

Statistical Difference Beyond the Polarizing Regime[★]

Itay Berman^{★★}, Akshay Degwekar^{★★}, Ron D. Rothblum^{★★★}, and Prashant Nalini Vasudevan[†]

Abstract. The polarization lemma for statistical distance (SD), due to Sahai and Vadhan (JACM, 2003), is an efficient transformation taking as input a pair of circuits (C_0, C_1) and an integer k and outputting a new pair of circuits (D_0, D_1) such that if $\text{SD}(C_0, C_1) \geq \alpha$ then $\text{SD}(D_0, D_1) \geq 1 - 2^{-k}$ and if $\text{SD}(C_0, C_1) \leq \beta$ then $\text{SD}(D_0, D_1) \leq 2^{-k}$. The polarization lemma is known to hold for any constant values $\beta < \alpha^2$, but extending the lemma to the regime in which $\alpha^2 \leq \beta < \alpha$ has remained elusive. The focus of this work is in studying the latter regime of parameters. Our main results are:

1. Polarization lemmas for different notions of distance, such as *Triangular Discrimination* (TD) and *Jensen-Shannon Divergence* (JS), which enable polarization for some problems where the statistical distance satisfies $\alpha^2 < \beta < \alpha$. We also derive a polarization lemma for statistical distance with any inverse-polynomially small gap between α^2 and β (rather than a constant).
2. The average-case hardness of the statistical difference problem (i.e., determining whether the statistical distance between two given circuits is at least α or at most β), for any values of $\beta < \alpha$, implies the existence of one-way functions. Such a result was previously only known for $\beta < \alpha^2$.
3. A (direct) constant-round interactive proof for estimating the statistical distance between any two distributions (up to any inverse polynomial error) given circuits that generate them. Proofs of closely related statements have appeared in the literature but we give a new proof which we find to be cleaner and more direct.

[★] This is an extended abstract. The full version of this paper, with formal statements of the theorems and proofs, may be found at: <https://eccc.weizmann.ac.il/report/2019/038/>.

^{★★} MIT. Emails: {itayberm, akshayd}@mit.edu. Research supported in part by NSF Grants CNS-1413920 and CNS-1350619, MIT-IBM Award, and by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under Vinod Vaikuntanathan's DARPA Young Faculty Award and contracts W911NF-15-C-0226 and W911NF-15-C-0236.

^{★★★} Technion. Email: rothblum@cs.technion.ac.il. This research was supported in part by the Israeli Science Foundation (Grant No. 1262/18).

[†] UC Berkeley. Email: prashvas@berkeley.edu. Research supported in part from DARPA/ARL SAFEWARE Award W911NF15C0210, AFOSR Award FA9550-15-1-0274, AFOSR YIP Award, a Hellman Award and research grants by the Okawa Foundation, Visa Inc., and Center for LongTerm Cybersecurity (CLTC, UC Berkeley).

1 Introduction

The STATISTICAL DIFFERENCE PROBLEM, introduced by Sahai and Vadhan [SV03], is a central computational (promise) problem in complexity theory and cryptography, which is also intimately related to the study of statistical zero-knowledge (SZK). The input to this problem is a pair of circuits C_0 and C_1 , specifying probability distributions (i.e., that are induced by feeding the circuits with a uniformly random string). YES instances are those in which the statistical distance¹ between the two distributions is at least $2/3$ and NO instances are those in which the distance is at most $1/3$. Input circuits that do not fall in one of these two cases are considered to be outside the promise (and so their value is left unspecified).

The choice of the constants $1/3$ and $2/3$ in the above definition is somewhat arbitrary (although not entirely arbitrary as will soon be discussed in detail). A more general family of problems can be obtained by considering a suitable parameterization. More specifically, let $0 \leq \beta < \alpha \leq 1$. The (α, β) parameterized version of the STATISTICAL DIFFERENCE PROBLEM, denoted $\text{SDP}^{\alpha, \beta}$, has as its YES inputs pairs of circuits that induce distributions that have distance at least α whereas the NO inputs correspond to circuits that induce distributions that have distance at most β .

Definition 1.1 (STATISTICAL DIFFERENCE PROBLEM). *Let $\alpha, \beta: \mathbb{N} \rightarrow [0, 1]$ with $\alpha(n) > \beta(n)$ for every n . The STATISTICAL DIFFERENCE PROBLEM with promise (α, β) , denoted $\text{SDP}^{\alpha, \beta}$, is given by the sets*

$$\begin{aligned} \text{SDP}_Y^{\alpha, \beta} &= \{(C_0, C_1) \mid \text{SD}(C_0, C_1) \geq \alpha(n)\} \text{ and} \\ \text{SDP}_N^{\alpha, \beta} &= \{(C_0, C_1) \mid \text{SD}(C_0, C_1) \leq \beta(n)\}, \end{aligned}$$

where n is the output length of the circuits C_0 and C_1 .²

(Here and below we abuse notation and use C_0 and C_1 to denote both the circuits and the respective distributions that they generate.)

The elegant *polarization lemma* of [SV03] shows how to polarize the statistical distance between two distributions. In more detail, for any constants α and β such that $\beta < \alpha^2$, the lemma gives a transformation that makes distributions that are at least α -far be extremely far and distributions that are β -close be

¹ Recall that the statistical distance between two distributions P and Q over a set $\mathcal{Y}$ is defined as $\text{SD}(P, Q) = \frac{1}{2} \sum_{y \in \mathcal{Y}} |P_y - Q_y|$, where P_y (resp., Q_y) is the probability mass that P (resp., Q) puts on $y \in \mathcal{Y}$.

² In prior works α and β were typically thought of as constants (and so their dependence on the input was not specified). In contrast, since we will want to think of them as parameters, we choose to let them depend on the output length of the circuit since this size seems most relevant to the distributions induced by the circuits. Other natural choices could have been the input length or the description size of the circuits. We remark that these different choices do not affect our results in a fundamental way.

extremely close. Beyond being of intrinsic interest, the polarization lemma is used to establish the SZK completeness of $\text{SDP}^{\alpha,\beta}$, when $\alpha^2 > \beta$, and has other important applications in cryptography such as the amplification of weak public key encryption schemes to full fledged ones [DNR04,HR05].

Sahai and Vadhan left the question of polarization for parameters α and β that do not meet the requirements of their polarization lemma as an open question. We refer to this setting of α and β as the *non-polarizing* regime. We emphasize that by *non-polarizing* we merely mean that in this regime polarization is not currently known and not that it is impossible to achieve (although some barriers are known and will be discussed further below). The focus of this work is studying the STATISTICAL DIFFERENCE PROBLEM in the non-polarizing regime.

1.1 Our Results

We proceed to describe our results.

1.1.1 Polarization and SZK Completeness for Other Notions of Distance

The statistical distance metric is one of the central information theoretic tools used in cryptography as it is very useful for capturing similarity between distributions. However, in information theory there are other central notions that measure similarity such as mutual information and KL divergence as well as others.

Loosely speaking, our first main result shows that polarization is possible even in *some* cases in which $\beta \geq \alpha^2$. However, this result actually stems from a more general study showing that polarization is possible for other notions of distance between distributions from information theory, which we find to be of independent interest.

When distributions are extremely similar or extremely dissimilar, these different notions of distance are often (but not always) closely related and hence interchangeable. This equivalence is particularly beneficial when considering applications of SZK—for some applications one distance measure may be easier to use than others. For example, showing that the average-case hardness of SZK implies one-way functions can be analyzed using statistical distance (e.g., [Vad99, Section 4.8]), but showing that every language in SZK has instance-dependent commitments is naturally analyzed using entropy (e.g., [OV08]).

However, as the gaps in the relevant distances get smaller (i.e., the distributions are only somewhat similar or dissimilar), the relation between different statistical properties becomes less clear (for example, the reduction from $\text{SDP}^{\alpha,\beta}$ to the ENTROPY DIFFERENCE PROBLEM of [GV99] only works when roughly $\alpha^2 > \beta$). This motivates studying the computational complexity of problems defined using different notions of distance in this small gap regime. Studying this question can be (and, as we shall soon see, indeed is) beneficial in two aspects. First, providing a wider bag of statistical properties related to SZK, which can

make certain applications easier to analyze. Second, the computational complexity of these distance notions might shed light on the computational complexity of problems involving existing distance notions (e.g., $\text{SDP}^{\alpha,\beta}$ when $\alpha^2 < \beta$).

We focus here on two specific distance notions—the *triangular discrimination* and the *Jensen-Shannon divergence*, defined next.

Definition 1.2 (Triangular Discrimination). *The Triangular Discrimination (a.k.a. Le Cam divergence) between two distributions P and Q is defined as*

$$\text{TD}(P, Q) = \frac{1}{2} \sum_{y \in \mathcal{Y}} \frac{(P_y - Q_y)^2}{P_y + Q_y},$$

where $\mathcal{Y}$ is the union of the supports of P and Q .

The TRIANGULAR DISCRIMINATION PROBLEM with promise (α, β) , denoted $\text{TDP}^{\alpha,\beta}$, is defined analogously to $\text{SDP}^{\alpha,\beta}$, but with respect to TD rather than SD.

The triangular discrimination is commonly used, among many other applications, in statistical learning theory for parameter estimation with quadratic loss, see [Cam86, P. 48] (in a similar manner to how statistical distance characterizes the 0-1 loss function in hypothesis testing). Jumping ahead, while the definition of triangular discrimination seems somewhat arbitrary at first glance, in Section 2 we will show that this distance notion characterizes some basic phenomena in the study of statistical zero-knowledge. Triangular discrimination has recently found usage in theoretical computer science, and even specifically in problems related to SZK. Yehudayoff [Yeh16] showed that using TD yields a tighter analysis of the pointer chasing problem in communication complexity. The work of Komargodski and Yogev [KY18] uses triangular discrimination to show that the average-case hardness of SZK implies the existence of distributional collision resistant hash functions.

Next, we define the *Jensen-Shannon Divergence*. To start with, recall that the KL-divergence between two distributions P and Q is defined³ as $\text{KL}(P||Q) = \sum_{y \in \mathcal{Y}} P_y \log(P_y/Q_y)$. Also, given distributions P_0 and P_1 we define the distribution $\frac{1}{2}P_0 + \frac{1}{2}P_1$ as the distribution obtained by sampling a random coin $b \in \{0, 1\}$ and outputting a sample y from P_b (indeed, this notation corresponds to arithmetic operations on the probability mass functions). The Jensen-Shannon divergence measures the mutual information between b and y .

Definition 1.3 (Jensen-Shannon Divergence). *The Jensen-Shannon divergence between two distributions P and Q is defined as*

$$\text{JS}(P, Q) = \frac{1}{2} \text{KL}\left(P \left\| \frac{P+Q}{2} \right.\right) + \frac{1}{2} \text{KL}\left(Q \left\| \frac{P+Q}{2} \right.\right).$$

The JENSEN-SHANNON DIVERGENCE PROBLEM with promise (α, β) , denoted $\text{JSP}^{\alpha,\beta}$, is defined analogously to $\text{SDP}^{\alpha,\beta}$, but with respect to JS rather than SD.

³ To be more precise, in this definition we view $0 \cdot \log \frac{0}{0}$ as 0 and define the KL-divergence to be ∞ if the support of P is not contained in that of Q .

The Jensen-Shannon divergence enjoys a couple of important properties (in our context) that the KL-divergence lacks: it is symmetric and bounded. Both triangular discrimination and Jensen-Shannon divergence (as well as statistical distance and KL-divergence) are types of f -divergences, a central concept in information theory (see [PW17, Section 6] and references therein). They are both non-negative and bounded by one.⁴ Finally, the Jensen-Shannon divergence is a metric, while the triangular discrimination is a square of a metric.

With these notions of distance and corresponding computational problems in hand, we are almost ready to state our first set of results. Before doing so, we introduce an additional useful technical definition.

Definition 1.4 (Separated functions). *Let $g: \mathbb{N} \rightarrow [0, 1]$. A pair of $\text{poly}(n)$ -time computable functions (α, β) , where $\alpha = \alpha(n) \in [0, 1]$ and $\beta = \beta(n) \in [0, 1]$, is g -separated if $\alpha(n) \geq \beta(n) + g(n)$ for every $n \in \mathbb{N}$.*

We denote by $(1/\text{poly})$ -separated the set of all pairs of functions that are $(1/p)$ -separated for some polynomial p . Similarly, we denote by $(1/\log)$ -separated the set of all pairs of functions that are $(1/(c \log))$ -separated for some constant $c > 0$.

We can now state our first set of results: that both TDP and JSP, with a noticeable gap, are SZK complete.

Theorem 1.5. *Let (α, β) be $(1/\text{poly})$ -separated functions such that there exists a constant $\varepsilon \in (0, 1/2)$ such that $2^{-n^{1/2-\varepsilon}} \leq \beta(n)$ and $\alpha(n) \leq 1 - 2^{-n^{1/2-\varepsilon}}$, for every $n \in \mathbb{N}$. Then, $\text{TDP}^{\alpha, \beta}$ is SZK complete.*

Theorem 1.6. *For (α, β) as in Theorem 1.5, the problem $\text{JSP}^{\alpha, \beta}$ is SZK complete.*

The restriction on $2^{-n^{1/2-\varepsilon}} \leq \beta(n)$ and $\alpha(n) \leq 1 - 2^{-n^{1/2-\varepsilon}}$ should be interpreted as a non-degeneracy requirement (which we did not attempt to optimize), where we note that some restriction seems inherent. Moreover, we can actually decouple the assumptions in Theorems 1.5 and 1.6 as follows. To show that $\text{TDP}^{\alpha, \beta}$ and $\text{JSP}^{\alpha, \beta}$ are SZK-hard, only the non-degeneracy assumption (i.e., $2^{-n^{1/2-\varepsilon}} \leq \beta(n)$ and $\alpha(n) \leq 1 - 2^{-n^{1/2-\varepsilon}}$) is needed. On the other hand, to show that these problems are in SZK we only require that (α, β) are $(1/\text{poly})$ -separated.

Note that in particular, Theorems 1.5 and 1.6 imply polarization lemmas for both TD and JS. For example, for triangular discrimination, since $\text{TDP}^{\alpha, \beta} \in \text{SZK}$ and $\text{TDP}^{1-2^{-k}, 2^{-k}}$ is SZK-hard, one can reduce the former to the latter.

Beyond showing polarization for triangular discrimination, Theorem 1.5 has implications regarding the question of polarizing statistical distance, which was our original motivation. It is known that the triangular discrimination is sandwiched between the statistical distance and its square; namely, for every two

⁴ In the literature these distances are sometimes defined to be twice as much as our definitions. In our context, it is natural to have the distances bounded by one.

distributions P and Q it holds that (see [Top00, Eq. (11)]):

$$\text{SD}(P, Q)^2 \leq \text{TD}(P, Q) \leq \text{SD}(P, Q) \quad (1.1)$$

Thus, the problem $\text{SDP}^{\alpha, \beta}$ is immediately reducible to $\text{TDP}^{\alpha^2, \beta}$, which Theorem 1.5 shows to be SZK-complete, as long as the gap between α^2 and β is noticeable. Specifically, we have the following corollary.

Corollary 1.7. *Let (α, β) be as in Theorem 1.5, with the exception that (α^2, β) are $(1/\text{poly})$ -separated (note that here α is squared). Then, the promise problem $\text{SDP}^{\alpha, \beta}$ is SZK complete.*

We highlight two implications of Theorem 1.5 and Corollary 1.7 (which were also briefly mentioned above).

Polarization with Inverse Polynomial Gap. Observe that Corollary 1.7 implies polarization of statistical distance in a regime in which α and β are functions of n , the output length of the two circuits, and α^2 and β are only separated by an inverse polynomial. This is in contrast to most prior works which focus on α and β that are constants. In particular, Sahai and Vadhan’s [SV03] proof of the polarization lemma focuses on constant α and β and can be extended to handle an inverse logarithmic gap, but does not seem to extend to an inverse polynomial gap.⁵ Corollary 1.7 does yield such a result, by relying on a somewhat different approach.

Polarization Beyond $\alpha^2 > \beta$. Theorem 1.5 can sometimes go beyond the requirement that $\alpha^2 > \beta$ for polarizing statistical distance. Specifically, it shows that any problem with noticeable gap in the triangular discrimination can be polarized. Indeed, there are distributions (P, Q) and (P', Q') with $\text{SD}(P, Q) > \text{SD}(P', Q') > \text{SD}(P, Q)^2$ but still $\text{TD}(P, Q) > \text{TD}(P', Q')$.⁶ Circuits generating such distributions were until now not known to be in the polarizing regime, but can now be polarized by combining Theorem 1.5 and Eq. (1.1).

1.1.2 From Statistical Difference to One-way Functions

We continue our study of the STATISTICAL DIFFERENCE PROBLEM, focusing on the regime where $\beta < \alpha$ (and in particular even when $\beta \geq \alpha^2$). We show that in this regime the $\text{SDP}^{\alpha, \beta}$ problem shares many important properties of SZK (although we fall short of actually showing that it lies in SZK—which is equivalent to polarization for any $\beta < \alpha$).

⁵ Actually, it was claimed in [GV11] that the [SV03] proof does extend to the setting of an inverse polynomial gap between α^2 and β but this claim was later retracted, see <http://www.wisdom.weizmann.ac.il/~oded/entropy.html>.

⁶ For example, for a parameter $\gamma \in [0, 1]$ consider the distributions R_0^γ and R_1^γ over $\{0, 1, 2\}$: R_b^γ puts γ mass on b and $1 - \gamma$ mass on 2. It holds that $\text{SD}(R_0^\gamma, R_1^\gamma) = \text{TD}(R_0^\gamma, R_1^\gamma) = \gamma$. If, say, $(P, Q) = (R_0^{1/2}, R_1^{1/2})$ and $(P', Q') = (R_0^{1/3}, R_1^{1/3})$, then $\text{SD}(P, Q) > \text{SD}(P', Q') > \text{SD}(P, Q)^2$ but $\text{TD}(P, Q) > \text{TD}(P', Q')$.

First, we show that similarly to SZK, the average-case hardness of $\text{SDP}^{\alpha,\beta}$ implies the existence of one-way functions. The fact that average-case hardness of SZK (or equivalently $\text{SDP}^{\alpha,\beta}$ for $\beta < \alpha^2$) implies the existence of one-way functions was shown by Ostrovsky [Ost91]. Indeed, our contribution is in showing that the weaker condition of $\beta < \alpha$ (rather than $\beta < \alpha^2$) suffices for this result.

Theorem 1.8. *Let (α, β) be $(1/\text{poly})$ -separated functions. Suppose that $\text{SDP}^{\alpha,\beta}$ is average-case hard. Then, there exists a one-way function.*

The question of constructing one-way functions from the (average-case) hardness of SDP is closely related to a result of Goldreich's [Gol90] showing that the existence of efficiently sampleable distributions that are statistically far but computationally indistinguishable implies the existence of one-way functions. Our proof of Theorem 1.8 allows us to re-derive the following strengthening of [Gol90], due to Naor and Rothblum [NR06, Theorem 4.1]: for any $(1/\text{poly})$ -separated (α, β) , the existence of efficiently sampleable distributions whose statistical distance is α but no efficient algorithm can distinguish between them with advantage more than β , implies the existence of one-way functions. See further discussion in Theorem 2.1.

1.1.3 Interactive Proof for Statistical Distance Approximation

As our last main result, we construct a new interactive protocol that lets a verifier estimate the statistical distance between two given circuits up to any noticeable precision.

Theorem 1.9. *There exists a constant-round public-coin interactive protocol between a prover and a verifier that, given as input a pair of circuits (C_0, C_1) , a claim $\Delta \in [0, 1]$ for their statistical distance, and a tolerance parameter $\delta \in [0, 1]$, satisfies the following properties:*

- **Completeness:** *If $\text{SD}(C_0, C_1) = \Delta$, then the verifier accepts with probability at least $2/3$ when interacting with the honest prover.*
- **Soundness:** *If $|\text{SD}(C_0, C_1) - \Delta| \geq \delta$, then when interacting with any (possibly cheating) prover, the verifier accepts with probability at most $1/3$.*
- **Efficiency:** *The verifier runs in time $\text{poly}(|C_0|, |C_1|, 1/\delta)$.*

(As usual the completeness and soundness errors can be reduced by applying parallel repetition. We can also achieve perfect completeness using a result from [FGM⁺89].)

Theorem 1.9 is actually equivalent to the following statement.

Theorem 1.10 ([BL13, Theorem 6], [BBF16, Theorem 2]). *For any (α, β) that are $(1/\text{poly})$ -separated, it holds that $\text{SDP}^{\alpha,\beta} \in \text{AM} \cap \text{coAM}$.⁷*

⁷ Recall that AM is the class of problems that have constant-round public-coin interactive proofs. coAM is simply the complement of AM.

It is believed that $\text{AM} \cap \text{coAM}$ lies just above SZK , and if we could show that $\text{SDP}^{\alpha, \beta}$ is in SZK , that would imply SD polarization for such α and β .

Since Theorem 1.9 can be derived from existing results in the literature, we view our main contribution to be the proof which is via a single protocol that we find to be cleaner and more direct than alternate approaches.

Going into a bit more detail, [BL13, BBF16]’s proofs are in fact a combination of two separate constant-round protocols. The first protocol is meant to show that $\text{SDP}^{\alpha, \beta} \in \text{AM}$ and follows directly by taking the interactive proof for SDP presented by Sahai and Vadhan (which has completeness error $(1 - \alpha)/2$ and soundness error $(1 + \beta)/2$), and applying parallel repetition (and the private-coin to public-coin transformation of [GS89]).

The second protocol is meant to show that $\text{SDP}^{\alpha, \beta} \in \text{coAM}$, and is based on a protocol by Bhatnagar, Bogdanov, and Mossel [BBM11]. Another approach for proving that $\text{SDP}^{\alpha, \beta} \in \text{coAM}$ is by combining results of [GVW02] and [SV03]. Goldreich, Vadhan and Wigderson [GVW02] showed that problems with laconic interactive proofs, that is proofs where the communication from the prover to the verifier is small, have coAM proofs. Sahai and Vadhan [SV03], as described earlier, showed that $\text{SDP}^{\alpha, \beta}$, and SZK in general, has an interactive proof where the prover communicates a single bit. Combining these results immediately gives a coAM protocol for $\text{SDP}^{\alpha, \beta}$ when (α, β) are $\Omega(1)$ -separated. As for (α, β) that are only $(1/\text{poly})$ -separated, while the [GVW02] result as-stated does not suffice, it seems that their protocol can be adapted to handle this case as well.⁸

As mentioned above, we give a different, and direct, proof of Theorem 1.9 that we find to be simpler and more natural than the above approach. In particular, our proof utilizes the techniques developed for our other results, which enable us to give a single and more general protocol—one that approximates the statistical difference (as in Theorem 1.9), rather than just deciding if that distance is large or small.

At a very high level, our protocol may be viewed as an application of the set-lower-bound-based techniques of Akavia et al [AGGM06] or Bogdanov and Brzuska [BB15] to our construction of a one-way function from the average-case hardness of SDP (i.e., Theorem 1.8), though there are technical differences in our setting. Both these papers show how to construct a coAM protocol for any language that can be reduced, to inverting a *size-verifiable* one-way function.⁹ While we do not know how to reduce solving SDP in the worst-case to inverting any specific function, we make use of the fact that associated with each instance of SDP , there is an *instance-dependent* function [OW93], that is size-verifiable on the average.

⁸ In more detail, the [GVW02] result is stated for protocols in which the gap between completeness and soundness is constant (specifically $1/3$). In case α and β are only $1/\text{poly}$ -separated, the [SV03] protocol only has a $1/\text{poly}$ gap (and we cannot afford repetition since it will increase the communication). Nevertheless, by inspecting the [GVW02] proof, it seems as though it can be adapted to cover any noticeable gap.

⁹ Informally, a function f is size-verifiable if given an output $y = f(x)$, there exists an AM protocol to estimate $|f^{-1}(y)|$.

1.2 Additional Related Works

Barriers to Improved Polarization. Holenstein and Renner [HR05] show that in a limited model dubbed “oblivious polarization”, the condition $\alpha^2 > \beta$ on the statistical distance is necessary for polarizing statistical distance.¹⁰ All the past polarization reductions fit in this framework and so do ours. Specifically, Holenstein and Renner show distributions where $\alpha^2 < \beta$ and cannot be polarized in this model. We show a condition that suffices for polarization, even for distributions where $\alpha^2 \leq \beta$. This does not contradict the [HR05] result because their distributions do not satisfy this condition.

In a more general model, [LZ17,CGVZ18] showed lower bounds for SZK-related distribution manipulation tasks. The model they consider allows the reduction arbitrary oracle access to the circuits that sample the distributions, as opposed to the more restricted model of oblivious polarization. In this model, Lovett and Zhang [LZ17] show that efficient entropy reversal is impossible¹¹, and Chen, Göös, Vadhan and Zhang [CGVZ18] showed that entropy flattening requires $\Omega(n^2)$ invocations to the underlying circuit. Showing lower bounds for polarization in this more general model remains an interesting open question.

Polarization for other Notions of Distance. In the process of characterizing zero-knowledge in the help model, Ben-Or and Gutfreund [BG03] and Chailloux et al. [CCKV08] gave a polarization procedure that considers two different distances for every $(1/\log)$ -separated $\alpha > \beta$: if the statistical distance is at most β , then it decreases to 2^{-k} ; and if the *mutual disjointness*¹² is at least α , then it increases to $1 - 2^{-k}$. Fehr and Vaudenay [FV17] raise the question of polarization for the fidelity measure¹³ but leave resolving it as an open problem (see Section 2.3.3 for details).

SDP and Cryptography. We show that average-case hardness of $\text{SDP}^{\alpha,\beta}$ implies one-way functions. In the reverse direction, Bitansky et al. [BDV17] show that one-way functions do not imply even worst-case hardness of $\text{SDP}^{\alpha,\beta}$ in a black-box manner for any $(1/\text{poly})$ -separated α, β .¹⁴

¹⁰ Roughly speaking, an oblivious polarization is a randomized procedure to polarize without invoking the circuits; it takes as input a bit σ and an integer k , and outputs a sequence of bits $(b_1^\sigma, \dots, b_\ell^\sigma)$ and a string r^σ . Given a pair of circuits (C_0, C_1) , such a procedure defines a pair of circuits (D_0, D_1) as follows: D_σ samples $(b_1^\sigma, \dots, b_\ell^\sigma)$ and r^σ and outputs $(C_{b_1^\sigma}, \dots, C_{b_\ell^\sigma}, r^\sigma)$. We are guaranteed that if $\text{SD}(C_0, C_1) \geq \alpha$, then $\text{SD}(D_0, D_1) \geq 1 - 2^{-k}$, and if $\text{SD}(C_0, C_1) \leq \beta$, then $\text{SD}(D_0, D_1) \leq 2^{-k}$.

¹¹ Entropy reversal refers to the task of given circuit C and parameter t output (C', t') such that when $H(C) > t$, then $H(C') < t' - 1$ and if $H(C) < t - 1$, then $H(C') > t'$.

¹² For an ordered pair of distributions P and Q , their disjointness is $\text{Disj}(P, Q) = \Pr_{y \sim P}[y \notin \text{Supp}(Q)]$, and their mutual disjointness is $\text{MutDisj}(P, Q) = \min(\text{Disj}(P, Q), \text{Disj}(Q, P))$.

¹³ For two distributions P and Q , their fidelity is defined as $\text{Fidelity}(P, Q) = \sum_y \sqrt{P_y \cdot Q_y}$.

¹⁴ While [BDV17] state the result for constant α, β , the construction and analysis extend to our setting.

2 Techniques

We begin in Section 2.1 by describing how to construct a one-way function from the average-case hardness of SD with any noticeable gap (Theorem 1.8). The techniques used there are also central in our interactive protocol for SD estimation (Theorem 1.9), which is described in Section 2.2, as well as in our proof that triangular discrimination and Jensen-Shannon divergence are SZK complete (Theorems 1.5 and 1.6), which are outlined in Section 2.3 below.

2.1 One-Way Function From Statistical Difference with Any Noticeable Gap

We first show the existence of *distributionally* one-way functions. Namely, an efficiently computable function f for which it is hard to sample a uniformly random pre-image for a random output y (rather than an arbitrary pre-image as in a standard one-way function). This suffices since Impagliazzo and Luby [IL89] showed how to convert a distributionally one-way function into a standard one.

Assume that we are given a distribution over a pair of circuits (C_0, C_1) such that it is hard to distinguish between the cases $\text{SD}(C_0, C_1) \geq \alpha$ or $\text{SD}(C_0, C_1) \leq \beta$, for some $\alpha > \beta + 1/\text{poly}$. A natural candidate for a one-way function is the (efficiently computable) function

$$f_{C_0, C_1}(b, x) = C_b(x). \quad (2.1)$$

Namely, f is parameterized by the circuits (C_0, C_1) (which are to be sampled according to the hard distribution), and the bit b chooses which of the two circuits would be evaluated on the string x . This function appears throughout the SZK literature (e.g., it corresponds to the verifier's message in the SDP protocol of [SV03]).

Assume that f is not distributionally one-way, and let A be an algorithm that given (C_0, C_1) and a random input y —sampled by first drawing a uniformly random bit b and a string x and then computing $y = C_b(x)$ —outputs a uniformly random element (b', x') from the set $f_{C_0, C_1}^{-1}(y) = \{(b, x) : C_b(x) = y\}$. For simplicity, we assume that A is a perfect distributional inverter, that is for *every fixed* (C_0, C_1, y) it outputs uniformly random elements of $f_{C_0, C_1}^{-1}(y)$.

Arguably, the most natural approach for distinguishing between the cases of high or low statistical distance given the two circuits and the inverter, is to choose x and b at random, invoke the inverter to obtain (b', x') , and check whether $b = b'$. Indeed, if $\text{SD}(C_0, C_1) = 1$, then $\Pr[b = b'] = 1$, and if $\text{SD}(C_0, C_1) = 0$, then $\Pr[b = b'] = \frac{1}{2}$. Thus, we can distinguish between the cases with constant advantage.

But what happens when the gap in the statistical distance is smaller? To analyze this case we want to better understand the quantity $\Pr[b = b']$. It turns out that this quantity is characterized by the triangular discrimination between

the circuits. Let P_b denote the output distribution of C_b . Using elementary manipulations (and the fact that $\frac{1}{2}(P_0 + P_1)$ is a distribution), it holds that¹⁵

$$\begin{aligned}
\Pr[b = b'] &= \frac{1}{2} \Pr_{y \sim P_0}[b' = 0] + \frac{1}{2} \Pr_{y \sim P_1}[b' = 1] \\
&= \frac{1}{2} \sum_y \frac{P_0(y)^2 + P_1(y)^2}{P_0(y) + P_1(y)} \\
&= \frac{1}{4} \sum_y \frac{(P_0(y) + P_1(y))^2}{P_0(y) + P_1(y)} + \frac{1}{4} \sum_y \frac{(P_0(y) - P_1(y))^2}{P_0(y) + P_1(y)} \\
&= \frac{1}{2} + \frac{1}{4} \sum_y \frac{(P_0(y) - P_1(y))^2}{P_0(y) + P_1(y)} \\
&= \frac{1 + \text{TD}(C_0, C_1)}{2}.
\end{aligned} \tag{2.2}$$

Based on the general bounds between triangular discrimination and statistical distance (Eq. (1.1)), which are known to be tight, all we are guaranteed is

$$\begin{aligned}
\text{SD}(C_0, C_1) \geq \alpha &\implies \Pr[b = b'] \geq \frac{1 + \alpha^2}{2} \\
\text{SD}(C_0, C_1) \leq \beta &\implies \Pr[b = b'] \leq \frac{1 + \beta}{2}.
\end{aligned}$$

So, this approach is limited to settings in which $\alpha^2 > \beta$.

To overcome this limitation we want to find a quantity that is more tightly characterized by the statistical distance of the circuits. This quantity, which we call *imbalance*, will be central in all of the proofs in this work. The imbalance measures how likely it is that an output string y was generated from C_1 versus C_0 . Formally,

$$\theta_y \triangleq \Pr[b = 1|y] - \Pr[b = 0|y] = \frac{P_1(y) - P_0(y)}{P_1(y) + P_0(y)}. \tag{2.3}$$

Elementary manipulations yields that

$$\begin{aligned}
\text{SD}(C_0, C_1) &= \frac{1}{2} \sum_y |P_1(y) - P_0(y)| \\
&= \sum_y \frac{1}{2} (P_1(y) + P_0(y)) \cdot \frac{|P_1(y) - P_0(y)|}{P_1(y) + P_0(y)} \\
&= \mathbb{E}_{y \sim (\frac{1}{2}P_0 + \frac{1}{2}P_1)} [|\theta_y|].
\end{aligned} \tag{2.4}$$

¹⁵ In Section 1 we used P_y to denote the probability mass a distribution P puts on an element y , while here we use $P(y)$. In the rest of this work we choose which notation to use based on readability and context.

(Recall that y is sampled by first drawing a uniform random bit b and a string x , and setting $y = C_b(x)$. Hence, using the notation that P_b denotes the output distributions of the circuit C_b , the marginal distribution of y is $\frac{1}{2}P_0 + \frac{1}{2}P_1$.)

Eq. (2.4) naturally gives rise to the following algorithm for approximating $\text{SD}(C_0, C_1)$:

Algorithm to estimate $\text{SD}(C_0, C_1)$ using the inverter A :

1. Sample polynomially many $y_1, \dots, y_t$ independently from $\frac{1}{2}P_0 + \frac{1}{2}P_1$.
2. For every y_i :
 - (a) Call $A(y_i)$ polynomially many times to get $b'_1, \dots, b'_k$.
 - (b) Let m be the number of ones in $b'_1, \dots, b'_k$.
 - (c) Set $p_1 = m/k$, $p_0 = (k - m)/k$ and $\hat{\theta}_i = p_1 - p_0$.
3. Return $\frac{1}{t} \sum_{i=1}^t |\hat{\theta}_i|$.

The quantities p_1 and p_0 are in fact the empirical distribution of b conditioned on y , computed using k samples. By choosing large enough k , we get that $(p_1, p_0) \approx (\Pr[b = 1|y], \Pr[b = 0|y])$ and so $\hat{\theta}_i \approx \theta_{y_i}$. By then choosing large enough t , we get that $\frac{1}{t} \sum_{i=1}^t |\hat{\theta}_i| \approx \text{SD}(C_0, C_1)$. Hence, we can distinguish between the cases $\text{SD}(C_0, C_1) \geq \alpha$ or $\text{SD}(C_0, C_1) \leq \beta$, for any $\alpha > \beta + 1/\text{poly}$.

Essentially the same proof continues to work if A is not a perfect distributional inverter, but is close enough to being so—that is, on input y its output distribution is close to being uniform over $f^{-1}(y)$ for most (but not all) tuples C_0, C_1, y .

The above proof strategy also yields a new proof for the strengthening of [Gol90] by Naor and Rothblum [NR06].¹⁶ See Theorem 2.1 below for a discussion about the differences between our techniques and those of [NR06].

Distributional Collision Resistant Hash Function. As a matter of fact, the above proof also shows that the average-case hardness of $\text{SDP}^{\alpha, \beta}$ also implies that the function $f_{C_0, C_1}(b, x) = C_b(x)$ is a distributional k -multi-collision¹⁷ resistant hash function, for $k = O\left(\frac{\log n}{(\alpha - \beta)^2}\right)$. That is, for a random output y of f , it is difficult to find k random preimages of y . This is because access to such a set of k random pre-images of random y_i 's is all we use the inverter A for in the above reduction, and it could handily be replaced with a k -distributional multi-collision finder.

Remark 2.1 (Comparison to [NR06]) *Naor and Rothblum's proof implicitly attempts to approximate the maximal likelihood bit of y ; that is, the bit b_{ml}*

¹⁶ Namely, that for any $(1/\text{poly})$ -separated (α, β) , the existence of efficiently sampleable distributions whose statistical distance is α but no efficient algorithm can distinguish between them with advantage more than β , implies the existence of one-way functions.

¹⁷ Multi-collision hash functions, recently considered in several works [KNY17, KNY18, BKP18, BDRV18], are hash functions for which it is hard to find multiple inputs that all hash to the same output.

such that $\Pr[b = b_{ml}|y] > \Pr[b = 1 - b_{ml}|y]$ (breaking ties arbitrarily). Indeed, the maximal likelihood bit, as shown by [SV03], is closely related to the statistical distance:

$$\Pr[b = b_{ml}] = \frac{1 + \text{SD}(C_0, C_1)}{2}. \quad (2.5)$$

To approximate b_{ml} , [NR06] make, like us, many calls to $A(y)$, and take the majority of the answered bits. The idea is that when the statistical distance is large, the majority is likely to be b_{ml} , and when the statistical distance is small, the majority is equally likely to be b_{ml} or $1 - b_{ml}$.

To formally prove this intuition, it must hold that if $\text{SD}(C_0, C_1)$ is large, then $\Pr[b = b_{ml}|y] - \Pr[b = 1 - b_{ml}|y]$ is sufficiently large; putting in our terminology and using Eq. (2.4), if $\mathbb{E}_y[|\theta_y|]$ is sufficiently large, then $|\theta_y|$ should be large for a random y (and the opposite should hold if $\text{SD}(C_0, C_1)$ is small). While these statements are true, in order to prove them, [NR06]’s analysis involves some work which results in a more complicated analysis.

We manage to avoid such complications by using the imbalance θ_y and its characterization of statistical distance (Eq. (2.4)). Furthermore, [NR06]’s approach only attempts to distinguish between the cases when $\text{SD}(C_0, C_1)$ is high or low, while our approach generalizes to approximate $\text{SD}(C_0, C_1)$. Lastly, Naor and Rothblum do not construct one-way functions based on the average-case hardness of $\text{SDP}^{\alpha, \beta}$ with any noticeable gap as we do. Using their technique to do so seems to require additional work—work that our analysis significantly simplifies.

2.2 Interactive Proof for Statistical Distance Approximation

We proceed to describe a constant-round public-coin protocol in which a computationally unbounded prover convinces a computationally bounded verifier that the statistical difference of a given pair of circuits is what the prover claims it to be, up to any inverse polynomial (additive) error. Such a protocol simultaneously establishes the inclusion of $\text{SDP}^{\alpha, \beta}$ in both AM and coAM for any $\alpha > \beta + 1/\text{poly}$.

Our starting point is the algorithm we described above that used a one-way function inverter to estimate the statistical distance. Specifically, that algorithm used the inverter to estimate θ_y for random y ’s, and then applied Eq. (2.4). We would like to use the prover, instead of the inverter, to achieve the same task.

In our protocol, the verifier draws polynomially many y ’s and sends them to the prover. The prover responds with values $\hat{\theta}_i$ ’s, which it claims are the genuine θ_{y_i} ’s. But how can the verifier trust that the prover sent the correct values? In the reduction in Section 2.1, we used k many samples of b conditioned on y to estimate b ’s true distribution. A standard concentration bound shows that as k grows, the number of ones out of $b_1, \dots, b_k$, all sampled from $(b|y)$, is very close to $\Pr[b = 1|y] \cdot k$. Similarly, the number of zeros is very close to $\Pr[b = 0|y] \cdot k$. Consider the following *typical set* for any fixed y and arbitrary value θ :

$$\mathcal{T}_y^{k, \theta} = \left\{ (b_1, x_1, b_2, x_2, \dots, b_k, x_k) \left| \begin{array}{l} C_{b_i}(x_i) = y \text{ for all } i, \\ \text{and } \frac{\sum_{i=1}^k b_i - \sum_{i=1}^k (1-b_i)}{k} \approx \theta \end{array} \right. \right\}.$$

Namely, $\mathcal{T}_y^{k,\theta}$ contains every k -tuple of (b_i, x_i) such that all map to y , and each tuple can be used to estimate θ well—the difference between the number of ones and the number of zeros, normalized by k , is close to θ . Also consider the *pre-image* set of y : $\mathcal{I}_y = \{(b, x) \mid C_b(x) = y\}$. Since as k grows the estimation of θ_y improves, we expect that $\mathcal{T}_y^{k,\theta_y}$ —the typical set of y with the value θ_y —to contain almost all tuples. Indeed, standard concentration bounds show that

$$\frac{|\mathcal{T}_y^{k,\theta_y}|}{|\mathcal{I}_y|^k} \geq 1 - e^{-\Omega(k)}. \quad (2.6)$$

On the other hand, the sets $\mathcal{T}_y^{k,\theta'}$, corresponding to values θ' that are far from θ_y , should be almost empty. Indeed, if $|\theta' - \theta_y| \geq \Omega(1)$, then,

$$\frac{|\mathcal{T}_y^{k,\theta'}|}{|\mathcal{I}_y|^k} \leq e^{-\Omega(k)}. \quad (2.7)$$

So, for the verifier to be convinced that the value $\hat{\theta}$ sent by the prover is close to θ_y , the prover can prove that the typical set $\mathcal{T}_y^{k,\hat{\theta}}$ is large. To do so, the parties will use the public-coin constant round protocol for set lower-bound of [GS89], which enables the prover to assert statements of the form “the size of the set $\mathcal{S}$ is at least s ”.

However, there is still one hurdle to overcome. The typical set $\mathcal{T}_y^{k,\theta_y}$ is only large *relative* to $|\mathcal{I}_y|^k$. Since we do not know how to compute $|\mathcal{I}_y|$ it is unclear what should be the size s that we run the set lower-bound protocol with. Our approach for bypassing this issue is as follows. First observe that the *expected value*, over a random y , of the logarithm of the size of $\mathcal{I}_y$ is the entropy¹⁸ of (b, x) given y . Namely,

$$\mathbb{E}_y[\log|\mathcal{I}_y|] = H(B, X|Y), \quad (2.8)$$

where the jointly distributed random variables (B, X, Y) take the values of randomly drawn (x, b, y) . Thus, if we draw t independent elements $y_1, \dots, y_t$, the average of $\log|\mathcal{I}_{y_i}|$ gets closer to $t \cdot H(B, X|Y)$, as t grows. Specifically,

$$\Pr \left[\prod_{i=1}^t |\mathcal{I}_{y_i}| \approx 2^{t \cdot H(B, X|Y)} \right] \geq 1 - e^{-\Omega(t/n^2)}, \quad (2.9)$$

where n denotes the output length of the given circuits. For large enough t , we can thus assume that the size of this product set is approximately $2^{t \cdot H(B, X|Y)}$, and run the set lower bound protocol for all the y_i ’s together. That is, we ask

¹⁸ Recall that the entropy of a random variable X over $\mathcal{X}$ is defined as $H(X) = \sum_{x \in \mathcal{X}} \Pr[X = x] \log(1/\Pr[X = x])$. The conditional entropy of X given Y is $H(X|Y) = \mathbb{E}_{y \sim Y}[H(X|Y = y)]$.

the prover to send t estimates $(\hat{\theta}_1, \dots, \hat{\theta}_t)$ for the values $(\theta_{y_1}, \dots, \theta_{y_t})$, and prove that the size of the product set $\mathcal{T}_{y_1}^{k, \hat{\theta}_1} \times \dots \times \mathcal{T}_{y_t}^{k, \hat{\theta}_t}$ is almost $2^{t \cdot H(B, X|Y)}$.

So far we have reduced knowing the size of $\mathcal{I}_y$ to knowing $H(B, X|Y)$, but again it seems difficult for the verifier to compute this quantity on its own. Actually, standard entropy manipulations show that

$$H(B, X|Y) = (m + 1) - H(Y),$$

where m denotes the input length of the given circuits. It thus suffices to approximate $H(Y)$. Recall that y is the output of the circuit that maps (x, b) to $C_b(x)$, so Y is drawn according to an output distribution of a known circuit. Luckily, Goldreich, Sahai and Vadhan [GSV99] showed that approximating the output entropy of a given circuit is in NISZK, and thus has a constant-round public-coin protocol (since $\text{NISZK} \subseteq \text{AM} \cap \text{coAM}$).

To conclude, we describe the entirety of our protocol, which proves Theorem 1.9.

Protocol to approximate $\text{SD}(C_0, C_1)$, given the circuits (C_0, C_1) as input:

1. First, the prover sends the verifier a claim $\hat{H}$ of the value of $H(Y)$.
 2. The parties execute [GSV99]’s protocol to convince the verifier that this claim—that $\hat{H} \approx H(Y)$ —is correct.
 3. The verifier uses $\hat{H}$ to compute $\hat{H}(B, X|Y)$ as $((m + 1) - \hat{H})$.
 4. The verifier samples $y_1, \dots, y_t$ from $\frac{C_0 + C_1}{2}$ and sends them to the prover.
 5. The prover responds with $\hat{\theta}_1, \dots, \hat{\theta}_t$ as claims for the values $\theta_{y_1}, \dots, \theta_{y_t}$.
 6. The parties run a set lower-bound protocol to prove that the set $\mathcal{T}_{y_1}^{\hat{\theta}_1, k} \times \dots \times \mathcal{T}_{y_t}^{\hat{\theta}_t, k}$ is almost as large as $(\mathcal{I}_{y_1} \times \dots \times \mathcal{I}_{y_t})^k$.
 - Here, they use $2^{tk\hat{H}(B, X|Y)}$ as a proxy for $(|\mathcal{I}_{y_1}| \cdot \dots \cdot |\mathcal{I}_{y_t}|)^k$.
 7. If the verifier has not rejected so far, it outputs $\frac{1}{t} \sum_{i=1}^t |\hat{\theta}_i|$.
-

2.3 TDP and JSP are SZK-Complete

We show that both $\text{TDP}^{\alpha, \beta}$ and $\text{JSP}^{\alpha, \beta}$ with $\alpha > \beta + 1/\text{poly}$ are SZK-complete. Since the proof of the former uses that of the latter we start by giving an outline that $\text{JSP}^{\alpha, \beta}$ is SZK-complete.

2.3.1 JENSEN-SHANNON DIVERGENCE PROBLEM is SZK-complete

We need to show that $\text{JSP}^{\alpha, \beta}$ with $\alpha > \beta + 1/\text{poly}$ is both in SZK and SZK-hard. In both parts we use the following characterization of the Jensen-Shannon divergence, which follows from its definition. Given a pair of circuits C_0 and C_1 , consider the jointly distributed random variables (B, X, Y) , where B is a

uniformly random bit, X is a uniformly random string and $Y = C_B(X)$. Then, it follows from some elementary manipulation that:

$$\text{JS}(C_0, C_1) = 1 - H(B|Y). \quad (2.10)$$

We use this characterization to tie JENSEN-SHANNON DIVERGENCE PROBLEM to another SZK-complete problem—the ENTROPY DIFFERENCE PROBLEM (EDP) with a gap function g . The input to EDP^g is also a pair of circuits C_0 and C_1 . YES instances are those in which the entropy gap $H(C_0) - H(C_1)$ is at least $g(n)$ (where n is the output length of the circuits) and NO instances are those in which the gap is at most $-g(n)$. Goldreich and Vadhan [GV99] showed that EDP^g is SZK-complete for any noticeable function g . Our proof that $\text{JSP}^{\alpha, \beta}$ is SZK-complete closely follows the reduction from the reverse problem of SDP (i.e., in which YES instances are distributions that are statistically *close*) to EDP [Vad99, Section 4.4].

$\text{JSP}^{\alpha, \beta}$ is in SZK: We reduce $\text{JSP}^{\alpha, \beta}$ to $\text{ED}^{(\alpha - \beta)/2}$. Given C_0 and C_1 , the reduction outputs a pair of circuits D_0 and D_1 such that D_1 outputs a sample from (B, Y) and D_0 outputs a sample from (B', Y) , where B' is an independent random bit with $H(B') = 1 - \frac{\alpha + \beta}{2}$. The chain rule for entropy¹⁹ implies that

$$H(D_0) - H(D_1) = 1 - \frac{\alpha + \beta}{2} - H(B|Y) = \text{JS}(C_0, C_1) - \frac{\alpha + \beta}{2},$$

where the second equality follows from Eq. (2.10). Thus, if $\text{JS}(C_0, C_1) \geq \alpha$, then $H(D_0) - H(D_1) \geq \frac{\alpha - \beta}{2}$; and if $\text{JS}(C_0, C_1) \leq \beta$, then $H(D_0) - H(D_1) \leq -\frac{\alpha - \beta}{2}$. And since $\text{ED}^{(\alpha - \beta)/2} \in \text{SZK}$, we get that $\text{JSP}^{\alpha, \beta} \in \text{SZK}$.

$\text{JSP}^{\alpha, \beta}$ is SZK-hard: We reduce $\text{SDP}^{1 - 2^{-k}, 2^{-k}}$ to the problem $\text{JSP}^{\alpha, \beta}$, for some large enough k . This is sufficient since $\text{SDP}^{1 - 2^{-k}, 2^{-k}}$ is known to be SZK-hard [SV03].²⁰ In the presentation of related results in his thesis, Vadhan relates the statistical distance of the circuits to the entropy of B given Y [Vad99, Claim 4.4.2]. For example, if $\text{SD}(C_0, C_1) = 0$ (i.e., the distributions are identical), then $B|Y$ is a uniformly random bit, and so $H(B|Y) = 1$; and if $\text{SD}(C_0, C_1) = 1$ (i.e., the distributions are disjoint), then B is completely determined by Y , and so $H(B|Y) = 0$. More generally, Vadhan showed that if $\text{SD}(C_0, C_1) = \delta$, then²¹

$$1 - \delta \leq H(B|Y) \leq h\left(\frac{1 + \delta}{2}\right). \quad (2.11)$$

¹⁹ For a jointly distributed random variables X and Y , it holds that $H(X, Y) = H(X) + H(Y|X)$.

²⁰ For the simplicity of presentation, we are ignoring subtle details about the relation of k to the output length of the circuits. See the full version for the formal proof.

²¹ The function h is the binary entropy function. That is, $h(p) = -p \log(p) - (1 - p) \log(1 - p)$ is the entropy of a Bernoulli random variable with parameter p .

By taking k to be large enough (as a function of α and β), and applying Eqs. (2.10) and (2.11), we have that if $\text{SD}(C_0, C_1) \geq 1 - 2^{-k}$, then $\text{JS}(C_0, C_1) \geq \alpha$; and if $\text{SD}(C_0, C_1) \leq 2^{-k}$, then $\text{JS}(C_0, C_1) \leq \beta$. Thus, the desired reduction is simply the identity function that outputs the input circuits.

2.3.2 TRIANGULAR DISCRIMINATION PROBLEM is SZK-complete.

We need to show that $\text{TDP}^{\alpha, \beta}$ with $\alpha > \beta + 1/\text{poly}$ is both in SZK and SZK-hard. Showing the latter is very similar to showing that $\text{JSP}^{\alpha, \beta}$ is SZK-hard, but using Eq. (1.1) to relate the triangular discrimination to statistical distance (instead of Eq. (2.11) that relates the Jensen-Shannon divergence to statistical distance). We leave the formal details to the body of this paper and focus here on showing that $\text{TDP}^{\alpha, \beta}$ is in SZK.

A natural approach to show that $\text{TDP}^{\alpha, \beta}$ is in SZK is to follow Sahai and Vadhan's proof that $\text{SDP}^{2/3, 1/3}$ is in SZK. Specifically, a main ingredient in that proof is to polarize the statistical distance of the circuits (to reduce the simulation error). Indeed, if we can reduce $\text{TDP}^{\alpha, \beta}$ to, say, $\text{TDP}^{0.9, 0.1}$ by polarizing the triangular discrimination, then Eq. (1.1) would imply that we also reduce $\text{TDP}^{\alpha, \beta}$ to $\text{SDP}^{2/3, 1/3}$, which we know is in SZK.

We are indeed able to show such a polarization lemma for triangular discrimination (using similar techniques to [SV03]'s polarization lemma). However, this lemma only works when the gap between α and β is roughly $1/\log$. Actually, the polarization lemma of [SV03] also suffers the same limitation with respect to the gap between α^2 and β .

Still, we would like to handle also the case that the gap between α and β is only $1/\text{poly}$. To do so we take a slightly different approach. Specifically, we reduce $\text{TDP}^{\alpha, \beta}$ to $\text{JSP}^{\alpha', \beta'}$, where α' and β' are also noticeably separated.

An important step toward showing this reduction is to characterize the triangular discrimination and the Jensen-Shannon divergence via the imbalance θ_y (see Eq. (2.3)), as we already did for statistical distance. Recall that given $Y = y$, the random variable B takes the value 1 with probability $\frac{1+\theta_y}{2}$, and 0 otherwise. Hence, Eq. (2.10) can also be written as

$$\text{JS}(C_0, C_1) = 1 - \mathbb{E}_{y \sim Y} \left[h \left(\frac{1 + \theta_y}{2} \right) \right]. \quad (2.12)$$

As for the triangular discrimination, it follows from the definition that

$$\text{TD}(C_0, C_1) = \mathbb{E}_{y \sim Y} [\theta_y^2]. \quad (2.13)$$

Furthermore, by Taylor approximation, for small values of θ , it holds that

$$h \left(\frac{1 + \theta}{2} \right) \approx 1 - \theta^2. \quad (2.14)$$

As we can see, the above equations imply that if all the θ_y 's were small, a gap in the triangular discrimination would also imply a gap in the Jensen-Shannon divergence. Thus, we would like an operation that reduces all the θ_y .

The main technical tool we use to reduce θ_y is to consider the *convex combination* of the two input circuits. Given a pair of circuits C_0 and C_1 , consider the pair of circuits D_0 and D_1 such that $D_b = \lambda \cdot C_b + (1 - \lambda) \cdot \frac{C_0 + C_1}{2}$.²² Let Q_b denote the output distribution of D_b , and recall that P_b denotes the output distribution of C_b . We also let θ'_y be defined similarly to θ_y , but with respect to D_0 and D_1 (rather than C_0 and C_1). Using this notation, we have that $\theta_y = \frac{P_1(y) - P_0(y)}{P_1(y) + P_0(y)}$, and it may be seen that

$$\theta'_y = \frac{Q_1(y) - Q_0(y)}{Q_1(y) + Q_0(y)} = \lambda \cdot \theta_y. \quad (2.15)$$

So, our reduction chooses a sufficiently small λ , and outputs the circuits D_0 and D_1 . Some care is needed when choosing λ . Eqs. (2.13) and (2.15) yield that $\text{TD}(D_0, D_1) = \lambda^2 \cdot \text{TD}(C_0, C_1)$. Hence, the convex combination also shrinks the gap in triangular discrimination. We show that by choosing $\lambda \approx \sqrt{\alpha - \beta}$, the approximation error in Eq. (2.14) is smaller than the aforementioned shrinkage, and the reduction goes through. The resulting gap in the Jensen-Shannon divergence is roughly $(\alpha - \beta)^2$, which is noticeable by the assumption that $\alpha > \beta + 1/\text{poly}$.

This shows that $\text{TDP}^{\alpha, \beta}$ is in SZK if $\alpha > \beta + 1/\text{poly}$. By the relationship between TD and SD (Eq. (1.1)), this implies that $\text{SDP}^{\alpha, \beta}$ is in SZK if $\alpha^2 > \beta + 1/\text{poly}$. This, in turn, by the SZK-hardness of $\text{SDP}^{2/3, 1/3}$ and the known polarization lemma that applies for the same, implies polarization for statistical distance for any (α, β) such that $\alpha^2 > \beta + 1/\text{poly}$.

2.3.3 Reflections and an Open Problem

Many f -divergences of interest can be expressed as an expectation, over $y \sim Y$, of a simple function of θ_y . That is, an expression of the form $\mathbb{E}_{y \sim Y}[g(\theta_y)]$, for some function $g : [-1, 1] \rightarrow [0, 1]$. For example:

- $\text{SD}(C_0, C_1) = \mathbb{E}_{y \sim Y}[|\theta_y|]$ (i.e., $g(z) = |z|$, see Eq. (2.4));
- $\text{TD}(C_0, C_1) = \mathbb{E}_{y \sim Y}[\theta_y^2]$ (i.e., $g(z) = z^2$, see Eq. (2.13)); and
- $\text{JS}(C_0, C_1) = \mathbb{E}_{y \sim Y}\left[1 - h\left(\frac{1 + \theta_y}{2}\right)\right]$ (i.e., $g(z) = 1 - h\left(\frac{1 + z}{2}\right)$, see Eq. (2.12)).

To reduce TDP to JSP, we took a convex combination of the two circuits and used the fact that $1 - h\left(\frac{1 + \theta_y}{2}\right) \approx O(\theta_y^2)$ for small values of θ_y . While this worked for polarization of TD (which corresponds to $g(z) = z^2$), it seems unlikely to yield a polarization lemma for SD for an arbitrarily small (but noticeable) gap. The reason is that the function $g(z) = |z|$ — the g -function corresponding to SD — is not differentiable at 0 and in particular does not act like z^2 for small

²² This definition of convex combination is more convenient to analyze than perhaps the more natural definition of $D_b = \lambda \cdot C_b + (1 - \lambda) \cdot C_{1-b}$.

values of z . As we find this similarity between the different notions of distance striking, and indeed our proofs leverage the relations between them, we provide in Fig. 2.1 a plot comparing the different choices for the function g .

Another popular f -divergence that we have not discussed thusfar²³ is the *squared Hellinger distance*, defined as $H^2(P, Q) = \frac{1}{2} \sum_y (\sqrt{P_y} - \sqrt{Q_y})^2$. It can be shown that $H^2(C_0, C_1) = \mathbb{E}_{y \sim Y} [1 - \sqrt{1 - \theta_y^2}]$, and so also this distance falls within the above framework (i.e., by considering $g(z) = 1 - \sqrt{1 - z^2}$).

Notably, the squared Hellinger distance also acts like JS (and TD) around 0; namely, $1 - \sqrt{1 - \theta_y^2} \approx O(\theta_y^2)$ for small values of θ_y . However, unlike $\text{TDP}^{\alpha, \beta}$, we do not know how to show that the HELLINGER DIFFERENCE PROBLEM, denoted $\text{HDP}^{\alpha, \beta}$ and defined analogously to $\text{TDP}^{\alpha, \beta}$ (while replacing the distance TD with H^2), is in SZK for all $(1/\text{poly})$ -separated (α, β) . We do mention that $H^2(P, Q) \leq \text{TD}(P, Q) \leq 2H^2(P, Q)$, and thus $\text{HDP}^{\alpha, \beta}$ is in SZK if α and $\beta/2$ are $(1/\text{poly})$ -separated. However, the proof described above does not go through if we try to apply it to the Hellinger distance—we cannot guarantee that the gap in the Hellinger distance after taking the convex combination is larger than the error in the Taylor approximation. Indeed, the question whether $\text{HDP}^{\alpha, \beta}$ is in SZK for any $(1/\text{poly})$ -separated (α, β) , first raised by Fehr and Vaudenay [FV17], remains open.

References

- AARV17. Benny Applebaum, Barak Arkis, Pavel Raykov, and Prashant Nalini Vasudevan. Conditional disclosure of secrets: Amplification, closure, amortization, lower-bounds, and separations. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 727–757. Springer, 2017.
- AGGM06. Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on np-hardness. In Jon M. Kleinberg, editor, *Symposium on Theory of Computing*, pages 701–710. ACM, 2006.
- AH91. William Aiello and Johan Hastad. Statistical Zero-knowledge Languages can be recognized in two rounds. *Journal of Computer and System Sciences*, 42(3):327–345, 1991.
- BB15. Andrej Bogdanov and Christina Brzuska. On basing size-verifiable one-way functions on np-hardness. In *TCC*, pages 1–6, 2015.
- BBF16. Zvika Brakerski, Christina Brzuska, and Nils Fleischhacker. On statistically secure obfuscation with approximate correctness. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II*, pages 551–578, 2016.

²³ Actually we will use the squared Hellinger distance to analyze triangular discrimination of direct product distributions (see the full version for details). Also, the squared Hellinger distance is closely related to the Fidelity distance: $\text{Fidelity}(P, Q) = 1 - H^2(P, Q)$.




Fig. 2.1: Comparison between the difference choices of the function g that were discussed. Since all functions are symmetric around 0, we restrict to the domain $[0, 1]$. Recall that $g_1(\theta) = |\theta|$ corresponds to SD, $g_2(\theta) = \theta^2$ to TD, $g_3(\theta) = 1 - h(\frac{1+\theta}{2})$ to JS and $g_4(\theta) = 1 - \sqrt{1 - \theta^2}$ to H^2 .

- BBM11. Nayantara Bhatnagar, Andrej Bogdanov, and Elchanan Mossel. The computational complexity of estimating MCMC convergence time. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 14th International Workshop, APPROX 2011, and 15th International Workshop, RANDOM 2011, Princeton, NJ, USA, August 17-19, 2011. Proceedings*, pages 424–435, 2011.
- BCH⁺17. Adam Bouland, Lijie Chen, Dhiraj Holden, Justin Thaler, and Prashant Nalini Vasudevan. On the power of statistical zero knowledge. In *FOCS*, 2017.
- BDRV18. Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Multi-collision resistant hash functions and their applications. In *EUROCRYPT*, 2018.
- BDV17. Nir Bitansky, Akshay Degwekar, and Vinod Vaikuntanathan. Structure vs hardness through the obfuscation lens. *CRYPTO*, 2017.
- BG03. Michael Ben-Or and Danny Gutfreund. Trading help for interaction in statistical zero-knowledge proofs. *J. Cryptology*, 16(2):95–116, 2003.
- BHZ87. Ravi B. Boppana, Johan Håstad, and Stathis Zachos. Does co-np have short interactive proofs? *Inf. Process. Lett.*, 25(2):127–132, 1987.
- BKP18. Nir Bitansky, Yael Tauman Kalai, and Omer Paneth. Multi-collision resistance: a paradigm for keyless hash functions. In *STOC*, 2018.
- BL13. Andrej Bogdanov and Chin Ho Lee. Limits of provable security for homomorphic encryption. In *Advances in Cryptology - CRYPTO 2013 - 33rd*

- Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, pages 111–128, 2013.
- Cam86. Lucien Le Cam. *Asymptotic Methods in Statistical Decision Theory*. Springer-Verlag, New York, NY, 1986.
- CCKV08. André Chailloux, Dragos Florin Ciocan, Iordanis Kerenidis, and Salil P. Vadhan. Interactive and noninteractive zero knowledge are equivalent in the help model. In *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008.*, pages 501–534, 2008.
- CGVZ18. Yi-Hsiu Chen, Mika Göös, Salil P. Vadhan, and Jiapeng Zhang. A tight lower bound for entropy flattening. In *CCC*, 2018.
- DNR04. Cynthia Dwork, Moni Naor, and Omer Reingold. Immunizing encryption schemes from decryption errors. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT*, volume 3027 of *Lecture Notes in Computer Science*, pages 342–360. Springer, 2004.
- 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.
- For89. Lance Fortnow. The complexity of perfect zero-knowledge. *Advances in Computing Research*, 5:327–343, 1989.
- FV17. Serge Fehr and Serge Vaudenay. Personal Communication, 2017.
- Gol90. Oded Goldreich. A note on computational indistinguishability. *Inf. Process. Lett.*, 34(6):277–281, 1990.
- Gol17. Oded Goldreich. *Introduction to Property Testing*. Cambridge University Press, 2017.
- GS89. Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. *Advances in Computing Research*, 5:73–90, 1989.
- GSV98. Oded Goldreich, Amit Sahai, and Salil Vadhan. Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In *STOC*, 1998.
- GSV99. Oded Goldreich, Amit Sahai, and Salil P. Vadhan. Can statistical zero knowledge be made non-interactive? or on the relationship of SZK and NISZK. In Michael J. Wiener, editor, *CRYPTO*, volume 1666 of *Lecture Notes in Computer Science*, pages 467–484. Springer, 1999.
- GV99. Oded Goldreich and Salil P. Vadhan. Comparing entropies in statistical zero knowledge with applications to the structure of SZK. In *CCC*, 1999.
- GV11. Oded Goldreich and Salil P. Vadhan. On the complexity of computational problems regarding distributions. In Oded Goldreich, editor, *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation - In Collaboration with Lidor Avigad, Mihir Bellare, Zvika Brakerski, Shafi Goldwasser, Shai Halevi, Tali Kaufman, Leonid Levin, Noam Nisan, Dana Ron, Madhu Sudan, Luca Trevisan, Salil Vadhan, Avi Wigderson, David Zuckerman*, volume 6650 of *Lecture Notes in Computer Science*, pages 390–405. Springer, 2011.
- GVW02. Oded Goldreich, Salil Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *Computational Complexity*, 11(1-2):1–53, 2002.
- HR05. Thomas Holenstein and Renato Renner. One-way secret-key agreement and applications to circuit polarization and immunization of public-key encryption. In *CRYPTO*, pages 478–493, 2005.
- IL89. Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography. In *STOC*, pages 230–235, 1989.

- KNY17. Ilan Komargodski, Moni Naor, and Eylon Yogev. White-box vs. black-box complexity of search problems: Ramsey and graph property testing. In *FOCS*, 2017.
- KNY18. Ilan Komargodski, Moni Naor, and Eylon Yogev. Collision resistant hashing for paranoids: Dealing with multiple collisions. In *EUROCRYPT*, pages 162–194, 2018.
- KY18. Ilan Komargodski and Eylon Yogev. On distributional collision resistant hashing. In *CRYPTO*, 2018.
- LZ17. Shachar Lovett and Jiapeng Zhang. On the impossibility of entropy reversal, and its application to zero-knowledge proofs. In *TCC*, 2017.
- NR06. Moni Naor and Guy N. Rothblum. Learning to impersonate. In *ICML*, pages 649–656, 2006.
- Ost91. Rafail Ostrovsky. One-way functions, hard on average problems, and statistical zero-knowledge proofs. In *Structure in Complexity Theory Conference*, pages 133–138, 1991.
- OV08. Shien Jin Ong and Salil P. Vadhan. An equivalence between zero knowledge and commitments. In *TCC*, pages 482–500, 2008.
- OW93. Rafail Ostrovsky and Avi Wigderson. One-way functions are essential for non-trivial zero-knowledge. In *ISTCS*, pages 3–17, 1993.
- PW17. Yury Polyanskiy and Yihong Wu. Lecture notes on information theory. Available at: http://people.lids.mit.edu/yp/homepage/data/itlectures_v5.pdf, 2017.
- SV03. Amit Sahai and Salil Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM (JACM)*, 50(2):196–249, 2003.
- Top00. Flemming Topsøe. Some inequalities for information divergence and related measures of discrimination. *IEEE Transactions on Information Theory*, 46(4):1602–1609, July 2000.
- Vad99. Salil Pravin Vadhan. *A study of statistical zero-knowledge proofs*. PhD thesis, Massachusetts Institute of Technology, 1999.
- Yeh16. Amir Yehudayoff. Pointer chasing via triangular discrimination. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:151, 2016.