}^{\geq D} = (\Pi_Y, \Pi_N)$ as follows:

$$\begin{aligned}\Pi_Y &:= \{ \langle C, C' \rangle : \|\rho_C - \rho_{C'}\|_{\text{tr}} \geq D \} , \\ \Pi_N &:= \{ \langle C, C' \rangle : \|\rho_C - \rho_{C'}\|_{\text{tr}} \leq d \} .\end{aligned}$$

In this definition, both $d = d(n)$ and $D = D(n)$ may be parameters depending on the input length $n = |\langle C, C' \rangle|$. (Here, the input description is a classical bit-string.)

Then, appealing to the result of [Wat02], we can use the following definition.

Definition 8.4. Let pr-QSZK be defined as the class of promise problems for which there is a many-to-one (classical, deterministic) polynomial-time reduction from Π to $\text{TD}_{\leq 1/3}^{\geq 2/3}$.

Theorem 8.5 ([Wat02]). pr-QSZK is closed under complement.

Theorem 8.6 ([Wat02]). $\text{pr-QSZK} \subseteq \text{pr-QIP}[2] \cap \text{pr-coQIP}[2]$.

For upper bounds on the complexity of $\text{TD}_{\leq d(n)}^{\geq D(n)}$, we have the following two results, analogous to Theorems 4.15 and 4.14.

Theorem 8.7 (Follows from [Wat02]). Suppose $0 \leq d = d(n) < D = D(n) \leq 1$ are polynomial-time computable, and satisfy $D > d + \frac{1}{\text{poly}(n)}$. Then, $\text{TD}_{\leq d}^{\geq D} \in \text{pr-QIP}[2]$.

If we drop the requirement that d, D be computable, but keep the gap requirement, then $\text{TD}_{\leq d}^{\geq D}$ is many-to-one reducible in non-uniform (classical, deterministic) polynomial time to a problem in $\text{pr-QIP}[2]$.

Theorem 8.7 follows from a “distinguishing protocol” analogous to that in Theorem 4.15.³⁸ Unlike the classical case, there is no known “non-uniform derandomization” result known for $\text{QIP}[2]$ (or for other quantum classes). However, we do have a satisfying analogue of Theorem 4.14:

Theorem 8.8 (Follows from [Wat02]). Suppose $0 \leq d = d(n) < D = D(n) \leq 1$ are polynomial-time computable, and satisfy $D^2 > d + \frac{1}{\text{poly}(n)}$. Then, $\text{TD}_{\leq d}^{\geq D} \in \text{pr-QSZK}$.

If we drop the requirement that d, D be computable, but keep the gap requirement, then $\text{TD}_{\leq d}^{\geq D}$ is many-to-one reducible in non-uniform (classical, deterministic) polynomial time to a problem in pr-QSZK .

³⁸In [Wat02] only the case where D, d are constants is studied, but the result extends easily to when they are functions of n . Also, the second case, where we merely have the gap requirement, is not explicitly analyzed, but follows by a trivial modification of the proof of [Wat02, Theorem 4].

8.4 Quantum distributional stability

We will use a quantum analogue of the distributional stability property:

Definition 8.9. Let $t, t', n \in \mathbb{N}^+$. Given a mapping $F : \{0, 1\}^{t \times n} \rightarrow \text{MS}_{t'}$, and a collection $\mathcal{D}_1, \dots, \mathcal{D}_T$ of mutually independent distributions over $\{0, 1\}^n$, for $j \in [t]$ let

$$\gamma_j := \mathbb{E}_{y \sim \mathcal{D}_j} \left[\left\| F(\mathcal{D}_1, \dots, \mathcal{D}_{j-1}, y, \mathcal{D}_{j+1}, \dots, \mathcal{D}_t) - F(\mathcal{D}_1, \dots, \mathcal{D}_t) \right\|_{\text{tr}} \right].$$

For $\delta \in [0, 1]$, say that F is δ -quantumly-distributionally stable (or δ -QDS) with respect to $\mathcal{D}_1, \dots, \mathcal{D}_t$ if

$$\frac{1}{t} \sum_{j=1}^t \gamma_j \leq \delta.$$

The next lemma is analogous to Lemma B.4.

Lemma 8.10. Let $t, t', n \in \mathbb{N}^+$. Let $F : \{0, 1\}^{t \times n} \rightarrow \text{MS}_{t'}$ be given.

Then, R is δ -QDS with respect to any input distributions $\mathcal{D}_1, \dots, \mathcal{D}_t$, where we may take either of the bounds

1. $\delta := \sqrt{\frac{\ln 2}{2} \cdot \frac{t'}{t}}$;
2. $\delta := 1 - 2^{-\frac{t'}{t}-2}$.

The slight improvement in the bounds comes from the fact that R outputs *exactly* t' qubits. The proof of Lemma 8.10 is very similar to that of Lemma 6.2, and is described in Appendix C.

8.5 Building quantum disguising distributions

Next we prove quantum analogues of our Disguising-Distribution Lemmas. First, we have the following analogue of Lemma 6.3:

Lemma 8.11. Let $t, t' \in \mathbb{N}^+$, and let $F(x^1, \dots, x^T) : \{0, 1\}^{T \times n} \rightarrow \text{MS}_{t'}$ be given. Suppose F is δ -QDS with respect to input distribution $\mathcal{D}^{\otimes T}$, for every distribution \mathcal{D} over $\{0, 1\}^n$.

Fix some distribution \mathcal{D} over $\{0, 1\}^n$, and let x^1, \dots, x^d be independently sampled from \mathcal{D} . Let $k^* \sim \mathcal{U}_{[d]}$. Let $\widehat{\mathcal{D}}$ denote the distribution defined by sampling uniformly from the multiset $\{x^k\}_{k \neq k^*}$. Define

$$\beta_j := \mathbb{E}_{k^*, x^1, \dots, x^d} \left[\left\| R\left(\widehat{\mathcal{D}}^{\otimes(j-1)}, x^{k^*}, \widehat{\mathcal{D}}^{\otimes(T-j)}\right) - R\left(\widehat{\mathcal{D}}^{\otimes T}\right) \right\|_{\text{tr}} \right],$$

where all the $\widehat{\mathcal{D}}$ s are to be mutually independent (for fixed values of x^1, \dots, x^k and k^*). Then,

$$\frac{1}{T} \sum_{j=1}^t \beta_j \leq \delta + 2T/d.$$

Proof. The proof is identical to that of Lemma 6.3, except that we replace statistical distance with trace distance³⁹ and appeal to Claim 8.1 to argue that applying R does not increase trace distance between states. \square

³⁹(where appropriate—the input distributions we manipulate still are to be compared in statistical distance)

After establishing a quantum analogue of Lemma 6.4, we have:

Lemma 8.12. *Suppose F obeys the assumptions of Lemma 8.11. Let $S \subseteq \{0, 1\}^n$, and fix $d > 0$. Given any $\varepsilon > 0$, let $s := \lceil (.5 \ln 2)n/\varepsilon^2 \rceil$. Then there exists a collection K_1, \dots, K_s of size- d multisets contained in S , such that for every $y \in S$ the following holds:*

$$\mathbb{E}_{a \sim \mathcal{U}_{[s]}, j^* \sim \mathcal{U}_{[t]}} \left[\left\| F \left(\mathcal{U}_{K_a}^{\otimes(j^*-1)}, y, \mathcal{U}_{K_a}^{\otimes(T-j^*)} \right) - F \left(\mathcal{U}_{K_a}^{\otimes T} \right) \right\|_{\text{tr}} \right] \leq \delta + 2T/(d+1) + \varepsilon .$$

The proof is identical to that of Lemma 6.5, but again replacing statistical distance with trace distance. Then, by a proof analogous to that of Lemma 6.6, we obtain:

Lemma 8.13 (Quantum Disguising-Distribution Lemma). *Let $R(x^1, \dots, x^t) : \{0, 1\}^{t \times n} \rightarrow \text{MS}_{t'}$ be any possibly-randomized mapping, where $n, t, t' \in \mathbb{N}^+$. Let $S \subseteq \{0, 1\}^n$, and fix $d > 0$. Given any $\varepsilon > 0$, let $s := \lceil (.5 \ln 2)n/\varepsilon^2 \rceil$. Let*

$$\widehat{\delta} := \min \left\{ \sqrt{\frac{\ln 2}{2} \cdot \frac{t'}{t}}, \quad 1 - 2^{-\frac{t'}{t}-2} \right\} .$$

Then there exists a collection K_1, \dots, K_s of size- d multisets contained in S , such that for every $y \in S$, we have

$$\mathbb{E}_{a \sim \mathcal{U}_{[s]}, j^* \sim \mathcal{U}_{[t]}} \left[\left\| R \left(\mathcal{U}_{K_a}^{\otimes(j^*-1)}, y, \mathcal{U}_{K_a}^{\otimes(t-j^*)} \right) - R \left(\mathcal{U}_{K_a}^{\otimes t} \right) \right\|_{\text{tr}} \right] \leq \widehat{\delta} + 2t/(d+1) + \varepsilon .$$

8.6 Complexity upper bounds from quantum compression schemes

Now we are ready to prove a quantum analogue of Theorem 7.1.

Theorem 8.14. *Let L be any language. Suppose there is a QPT-OR-compression reduction $R(x^1, \dots, x^t) : \{0, 1\}^{t_1(n) \times n} \rightarrow \text{MS}_{t_2(n)}$ for L with (not necessarily computable) parameters $t_1(n), t_2(n) : \mathbb{N}^+ \rightarrow \mathbb{N}^+$, and with error bound $\xi(n) < .5$. Let*

$$\widehat{\delta} := \min \left\{ \sqrt{\frac{\ln 2}{2} \cdot \frac{t'}{t}}, \quad 1 - 2^{-\frac{t'}{t}-2} \right\} .$$

1. *If for some $c > 0$ we have*

$$(1 - 2\xi(n)) - \widehat{\delta} \geq \frac{1}{n^c} , \tag{18}$$

then there is a non-uniform (classical, deterministic) polynomial-time many-to-one reduction from L to a problem in pr-QIP[2].

2. *If we have the stronger bound*

$$(1 - 2\xi(n))^2 - \widehat{\delta} \geq \frac{1}{n^c} , \tag{19}$$

then L has a non-uniform (classical, deterministic) polynomial-time many-to-one reduction to a problem in pr-QSZK.

Proof of Theorem 8.14. The proof is closely analogous to that of Theorem 7.1, except that our non-uniform reduction, on input y , outputs a description $\langle C, C' \rangle$ of a pair of *quantum* circuits. If $y \in L$, then $\|\rho_C - \rho_{C'}\|_{\text{tr}} \geq D(n) := 1 - 2\xi(n)$; while if $y \notin L$, we have $\|\rho_C - \rho_{C'}\|_{\text{tr}} \leq d(n) := 2\hbar - 1 + \frac{1}{n^\epsilon}$. Applying Theorems 8.7 and 8.8 gives us the complexity upper bounds in items 1 and 2. \square

Using Theorem 8.14, we can prove quantum versions of Theorems 7.4 and 7.5, giving evidence against efficient quantum OR- and AND-compression for NP-complete languages, under the assumption that such languages are not non-uniformly reducible to problems in pr-QIP[2], or alternatively, in pr-QSZK. A quantum analogue of Theorem 7.3, item 1 can be proved. We can also give an analogue of Theorem 7.7 regarding quantum compression for “expressive” parametrized problems. All of these quantum results treat compression reductions where the output state is of size determined by the various input parameters.

9 Questions for further study

1. Can we extend the limitations we show on efficient compression for AND(SAT) and OR(SAT), to give corresponding lower bounds on the cost of solving these problems in the *oracle communication model* studied by Dell and Van Melkebeek [DvM10]? These authors were able to extend the lower bounds of [FS11] for OR(SAT) to this more general setting. Proceeding by straightforward analogy in our case seems to fail, however.
2. Using our results on the infeasibility of compression for AND(SAT), can we extend the work of [DvM10] to prove new kernel-size lower bounds for interesting problems *with* polynomial kernels, under the assumption $\text{NP} \not\subseteq \text{coNP/poly}$?
3. Can we obtain a better quantitative understanding of the limits to efficient f -compression of NP-complete languages, where f is a combining function other than OR or AND? The case $f = \bigvee_{i=1}^m \left(\bigwedge_{j=1}^m x^{i,j} \right)$ is an interesting candidate for study.

As noted in Section 1.3.1, our methods imply infeasibility of strong f -compression to any target language, for any *monotone* combining f depending on all variables. We can also handle non-monotone f , provided the *sensitivity* $s(f)$ satisfies $s(f) = n^{\Omega(1)}$. (See [BdW02] for background on sensitivity.)

If f is non-monotone and depends on all variables, but has low sensitivity, we seem to need the additional requirement that the target language L' be in NP in order to prove compression lower bounds (under the assumption $\text{NP} \not\subseteq \text{coNP/poly}$). As a concrete example of a question this leaves open: is there a strong f -compression reduction from SAT to *any* target language L' (beyond NP), when f is the “index function”

$$f(x_1, \dots, x_{\log n}, y_1, \dots, y_n) := y_x ?$$

4. Can we find other applications for the Disguising-Distribution Lemma?

Acknowledgements

I thank Hans Bodlaender, Holger Dell, Lance Fortnow, Russell Impagliazzo, James Lee, Dieter van Melkebeek, Ashwin Nayak, Karolina Soltys, Salil Vadhan, Avi Wigderson, Ryan Williams, and

several anonymous reviewers for helpful comments. Thanks especially to Russell, Ashwin, and Salil for allowing me to include their alternative proof suggestions.

References

- [AB09] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.
- [Adl78] Leonard M. Adleman. Two theorems on random polynomial time. In *19th IEEE FOCS*, pages 75–83, 1978.
- [AH91] William Aiello and Johan Håstad. Statistical zero-knowledge languages can be recognized in two rounds. *J. Comput. Syst. Sci.*, 42(3):327–345, 1991.
- [Alt94] Ingo Althöfer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 199, Supplement 1(0):339 – 355, 1994.
- [BDFH09] Hans L. Bodlaender, Rodney G. Downey, Michael R. Fellows, and Danny Hermelin. On problems without polynomial kernels. *J. Comput. Syst. Sci.*, 75(8):423–434, 2009. Earlier version in ICALP '08.
- [BdW02] Harry Buhrman and Ronald de Wolf. Complexity measures and decision tree complexity: a survey. *Theor. Comput. Sci.*, 288(1):21–43, 2002.
- [BG81] Charles H. Bennett and John Gill. Relative to a random oracle A , $P^A \neq NP^A \neq \text{co-NP}^A$ with probability 1. *SIAM J. Comput.*, 10(1):96–113, 1981.
- [BH08] Harry Buhrman and John M. Hitchcock. NP-hard sets are exponentially dense unless $\text{coNP} \subseteq \text{NP/poly}$. In *23rd IEEE Conference on Computational Complexity*, pages 1–7, 2008.
- [BJK11a] Hans L. Bodlaender, Bart M. P. Jansen, and Stefan Kratsch. Cross-composition: A new technique for kernelization lower bounds. In *STACS*, pages 165–176, 2011.
- [BJK11b] Hans L. Bodlaender, Bart M. P. Jansen, and Stefan Kratsch. Kernel bounds for path and cycle problems. In *IPEC*, pages 145–158, 2011.
- [BJK11c] Hans L. Bodlaender, Bart M. P. Jansen, and Stefan Kratsch. Preprocessing for treewidth: A combinatorial analysis through kernelization. In *38th ICALP*, pages 437–448, 2011.
- [BTY11] Hans L. Bodlaender, Stéphan Thomassé, and Anders Yeo. Kernel bounds for disjoint cycles and disjoint paths. *Theor. Comput. Sci.*, 412(35):4570–4578, 2011. Earlier version in ESA '09.
- [Buh] Harry Buhrman. Personal communication.
- [CCDF97] Liming Cai, Jianer Chen, Rodney G. Downey, and Michael R. Fellows. Advice classes of parameterized tractability. *Annals of Pure and Applied Logic*, 84(1):119 – 138, 1997.

- [CFM11] Yijia Chen, Jörg Flum, and Moritz Müller. Lower bounds for kernelizations and other preprocessing procedures. *Theory Comput. Syst.*, 48(4):803–839, 2011.
- [CT06] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2nd edition, 2006.
- [DF99] R. G. Downey and M.R. Fellows. *Parametrized Complexity*. Springer (Monographs in Computer Science), 1st edition, 1999.
- [DLS09] Michael Dom, Daniel Lokshtanov, and Saket Saurabh. Incompressibility through colors and IDs. In *36th ICALP*, pages 378–389, 2009.
- [DM12] Holger Dell and Dániel Marx. Kernelization of packing problems. In *23rd ACM-SIAM SODA*, pages 68–81, 2012.
- [DvM10] Holger Dell and Dieter van Melkebeek. Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses. In *42nd ACM STOC*, pages 251–260, 2010.
- [FGM⁺89] Martin Fürer, Oded Goldreich, Yishay Mansour, Michael Sipser, and Stathis Zachos. On completeness and soundness in interactive proof systems. *Advances in Computing Research*, 5:429–442, 1989.
- [FHT03] Alexei A. Fedotov, Peter Harremoës, and Flemming Topsøe. Refinements of pinsker’s inequality. *IEEE Transactions on Information Theory*, 49(6):1491–1498, 2003.
- [For87] Lance Fortnow. The complexity of perfect zero-knowledge (extended abstract). In Alfred V. Aho, editor, *19th ACM STOC*, pages 204–209, 1987.
- [FS11] Lance Fortnow and Rahul Santhanam. Infeasibility of instance compression and succinct PCPs for NP. *J. Comput. Syst. Sci.*, 77(1):91–106, 2011.
- [GN07] Jiong Guo and Rolf Niedermeier. Invitation to data reduction and problem kernelization. *SIGACT News*, 38(1):31–45, 2007.
- [GS86] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In *18th ACM STOC*, pages 59–68, 1986.
- [GSV98] Oded Goldreich, Amit Sahai, and Salil P. Vadhan. Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In *30th ACM STOC*, pages 399–408, 1998.
- [GV11] Oded Goldreich and Salil P. Vadhan. On the complexity of computational problems regarding distributions (a survey). *Electronic Colloquium on Computational Complexity (ECCC)*, TR11-004:4, 2011.
- [HN10] Danny Harnik and Moni Naor. On the compressibility of NP instances and cryptographic applications. *SIAM J. Comput.*, 39(5):1667–1713, 2010.
- [HW12] Danny Hermelin and Xi Wu. Weak compositions and their applications to polynomial lower bounds for kernelization. In *23rd ACM-SIAM SODA*, pages 104–113, 2012.

- [JJUW11] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP = PSPACE. *J. ACM*, 58(6):30, 2011. Earlier version in STOC '10.
- [KdW04] Iordanis Kerenidis and Ronald de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. *J. Comput. Syst. Sci.*, 69(3):395–420, 2004. Earlier version in STOC '03.
- [KNTSZ07] Hartmut Klauck, Ashwin Nayak, Amnon Ta-Shma, and David Zuckerman. Interaction in quantum communication. *IEEE Transactions on Information Theory*, 53(6):1970–1982, 2007.
- [Kra12] Stefan Kratsch. Co-nondeterminism in compositions: a kernelization lower bound for a Ramsey-type problem. In *23rd ACM-SIAM SODA*, pages 114–122, 2012.
- [KW00] Alexei Kitaev and John Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. In *32nd ACM STOC*, pages 608–617, 2000.
- [KW12] Stefan Kratsch and Magnus Wahlström. Compression via matroids: a randomized polynomial kernel for odd cycle transversal. In *23rd ACM-SIAM SODA*, pages 94–103, 2012.
- [LMM03] Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *4th ACM Conference on Electronic Commerce*, pages 36–41, 2003.
- [LY94] Richard J. Lipton and Neal E. Young. Simple strategies for large zero-sum games with applications to complexity theory. In *26th ACM STOC*, pages 734–740, 1994.
- [Nay99a] Ashwin Nayak. *Lower bounds for Quantum Computation and Communication*. PhD thesis, University of California, Berkeley, 1999.
- [Nay99b] Ashwin Nayak. Optimal lower bounds for quantum automata and random access codes. In *40th IEEE FOCS*, pages 369–377, 1999.
- [NC00] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [Nis91] Noam Nisan. Crew prams and decision trees. *SIAM J. Comput.*, 20(6):999–1007, 1991. Earlier version in STOC '89.
- [Oka00] Tatsuaki Okamoto. On relationships between statistical zero-knowledge proofs. *J. Comput. Syst. Sci.*, 60(1):47–108, 2000. Earlier version in STOC '96.
- [OP04] Masanori Ohya and Denes Petz. *Quantum Entropy and its Use*. Texts and Monographs in Physics. Springer-Verlag, Heidelberg, 2nd edition, 2004.
- [Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998. Earlier version in STOC '95.
- [Reg11] Oded Regev. Entropy-based bounds on dimension reduction in L_1 . arXiv:1108.1283v3, 2011.

- [RW09] Mark D. Reid and Robert C. Williamson. Generalised pinsker inequalities. In *COLT*, 2009.
- [Sha10] Ronen Shaltiel. Derandomized parallel repetition theorems for free games. In *IEEE Conference on Computational Complexity*, pages 28–37, 2010.
- [SV03] Amit Sahai and Salil P. Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 50(2):196–249, 2003.
- [SV08] Pranab Sen and Srinivasan Venkatesh. Lower bounds for predecessor searching in the cell probe model. *J. Comput. Syst. Sci.*, 74(3):364–385, 2008. Earlier version in CCC '03.
- [vL99] J. H. van Lint. *Introduction to coding theory*, volume 86 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 3rd edition, 1999.
- [Wat02] John Watrous. Limits on the power of quantum statistical zero-knowledge. In *43rd IEEE FOCS*, pages 459–468, 2002.
- [Wat03] John Watrous. PSPACE has constant-round quantum interactive proof systems. *Theor. Comput. Sci.*, 292(3):575–588, 2003. Earlier version in FOCS '99.
- [Wat09] John Watrous. Zero-knowledge against quantum attacks. *SIAM J. Comput.*, 39(1):25–58, 2009. Earlier version in STOC '06.
- [Yap83] Chee-Keng Yap. Some consequences of non-uniform conditions on uniform classes. *Theor. Comput. Sci.*, 26:287–300, 1983.

A Alternative proofs of distributional stability

A.1 A proof based on Raz’s lemma

R. Impagliazzo and S. Vadhan noted a similarity between distributional stability lemmas and a probabilistic lemma implicit in work of Raz [Raz98]. Vadhan pointed us to the following convenient form, given by Shaltiel in [Sha10, Lemma 3.1]:

Lemma A.1. *There is a $c > 0$ for which the following holds. Let X^1, \dots, X^t be i.i.d. random variables, and T an event with $\Pr[T] \geq 2^{-\Delta}$. Then for the conditioned variables $X_{[T]}^j$ we have*

$$\mathbb{E}_{j \sim \mathcal{U}_{[t]}}[\|X_{[T]}^j - X^j\|_{\text{stat}}] \leq \varepsilon := \sqrt{\frac{c\Delta}{t}}.$$

With this lemma we can derive a distributional stability result as follows. Suppose $R : S^t \rightarrow \{0, 1\}^{\leq t'}$ is given, and consider independent inputs $X^j \sim \mathcal{D}$ to R . We let \mathbf{R} denote the random output value. For any output z of R , Lemma A.1 above implies that

$$\mathbb{E}_{i \sim \mathcal{U}_{[n]}}[\|X_{[\mathbf{R}=z]}^j - X^j\|_{\text{stat}}] \leq \sqrt{c \log_2(1/\Pr[\mathbf{R} = z])/t}.$$

Taking expectations over $z \sim \mathbf{R}$ and using Jensen's inequality,

$$\begin{aligned} \mathbb{E}_{z \sim \mathbf{R}, i \sim \mathcal{U}_{[n]}} [\|X_{[\mathbf{R}=z]}^j - X^j\|_{\text{stat}}] &\leq \mathbb{E}_z \left[\sqrt{c \log_2(1/\Pr[\mathbf{R}=z])}/t \right] \\ &\leq \sqrt{\mathbb{E}_z [c \log_2(1/\Pr[\mathbf{R}=z])/t]} \\ &= \sqrt{c \cdot H_{rv}(t')/t} \\ &\text{(by the definition of Shannon entropy)} \\ &< \sqrt{c(t'+1)/t} . \end{aligned}$$

Let $\mathbf{R}' = R(Y^1, \dots, Y^t)$ denote a sample of \mathbf{R} based on inputs $Y^1, \dots, Y^t \sim \mathcal{D}^{\otimes t}$ that are independent of X^1, \dots, X^t . Now the crucial observation is that, for each $j \in [t]$, we have the chain of equalities

$$\mathbb{E}_z [\|X_{[\mathbf{R}=z]}^j - X^j\|_{\text{stat}}] = \| (X^j, \mathbf{R}) - (X^j, \mathbf{R}') \|_{\text{stat}} = \mathbb{E}_{x^j \sim \mathcal{D}} [\| \mathbf{R}_{[X^j=x^j]} - \mathbf{R} \|_{\text{stat}}] .$$

Each equality follows from the ‘‘distinguishability interpretation’’ of statistical distance. Combining we get

$$\mathbb{E}_{x^j \sim \mathcal{D}} [\| \mathbf{R}_{[X^j=x^j]} - \mathbf{R} \|_{\text{stat}}] < \sqrt{c(t'+1)/t} ,$$

which is comparable to what we got from the previous approach (up to constant factors; here we have assumed i.i.d. variables X^j , but this is not essential for this approach).

A.2 A proof based on the Average Encoding Theorem

The Average Encoding Theorem of [KNTSZ07] is a tool in quantum information theory, that has the following classical analogue. (See [SV08, Fact 5], where a purely classical proof is given. We restate the result slightly, converting from ℓ^1 distance to statistical distance.)

Theorem A.2. *Let X, M be random variables. Let Π be the distribution governing M . Then for the conditioned distributions $\Pi_{[X=x]}$ we have*

$$\sum_x \Pr[X=x] \cdot \| \Pi_{[X=x]} - \Pi \|_{\text{stat}} \leq \sqrt{\frac{\ln 2 \cdot I(X; M)}{2}} .$$

Using Theorem A.2 along with techniques suggested by Nayak and similar to those in the proof of [KNTSZ07, Theorem 5.4],⁴⁰ we can derive a distributional-stability result as follows. Again say we are given $R : S^t \rightarrow \{0, 1\}^{t'}$, and consider independent inputs $X^j \sim \mathcal{D}$ to R , giving an output distribution denoted \mathbf{R} . Applying Theorem A.2 to $X := X^j, M := \mathbf{R}$, we have

$$\mathbb{E}_{x^j \sim X^j} [\| \mathbf{R}_{[X^j=x]} - \mathbf{R} \|_{\text{stat}}] \leq \sqrt{\frac{\ln 2 \cdot I(X^j; \mathbf{R})}{2}} .$$

⁴⁰(That proof uses a version of the Average Encoding Theorem that treats Hellinger distance rather than statistical distance.)

Averaging over $j \in [t]$ and applying Jensen’s inequality and Lemma 4.4, we obtain

$$\begin{aligned} \mathbb{E}_{j \sim \mathcal{U}_{[t]}, x^j \sim X^j} [|\mathbf{R}_{[X^j=x]} - \mathbf{R}|_{\text{stat}}] &\leq \mathbb{E}_{j \sim \mathcal{U}_{[t]}} \left[\sqrt{\frac{\ln 2 \cdot I(X^j; \mathbf{R})}{2}} \right] \\ &\leq \sqrt{\frac{\ln 2 \cdot \mathbb{E}_{j \sim \mathcal{U}_{[t]}} [I(X^j; \mathbf{R})]}{2}} \\ &\leq \sqrt{\frac{\ln 2 \cdot (t' + 1)}{2t}} . \end{aligned}$$

B Our original distributional stability lemma

In this section we include our original proof of a distributional stability lemma, based on coding-theoretic ideas. This lemma proves a distributional stability result for mappings $R : \{0, 1\}^{t \times n} \rightarrow \{0, 1\}^{\leq t'}$, where $t' + 2 \leq t$. In an earlier draft, we used more complicated ideas to prove complexity upper bounds from AND-compression reductions where $t' = O(t \log t)$. This latter case can now be handled in the same way as the case $t' \ll t$, using the alternative bound on distributional stability provided by Lemma 6.2, item 2.

First, we provide some further needed background.

B.1 Entropy and the unreliability of compressive encodings

It is a basic principle of information theory that one cannot reliably encode a uniformly-generated t -bit message by an encoding of length $t - 2$ or less. (We can save essentially one bit by using a variable-length output.) Below we state and prove a standard claim that generalizes this fact, giving quantitative bounds on the reliability of compressive encoding methods.

Lemma B.1. *Let $t \in \mathbb{N}^+$, and let U be some finite universe. Say we are given a possibly-randomized “encoding” function*

$$\text{Enc}(x, y) : \{0, 1\}^t \times \{0, 1\}^N \rightarrow U ,$$

depending on a “message” input $x \in \{0, 1\}^t$ along with a “public randomness” input $y \in \{0, 1\}^N$. (Enc may also have additional internal randomness.) Say we are also given a (possibly-randomized, possibly-unreliable) “decoding” function

$$\text{Dec}(x, y) : U \times \{0, 1\}^N \rightarrow \{0, 1\}^t ,$$

that also has access to the public randomness y .

Suppose X, Y are two independent random variables over $\{0, 1\}^t, \{0, 1\}^N$ respectively. For $j \in [t]$, let

$$p_j := \Pr_{X, Y} [\text{Dec}_j(\text{Enc}(X, Y), Y) = X_j] ,$$

where $X = (X_1, \dots, X_t)$, and where Dec_j is the j^{th} output bit of Dec .

Let $p_{\text{avg}} := \frac{1}{t} \sum_{j=1}^t p_j$. Then, we must have

$$H(p_{\text{avg}}) \geq \frac{1}{t} (H_{\text{rv}}(X) - \log_2(|U|)) .$$

Our proof is closely modeled on the proof of a corresponding, but deeper, quantum result [KdW04, Appendix B].⁴¹) To prove Lemma B.1, we will use another basic information-theoretic fact, Fano’s inequality:

Lemma B.2 (Fano). [CT06, Chapter 2] Suppose Z_{in}, Z_{out} are two random variables: Z_{in} an “input message” over some alphabet Σ , and Z_{out} an “output message” over any domain. Let \tilde{Z}_{in} be a (possibly-randomized) function of Z_{out} , that attempts to recover the value Z_{in} . Let

$$p_{err} := \Pr[\tilde{Z}_{in} \neq Z_{in}] .$$

Here the randomness is over the entire experiment. Then, we have

$$H(p_{err}) + p_{err} \cdot \log_2(|\Sigma| - 1) \geq H_{rv}(Z_{in}|Z_{out}) .$$

We only use the case $|\Sigma| = 2$, so the second term on the left-hand side vanishes.

Proof of Lemma B.1. First we ask whether $p_{avg} \geq .5$. If not, we simply negate all of the decoding functions Dec_j , giving the modified average success probability $p'_{avg} = 1 - p_{avg}$, for which $H(p'_{avg}) = H(p_{avg})$. Next, note that the success probabilities p_j are taken over the randomness both in X and in Y (as well as in Enc, Dec). As Y is independent of X , we may non-uniformly fix some setting to Y that maximizes the sum of the conditional success probabilities. Then the re-modified average success probability satisfies $p''_{avg} \geq p'_{avg} \geq .5$. As $H(\cdot)$ is decreasing on $[.5, 1]$, any lower bound proved for $H(p''_{avg})$ will also lower-bound the $H(p_{avg})$ for the original encoding scheme with public randomness. Thus in the remainder of the proof, we assume $p_{avg} \geq .5$ and that the scheme uses no public randomness: our encoding Enc applies to X alone, and our decoding functions apply to the message $Enc(X)$ alone.

The chain rule for conditional entropy and the subadditivity of entropy imply that

$$H_{rv}(X|Enc(X)) = \sum_{j=1}^t H_{rv}(X_j|X_1, \dots, X_{j-1}, Enc(X)) \leq \sum_{j=1}^t H_{rv}(X_j|Enc(X)) . \quad (20)$$

Next, we apply Fano’s inequality, with $Z_{in} := X_j$, and $Z_{out} := Enc(X)$. Thus in this analysis we simply view $(X_{j'})_{j' \neq j}$ as additional sources of randomness in the encoding process. We let $\tilde{Z}_{in} := Dec_j(Enc(X))$. X_j is binary— $\Sigma = \{0, 1\}$ —so Fano’s inequality gives

$$H(p_j) = H(1 - p_j) \geq H_{rv}(X_j|Enc(X)) .$$

Summing over j and using Eq. (20),

$$\begin{aligned} \sum_{j=1}^t H(p_j) &\geq H_{rv}(X|Enc(X)) \\ &\geq H_{rv}(X) - H_{rv}(Enc(X)) , \end{aligned} \quad (21)$$

⁴¹(or, Appendix A in the arxiv version. This part of [KdW04] is itself a rederivation of a result from [Nay99b]; a similar result and proof appears in Nayak’s thesis [Nay99a, Theorem 3.2.8].

again using subadditivity. Now, $Enc(X)$ is a message over U , so $H_{rv}(Enc(X)) \leq \log_2(|U|)$. Also, the function H is concave on $[0, 1]$. Applying these observations to Eq. (21), and using Jensen's inequality, we have

$$H(p_{avg}) \geq \frac{1}{t} \sum_{j=1}^t H(p_j) \geq \frac{1}{t} (H_{rv}(X) - \log_2(|U|)) .$$

□

B.2 Bounds on the inverse entropy function

For $\alpha \in [0, 1]$, we will denote by $H_+^{-1}(\alpha)$ the unique H -preimage of α in the range $[\cdot 5, 1]$. Similarly, let $H_-^{-1}(\alpha)$ denote the unique H -preimage of α in the range $[0, \cdot 5]$. The following bounds on the inverse entropy function are useful in understanding the bounds provided by our original distributional stability lemma (Lemma B.4, to be presented shortly). These bounds on $H_+^{-1}(\cdot)$ are meant to be simple and illustrative, and are not quite best-possible.

Lemma B.3. *We have the following facts:*

1. If $m > 0$ is sufficiently large, then $H_+^{-1}(1/m) < 1 - 1/(4m \log_2 m)$.
2. $H_+^{-1}(1 - \delta) \leq \cdot 5 + \frac{\sqrt{\ln 2}}{2} \sqrt{\delta} + O(\delta^{3/2})$.

Proof. (1.) First consider any value $p \in (0, 1/2)$. We can upper-bound $H(p)$ in the following way:

$$\begin{aligned} H(p) &= p \log_2(1/p) + \underbrace{(1-p)}_{\leq 1} \log_2 \underbrace{(1/(1-p))}_{\leq 1+2p} \\ &< p \log_2(1/p) + \underbrace{2p}_{\text{(using } \log(1+c) < c \text{ for } c > 0)} \\ &< 3p \log_2(1/p) . \end{aligned}$$

From this, one can easily verify that for $m \geq 10^6$, we have

$$H\left(\frac{1}{4m \log_2 m}\right) < \frac{1}{m} .$$

Thus for such m ,

$$H_-^{-1}\left(\frac{1}{m}\right) > \frac{1}{4m \log_2 m} ,$$

so that

$$H_+^{-1}\left(\frac{1}{m}\right) = 1 - H_-^{-1}\left(\frac{1}{m}\right) < 1 - \frac{1}{4m \log_2 m} ,$$

giving item 1.

(2.) The binary entropy function H is infinitely differentiable on $(0, 1)$, with

$$H(\cdot 5) = 1, \quad H'(\cdot 5) = 0, \quad H''(\cdot 5) = -4(\ln 2)^{-1} .$$

Thus for $\beta \in [0, .25)$ we have

$$H(.5 + \beta) \leq 1 - 4(\ln 2)^{-1}\beta^2 + O(\beta^3) .$$

By considering settings $\beta := \frac{\sqrt{\ln 2}}{2}\sqrt{\delta} \pm O(\delta^{3/2})$ and using that $H(\cdot)$ is decreasing on $(.5, 1)$, we verify item 2. \square

B.3 The lemma

Lemma B.4. *Let $R(x^1, \dots, x^t) : \{0, 1\}^{t \times n} \rightarrow \{0, 1\}^{\leq t'}$ be any possibly-randomized mapping, where $t, t' \in \mathbb{N}^+$ satisfy $t' + 2 \leq t$.*

Then, R is δ -distributionally stable with respect to any input distributions $\mathcal{D}_1, \dots, \mathcal{D}_t$, where

$$\delta := 2H_+^{-1} \left(1 - \frac{t' + 1}{t} \right) - 1 .$$

We will prove Lemma B.4 by a reduction to an encoding/decoding task that allows us to apply Lemma B.1.

Proof of Lemma B.4. Let $\mathcal{D}_1, \dots, \mathcal{D}_t$ be independent distributions over $\{0, 1\}^n$, and for $j \in [t]$, let γ_j be as in Definition 6.1 for $F := R$.

Consider the following encoding/decoding experiment involving t ‘‘Receivers.’’ In the experiment, we will use R as a communication channel to attempt to transmit t bits b_1, \dots, b_t . Receiver $j \in [t]$ will be responsible for attempting to recover the value b_j . Formally:

1. For $j \in [t]$, let $y^j, w^j \sim \mathcal{D}_j$ (here y^j, w^j are independent of each other and of all other $y^{j'}, w^{j'}$);
2. Also, and independently, for $j \in [t]$ let $b_j \sim \mathcal{U}_{\{0,1\}}$;
3. If $b_j = 0$, let $x^j := y^j$. Otherwise, let $x^j := w^j$;
4. Let $z := R(x^1, \dots, x^t)$ (a possibly-randomized value), and let z be sent to the Receivers. Let $\{y^j\}_{j \in [t]}$ be visible to the receivers as public randomness;
5. Each Receiver j outputs a guess \tilde{b}_j for b_j , based on the values of the two random variables y^j and z . Specifically, Receiver j uses the maximum-likelihood rule $\tilde{b}_j := ML(b|y^j, z)$, described in Section 2.1.

Note that in making the guess \tilde{b}_j , Receiver j does *not* inspect the values $y^{j'}, j' \neq j$.

We analyze this experiment. First observe that, *conditioned* on a value y^j seen by Receiver j and on the value $b_j \in \{0, 1\}$ (which Receiver j does not see in the actual experiment), but leaving the other values $\{y^{j'}\}_{j' \neq j}$ unconditioned, the conditional distribution on z is that $z \sim R(\mathcal{D}_1, \dots, \mathcal{D}_{j-1}, y^j, \mathcal{D}_{j+1}, \dots, \mathcal{D}_t)$ if $b_j = 0$, and $z \sim R(\mathcal{D}_1, \dots, \mathcal{D}_j, \dots, \mathcal{D}_t)$ if $b_j = 1$.

Also, b_j is unbiased and independent of y^j . Thus, by the distinguishability interpretation of statistical distance (see Section 2.1), Receiver j ’s success probability in guessing b_j , conditioned exclusively on an observed value y^j , equals

$$\frac{1}{2} \left(1 + \left\| R(\mathcal{D}_1, \dots, \mathcal{D}_{j-1}, y^j, \mathcal{D}_{j+1}, \dots, \mathcal{D}_t) - R(\mathcal{D}_1, \dots, \mathcal{D}_j, \dots, \mathcal{D}_t) \right\|_{\text{stat}} \right) .$$

Thus Receiver j 's *overall* success probability in the experiment is precisely $\frac{1}{2}(1 + \gamma_j)$, where γ_j is as in the definition of distributional stability for R with respect to $\mathcal{D}_1, \dots, \mathcal{D}_t$.

In our present setup, we can regard

$$z = R(x^1, \dots, x^t) =: \text{Enc}(b_1, \dots, b_t, y^1, \dots, y^t)$$

as a randomized encoding function of b_1, \dots, b_t , with public randomness y^1, \dots, y^t and additional private randomness w^1, \dots, w^t . Similarly, we can view

$$(\tilde{b}_1, \dots, \tilde{b}_t) =: \text{Dec}(z, y^1, \dots, y^t)$$

as a (deterministic) decoding function. The success probability of our encoding/decoding experiment in successfully decoding b_j is $\frac{1}{2}(1 + \gamma_j)$.

Now $H_{rv}(b_1, \dots, b_t) = t$, and $H_{rv}(\text{Enc}(b_1, \dots, b_t, y^1, \dots, y^t)) \leq \log_2 \left(\left| \{0, 1\}^{\leq t'} \right| \right) < t' + 1$. Applying Lemma B.1, we find that

$$H \left(\frac{1}{2} \left(1 + \frac{1}{t} \sum_{j=1}^t \gamma_j \right) \right) \geq \frac{1}{t} (t - (t' + 1)) = 1 - \frac{t' + 1}{t},$$

which implies that

$$\frac{1}{t} \sum_{j=1}^t \gamma_j \leq 2H_+^{-1} \left(1 - \frac{t' + 1}{t} \right) - 1 = \delta.$$

Thus, R is δ -DS with respect to $\mathcal{D}_1, \dots, \mathcal{D}_t$. As these distributions were arbitrary, this proves Lemma B.4. \square

It is also possible to give a slightly shorter and more direct proof of Lemma B.4, using the same basic information-theoretic steps from the proof of Lemma B.1; one can apply these to steps 1-4 of the experiment described above, without defining the decoding process.⁴² However, we feel that the reduction to encoding/decoding is more intuitive, and more clearly illustrates a general idea that has been useful in various settings, e.g., in [Reg11].

C Proof of quantum distributional stability

Here we describe how Lemma 8.10 is proved. We use various concepts and results of quantum information theory. In particular, we assume familiarity with the notion of bipartite and reduced states. We let $S(\rho) := -\text{Tr}(\rho \log_2 \rho)$ denote the *Von Neumann entropy* of a quantum state (here, identifying ρ with its density matrix). For a bipartite state $\rho_{A,B}$ on subsystems A, B , we let ρ_A (resp. ρ_B) denote the reduced state over A (resp. B)

We define the *quantum mutual information* between subsystems A, B of a bipartite state ρ_{AB} as

$$I(A; B) := S(\rho_A) + S(\rho_B) - S(\rho_{AB}).$$

See [NC00], and Chapter 11 in particular, for more background in quantum information.

⁴²I thank James Lee and Avi Wigderson for suggesting to find a more direct information-theoretic proof.

We will sometimes speak of quantum systems containing a subsystem that is a classical random variable X . By this we mean a state of form

$$\sum_{x \in \text{supp}(X)} \Pr[X = x] \cdot |x\rangle\langle x| \otimes \rho_x ,$$

for some collection of quantum states $\{\rho_x\}$ on a fixed number of qubits.

Not all properties of classical entropy and mutual information are inherited by their quantum counterparts. However, we have [Nay99a, p. 33 and Appendix A]:

Fact C.1. *Quantum mutual information obeys the following properties, for all X, Y, Z :*

1. $I(X; Y) = I(Y; X)$;
2. $I(X; (Y, Z)) = I(X; Y) + I((X, Y); Z) - I(Y; Z)$;
3. (*Strong subadditivity*) $I(X; (Y, Z)) \geq I(X; Y)$;
4. $I(X; Z) = 0$ if the subsystems X, Z are independent classical random variables.

Item 3 is a nontrivial fact in the quantum setting, with multiple equivalent formulations; see [NC00, Chapter 11].

With these facts in hand, the proof of Lemma C.2 below exactly follows that of Lemma 4.4.

Lemma C.2. *If X^1, \dots, X^t are independent classical random variables and Y a quantum subsystem, then*

$$I(Y; (X^1, \dots, X^t)) \geq \sum_{j \in [t]} I(Y; X^j) .$$

Next, we need quantum analogues of Pinsker's and Vajda's inequalities. For mixed states ρ, σ over the same number of qubits, define the *relative entropy* (a quantum analogue of Kullback-Leibler divergence) as

$$S(\rho||\sigma) := \text{Tr}(\rho \log_2(\sigma)) - S(\rho) .$$

We also have the following analogue of Fact 4.6 [KNTSZ07, p. 10]:

Fact C.3. $I(A; B) = S(\rho_{AB}||\rho_A \otimes \rho_B)$.

A quantum Pinsker inequality was explicitly proved in [KNTSZ07, Theorem III.1].⁴³ However, that proof actually demonstrates a more general principle:

Theorem C.4 ([KNTSZ07]). *Suppose that for some $\alpha, \beta \geq 0$, the (classical) statistical distance and Kullback-Leibler divergence obey the relationship*

$$\|X - Y\|_{\text{stat}} \geq \alpha \implies D_{\text{KL}}(X||Y) \geq \beta .$$

for every pair of classical distributions X, Y .

Then, for any pair ρ, σ of quantum states,

$$\|\rho - \sigma\|_{\text{tr}} \geq \alpha \implies S(\rho||\sigma) \geq \beta .$$

⁴³An earlier version appears in [OP04].

Combining this principle with the classical Pinsker and Vajda inequalities, we obtain:

Corollary C.5 (Quantum Pinsker inequality). *For any states ρ, σ ,*

$$S(\rho||\sigma) \geq \frac{2}{\ln 2} \cdot \|\rho - \sigma\|_{\text{tr}}^2$$

Corollary C.6 (Quantum Vajda inequality). *For any states ρ, σ ,*

$$S(\rho||\sigma) \geq \frac{1}{\ln 2} \left(\ln \left(\frac{1}{1 - \|\rho - \sigma\|_{\text{tr}}} \right) - 1 \right).$$

With these tools, the proof of Lemma 8.10 is essentially identical to that of Lemma 6.2. We omit the details.