

Verification of Statistical Properties: Redefining the Possible

Clément Canonne* Sam Polgar[†] Aditya Vikram Singh[‡]
 Sri Aravindan Thyagarajan[§] Joy Qiping Yang[¶]

December 4, 2025

Abstract

We revisit the setting of Interactive Proof Systems for Distribution Testing, introduced by Chiesa and Gur (2018), showing that a simple twist on the task requirements may lead to dramatic improvements, allowing verifiers with *constant* sample complexity. We define and investigate the multi-prover and zero-knowledge versions of these interactive proof systems, using as flagship example the task of farness verification – the “dual” version of closeness testing. We hope that our results will inspire others to investigate and analyze the power and limitations of multiple provers for distribution verification.

1 Introduction

Analyzing (very) large datasets to build accurate models is the workhorse of machine learning and underlies most of the advances in machine learning over the past decade. These datasets are increasingly seen as valuable assets, *e.g.*, due to the difficulty in obtaining them (sensitive, regulated, or carefully curated user data), generating them (computation-heavy processes), or trusting them (poisoning attacks). For companies owning such datasets, this leads to a thorny issue: how to convince interested customers that a dataset has the application-specific statistical properties they need, while revealing as little data as possible?

One avenue forward, inspired by the success of interactive proofs in the functional setting,¹ is that of *proofs of proximity for probability distributions*, proposed by Chiesa and Gur [CG18]. In this setting, which brings together interactive proofs and (property) distribution testing [Bat+00a; GGR98], one sees the “valuable dataset” as a probability distribution, and has an interested customer (Arthur) who seeks to leverage interaction with an all-powerful yet untrusted prover (Merlin) to verify whether this distribution $\mathbf{p}$ has a property of interest (or is far from it). This direction of research has recently seen a surge of interest [HR22; HR23; HR24; HR25; GR25], with results showing how to perform a range of fundamental testing tasks – for instance, testing any

*The University of Sydney. Email: clement.canonne@sydney.edu.au.

[†]The University of Sydney. Email: sam.polgar@sydney.edu.au.

[‡]The University of Sydney. Email: adityavikram.singh@sydney.edu.au.

[§]The University of Sydney. Email: aravind.thyagarajan@sydney.edu.au.

[¶]The University of Sydney. Email: qiping.yang@sydney.edu.au.

¹By functional setting, we mean the “standard” setting where one seeks to assess whether a fully specified input, encoded as a string, belongs to a given language.

label-invariant property – *much* more efficiently, data-wise, than without access to a prover. Yet, in spite of these impressive advances, two key limitations remain: first, these protocols do not “only” convince the verifier Arthur that the desired property holds; they actually leak a *lot* of information about Merlin’s dataset to him, *enough to learn a the whole frequency profile of the data*. Second, perhaps even more problematic, they all require Arthur to *have a significant number of data points on his own* – i.e., to himself hold a dataset of size at least $\Omega(\sqrt{k})$, where k is the domain size, *from the same distribution as Merlin*. This can be seen as deeply unrealistic in many scenarios: How can each potential customer have access to this (sublinear, but still considerable) amount of data from the very dataset they are trying to buy? If they obtain it from some secure, preliminary commitment on the prover’s side, is Merlin disclosing large chunks of his data to buyers, one after the other?

Thus, a natural hope is to remove this strong and unrealistic assumption on Arthur, and seek protocols which only require very little – ideally a *constant* number – of data point on the verifier’s side. (Indeed, as we elaborate later, once this number is low enough, one can indeed justify and envision such a “secure, preliminary commitment on the prover’s side” to provide Arthur with these very few trusted samples.) Unfortunately, Chiesa and Gur [CG18] quashed this hope early on, showing by a very simple and elegant simulation argument that this $\Omega(\sqrt{k})$ sample complexity requirement on Arthur’s side was necessary even to interactively verify² a property as basic as uniformity. This early information-theoretic lower bound seemingly justifies the assumption made in follow-up works, which then focused on getting more bang for this $O(\sqrt{k})$ sample complexity buck. Yet, this strong and often unrealistic limitation, even justified by an information-theoretic lower bound, remains a strong and often unrealistic limitation: our work seeks to remove it.

Is there *any* interesting distribution verification task Arthur can perform with $O(1)$ samples?

Perhaps unsurprisingly, given the above framing, we argue that the answer is not only *yes*, but also that there are *many* such tasks. Our starting point is the following observation: while uniformity testing is a fundamental distribution testing task, and a building block for many other results [BKR04; Can+18; Gol20], *it is not in itself a “goal” in the distribution verification setting*. That is, it is unlikely that Arthur, a party eager to obtain relevant and hard-to-obtain data for their machine learning or data analysis purposes, hopes to be convinced the dataset on sale is actually “uniform noise,” or even from a fixed distribution he could easily sample from by himself – and that the seller Merlin’s pitch is that he will provide this easily samplable data to anyone willing to buy it.

Instead, one can much more easily motivate situations where Merlin wants to convince Arthur his dataset is *interesting* in some way: some clusters or high-probability records Arthur would be interested in learning more about, or some statistical artifacts worth paying a premium to investigate. In other words, Merlin wants to prove to Arthur the probability distribution he holds is very much *not* uniform. Now, while in the (non-interactive) distribution testing world these questions are equivalent, in the interactive proof setting swapping completeness and soundness might make all the difference: as we show (Theorem 1.1), there is a very simple protocol allowing to verify *non-uniformity* with a *constant* number of samples! And this is not only the case for (non-)uniformity: as we demonstrate, there exist other natural “dual” statistical tasks which allow for constant-sample verification protocols (yet for which non-interactive testing remains hard).

²Throughout, we will use “verify” instead of “test” to emphasize the difference with the non-interactive distribution testing setting, and to allow for generalizations of the tasks beyond the usual property testing formulation.

Moreover, as we show, this verification can be made without revealing too much to Arthur – in a sense which, following the considerable literature on Zero-Knowledge proofs, we make formal in the distribution setting.

1.1 Our contributions

“Lower bound arguments do not tell researchers the impossibility results exist. Researchers already know that impossibility results exist. Lower bound arguments tell researchers how the impossibility results can be bypassed.”
G.K. Chesterton (mildly paraphrased)

The main contribution of this paper is conceptual: namely, revisiting interactive proof systems for distribution testing (Definition 1.3), as introduced in [CG18], to enable a wider range of promise problems – and show that many of these, which remain “hard” even in the Arthur-Merlin setting, have a natural dual formulation which is not only practically relevant and motivated, *but also becomes easy*.

Theorem 1.1 (Verifying fairness from uniformity (Informal version of Theorem 4.1)). *There exists a simple one-round, communication-efficient protocol with the following guarantees: given a probability distribution $\mathbf{p}$ known to Merlin and provided via sample access to Arthur,*

- (Completeness) *If $\mathbf{p}$ is far from uniform, then there exists a strategy for Merlin such that Arthur accepts with probability at least $2/3$;*
- (Soundness) *If $\mathbf{p}$ is close to uniform, then no matter what strategy Merlin uses, Arthur accepts with probability at most $1/3$.*

Moreover, Arthur only takes a constant number of samples from $\mathbf{p}$.

To put this in perspective, recall that this task has sample complexity $\Omega\left(\frac{k}{\log k}\right)$ in the distribution testing setting (no Merlin) [VV11]; and $\Omega(\sqrt{k})$ samples on Arthur’s side in the interactive setting *when swapping completeness and soundness* (via a rather involved protocol [HR22]).

As a direct corollary of this positive result, we obtain a negative one (Corollary 4.6): namely, that an efficient reduction between two tasks, *Entropy Difference* and *Statistical Difference*, which underlies one of the completeness results for the class SZK (Statistical Zero-Knowledge), is impossible in the blackbox-sampling model.

Defining MIP distribution verification. Our next contribution is definitional: inspired by the Multi-prover Interactive Proof (MIP) setting [Ben+88], and motivated by well-studied distribution testing questions such as closeness testing [Bat+10; Cha+13; DK16], independence [Bat+01; ADK15], and testing properties of collections of distributions [LRR13; LRR12; DK16], we define the multi-prover version of IP distribution verification, whereby each of ℓ distinct non-interacting provers has full knowledge of (only) their own probability distribution $\mathbf{p}_i$, and Arthur, by interacting with these ℓ untrusted “Merlins”, seeks to verify some property of the joint input $(\mathbf{p}_1, \dots, \mathbf{p}_\ell)$. This MIP distribution verification, which is formalized in Definition 1.4, captures scenarios where distinct (possibly competing) data holders each hold their own dataset, and seek to convince a potential buyer Arthur that it is of interest for him to do so: this justifies the assumption that each “Merlin” only has knowledge of *one* distribution $\mathbf{p}_i$, and not the full tuple – in contrast to the functional MIP setting of [Ben+88].

Farness verification. With this definition in hand, we turn to a second verification task: the dual of closeness testing, which we term *farness verification*. Here, two provers – Merlin and Nimue – each have full knowledge of only their own probability distribution, $\mathbf{p}$ and $\mathbf{q}$, and seek to convince Arthur that $\mathbf{p}$ and $\mathbf{q}$ are statistical *far* from each other. One could see this arising from a scenario where Arthur is willing to get diverse, complementary data from different sources, to have a more representative training set: and so two datasets coming from the same distribution are of little appeal, while two datasets statistically far from each other will meet his goal. It is known that, in the non-interactive testing setting, $\Omega(k^{2/3})$ samples are required from at least one of the two distributions [Bat+10; Val11; Cha+13; DK16] for this farness testing task (which is then equivalent to closeness testing); and that by the same simulation argument as for uniformity, testing closeness in the interactive proof setting necessitates $\Omega(\sqrt{k})$ samples from Arthur [CG18].

Theorem 1.2 (Verifying farness (Informal version of Theorem 5.1)). *There exists a simple one-round, communication-efficient protocol with the following guarantees: given two probability distributions $\mathbf{p}$ and $\mathbf{q}$, with $\mathbf{p}$ known to Merlin and $\mathbf{q}$ to Nimue, respectively, and provided via sample access to Arthur,*

- (Completeness) *If $\mathbf{p}$ is far from $\mathbf{q}$, then there exists a strategy for Merlin and Nimue such that Arthur accepts with probability at least $2/3$;*
- (Soundness) *If $\mathbf{p} = \mathbf{q}$, then no matter what strategy Merlin and Nimue use, Arthur accepts with probability at most $1/2$.*

Moreover, Arthur only takes a constant number of samples from $\mathbf{p}$ and $\mathbf{q}$.

Along the way, we also establish some variants of the above theorem, allowing for stronger soundness guarantees under more stringent assumptions on $\mathbf{p}, \mathbf{q}$.

Zero-Knowledge. Having (re)defined IP distribution verification, and introduced its multi-prover generalization, it is natural to ask how to address this first “key limitation” alluded to earlier: that is, the fact that all existing IP distribution testing protocols in the literature leak a considerable amount of information from Merlin to Arthur. Our next contribution is to define the natural analogue of *Zero-Knowledge proofs* to the (single- and multi-prover) distribution verification setting (Definition 1.6), both statistical and perfect. We also show that these definitions are *achievable*: specifically, we show in Theorem 4.1 that the non-uniformity verification protocol of Theorem 1.1 satisfies our Statistical Zero-Knowledge definition, and, later on (Theorem 5.9), that a variant of our farness verification protocol (under stronger assumptions) achieves Perfect Zero-Knowledge.

In view of our motivation for the multi-prover setting, we also introduce a new, weaker notion of Zero-Knowledge proof, which is easier to attain yet does capture an important desirable property: that is, “any given prover should not learn much about the *other* provers’ data.” This definition, which we formalize in Definition 1.7, can be interpreted as follows: given an MIP distribution verification protocol Π with ℓ provers, for every $i \in [\ell]$, there exists an algorithm which simulates the view of prover M_i from Π in the completeness case (*i.e.*, where everyone is honest), yet only is given input $\mathbf{p}_i$ and the same sample access to $\mathbf{p}_1, \dots, \mathbf{p}_{i-1}, \mathbf{p}_{i+1}, \dots, \mathbf{p}_\ell$ as Arthur. In other words, this definition ensures that whatever M_i learns about $\mathbf{p}_1, \dots, \mathbf{p}_{i-1}, \mathbf{p}_{i+1}, \dots, \mathbf{p}_\ell$ when running a (valid) instance of the protocol Π on a “yes-input”, it could have learned in an “isolated execution” involving only himself and Arthur, pretending to run the protocol Π with *none of the other provers*. We show that in Theorem 5.1 that the non-uniformity verification protocol of Theorem 1.2 achieves this “One-Merlin-Out” Statistical Zero-Knowledge definition.

Repetition theorems. The completeness and soundness parameters set above to $2/3$ and $1/3$, respectively, may seem arbitrary: as “standard” in randomized computation and the interactive proof literature, one expects to be able to amplify them *via* repetition. In Section 6, we show that this is indeed the case, establishing the analogues of repetition theorems from the (function) IP and ZK literature for the distribution verification setting. In particular, the completeness parameter can be made exponentially close to 1 and the soundness parameter exponentially close to 0, thanks to guarantees provided by sequential repetition theorem (Theorem 6.1). Whether the same can be achieved using *parallel* repetition is something we leave as a future direction.

1.2 Definitions: (Multi)-Prover distribution verification, and capturing Zero-Knowledge

We begin by recalling (and suitably extending) the definition of Chiesa and Gur:

Definition 1.3 (IP distribution verification (Slight generalization of [CG18, Definition 5.2])). Given two disjoint properties $\mathcal{P}_0, \mathcal{P}_1 \subseteq \Delta(\mathcal{X})$ over a known domain $\mathcal{X}$ of size $k := |\mathcal{X}|$,³ an *interactive proof system for distribution verification* Π of properties $\mathcal{P}_0, \mathcal{P}_1 \subseteq \Delta(\mathcal{X})$ consists of an algorithm A which interactively exchanges messages with a prover M such that at the end of the interaction the following two conditions hold.

- (Completeness) for all $\mathbf{p} \in \mathcal{P}_0$, there exists a prover strategy M such that

$$\Pr[\Pi(M(\mathbf{p}), A^{\mathbf{p}}) = \text{accept}] \geq 2/3.$$

- (Soundness) for all $\mathbf{p} \in \mathcal{P}_1$, for all prover strategy M^* ,

$$\Pr[\Pi(M^*(\mathbf{p}), A^{\mathbf{p}}) = \text{accept}] \leq 1/3,$$

where $A^{\mathbf{p}}$ indicates that A has sample access to $\mathbf{p}$. The *sample complexity* of Π is the (worst-case) number of samples A draws from the distribution, the *communication complexity* of Π is the (worst-case) total number of bits exchanged between the parties, and the *round complexity* of Π is the (worst case) number of rounds of interaction, where each round consists of a message from one party to the other and its reply. We say that Π is *efficient* when the algorithm A is running in time polynomial in k and its sample complexity.

The above choice of $2/3$ for the probability of success of the proof system may seem arbitrary, and – as is the case in the standard non-interactive setting – easy to amplify by standard arguments. However, amplification by repetition theorems are much less trivial in the interactive setting (see, e.g., [Gol98, Appendix C.1] and [Gol25]). In Section 6, we expand on this, checking that the argument for parallel and sequential repetition in the “functional” (single-prover) IP setting do carry over to the (single-prover) distribution verification IP one.

“Functional” multi-prover interactive protocols (MIP) [Ben+88]. In the functional setting, an ℓ -prover proof system for a language L is a protocol Π between ℓ (computationally unbounded) provers $(M_1, \dots, M_\ell)$ and (probabilistic polynomial-time) verifier A such that (1) $M_1, \dots, M_\ell, A$ are all given the input x , but for every $i \neq j$, M_i does not see the messages between M_j and A ; (2) $M_1, \dots, M_\ell$ all have access to a read-only shared uniformly random string $r \in \{0, 1\}^*$, independent of the input x ; and, (3) Π satisfies the following correctness conditions:

³As common in the property testing literature, we slightly abuse notation and refer to a property $\mathcal{P}$ corresponding to a fixed size parameter k , instead of an (infinite) family $\mathcal{P} = \{\mathcal{P}_k\}_{k \in \mathbb{N}}$.

- (Completeness) for all $x \in L$, $\Pr[\Pi(M_1, \dots, M_\ell, A)(x) = \text{accept}] \geq 2/3$;
- (Soundness) for all $x \notin L$, for all $P'_1, \dots, P'_\ell$, $\Pr[\Pi(P'_1, \dots, P'_\ell, A)(x) = \text{accept}] \leq 1/3$;

where the probability is taken over r and the provers and verifier's own private randomness. Importantly, in contrast to the single-prover case, while sequential repetition theorems are known to increase the probability of success to be arbitrarily close to 1, success probability amplification cannot be done in parallel in the blackbox way (see, e.g., [FRS94, Section 6]).

Motivating the multi-prover model. In our case, for closeness testing:

1. The input is a pair of distributions $(\mathbf{p}, \mathbf{q})$.
2. There are two powerful provers M and N , and a sample-bounded verifier A .
3. M has access to $\mathbf{p}$ (but not $\mathbf{q}$), N has access to $\mathbf{q}$ (but not $\mathbf{p}$), and A has (sample-)access to $(\mathbf{p}, \mathbf{q})$.
4. Completeness: If $d_{TV}(\mathbf{p}, \mathbf{q}) \geq 9/10$, then $\Pr[\Pi(M, N, A) = \text{accept}] \geq 2/3$.
5. Soundness: If $d_{TV}(\mathbf{p}, \mathbf{q}) \leq 1/10$, then for all M', N' $\Pr[\Pi(M', N', A) = \text{accept}] \leq 1/3$.

We will in the remainder of this paper focus on either the single-prover setting (between Arthur and Merlin, abbreviated as A and M) or the 2-prover setting (involving Arthur, Merlin, and Nimue, abbreviated as A , M , and N). However, for the sake of generality, and as we believe this extension to be of interest, we define the ℓ -prover case, for arbitrary $\ell \geq 2$.

Definition 1.4 (MIP distribution verification). Given two disjoint properties $\mathcal{P}_0, \mathcal{P}_1 \subseteq \Delta(\mathcal{X}_1 \times \dots \times \mathcal{X}_\ell)$ over known domains $\mathcal{X}_1, \dots, \mathcal{X}_\ell$ of size $k_1, \dots, k_\ell$, a *multi-prover interactive proof system for distribution verification* Π of properties $\mathcal{P}_0, \mathcal{P}_1$ consists of an algorithm A which interactively exchanges messages with provers $M_1, \dots, M_\ell$ (which share a read-only uniformly random string $r \in \{0, 1\}^*$, independent of everything else, but do not communicate with each other) such that at the end of the interaction the following two conditions hold.

- (Completeness) for all $(\mathbf{p}_1, \dots, \mathbf{p}_\ell) \in \mathcal{P}_0$, there exists a prover strategy $(M_1, \dots, M_\ell)$ such that
$$\Pr[\Pi(M_1(\mathbf{p}_1), \dots, M_\ell(\mathbf{p}_\ell), A^{\mathbf{p}_1, \dots, \mathbf{p}_\ell}) = \text{accept}] \geq 2/3.$$
- (Soundness) for all $(\mathbf{p}_1, \dots, \mathbf{p}_\ell) \in \mathcal{P}_1$, for every prover strategy $(M_1^*, \dots, M_\ell^*)$,
$$\Pr[\Pi(M_1^*(\mathbf{p}_1), \dots, M_\ell^*(\mathbf{p}_\ell), A^{\mathbf{p}_1, \dots, \mathbf{p}_\ell}) = \text{accept}] \leq 1/3,$$

where $A^{\mathbf{p}_1, \dots, \mathbf{p}_\ell}$ indicates that A has (independent) sample access to $\mathbf{p}_1, \dots, \mathbf{p}_\ell$. The *sample complexity* of Π is the (worst-case) total number of samples A draws from the distributions, the *communication complexity* of Π is the (worst-case) total number of bits exchanged between the parties, and the *round complexity* of Π is the (worst case) number of rounds of interaction, where each round consists of a message from one party to the other and its reply. We say that Π is *efficient* when the algorithm A is running in time polynomial in $k_1, \dots, k_\ell$ and its sample complexity.

Remark 1.5. In contrast with the “functional” MIP, here $M_1, M_2, \dots, M_\ell$ only have *partial* access to the input $(\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_\ell)$. On the one hand, one can hope this weaker access to the input allows us to obtain soundness guarantees with a *constant* number of samples; on the other, it makes proving completeness more challenging, as even the “all-powerful provers” only have limited information. This also makes for a richer landscape when it comes to defining notions of zero knowledge since we can impose zero knowledge on the provers $M_1, \dots, M_\ell$ too: as we will now see.




Figure 1: The MIP distribution verification setting. The ℓ provers (“Merlins”) share a read-only random tape, independent of their respective inputs, but are not allowed to communicate otherwise. The i -th Merlin is given full knowledge of their part of the input, namely, distribution $\mathbf{p}_i$, and does not know or have access to $\mathbf{p}_j$ for $j \neq i$.

Defining Zero Knowledge for MIP distribution verifiers. We now introduce two different notions of Zero-Knowledge for MIP distribution verification, starting with the direct analogue of the (Statistical) Zero-Knowledge in the functional setting.

Definition 1.6 (ZK for MIP distribution verification). An MIP proof system $\Pi = (M_1, \dots, M_\ell, A)$ for distribution verification properties $\mathcal{P}_0, \mathcal{P}_1$ is said to be (*honest-verifier*) *statistical zero-knowledge* if there is a randomized algorithm S and a non-increasing function $\delta: \mathbb{N} \rightarrow [0, 1]$ such that, for every $(\mathbf{p}_1, \dots, \mathbf{p}_\ell) \in \mathcal{P}_0$ and all $t \in \mathbb{N}$,

$$d_{\text{TV}}\left(S^{\mathbf{p}_1, \dots, \mathbf{p}_\ell}(t), \left\langle A^{\mathbf{p}_1, \dots, \mathbf{p}_\ell}, M_1(\mathbf{p}_1), \dots, M_\ell(\mathbf{p}_\ell) \right\rangle_A(t)\right) \leq \delta(t),$$

where S has the same sample complexity as Π (up to constant factors), and $\left\langle A, M_1, \dots, M_\ell \right\rangle$ denotes the *view* of the interaction (*i.e.*, the transcript) between all parties, from Arthur’s perspective (including its own randomness and samples from $\mathbf{p}_1, \dots, \mathbf{p}_\ell$). When $\delta = 0$ (constant function), we say that Π is (*honest-verifier*) *perfect zero-knowledge*.

As in the functional case, we emphasize that the simulator is only required to work in the completeness case (when the input belongs to $\mathcal{P}_0$). However, there is an important difference with the functional case: namely, that we do not enforce computational efficiency, but instead require that the simulator only be given *sample access* to the input distributions (*i.e.*, the same access as Arthur) *and* uses a comparable number of samples.⁴ This sample complexity requirement is key, as otherwise the simulator would be allowed to draw a number of samples sufficient to, *e.g.*, learn the input distributions.

Having defined ZK, which captures that Arthur should not learn much from the interaction, we turn to the weaker, new notion we discussed earlier: a notion of Zero-Knowledge which captures that any given *prover* should not learn much about the *other* provers’ inputs.

⁴Note that we allow some constant-factor slack here, to allow sufficient robustness in the definition.

Definition 1.7 (HIPOZK for MIP distribution verification). An MIP proof system $\Pi = (M_1, \dots, M_\ell, A)$ for distribution verification properties $\mathcal{P}_0, \mathcal{P}_1$ is said to be *Statistical Honest Interactive Prover Only (HIPO) zero-knowledge* if there is a randomized algorithm S and a non-increasing function $\delta: \mathbb{N} \rightarrow [0, 1]$ such that, for every $(\mathbf{p}_1, \dots, \mathbf{p}_\ell) \in \mathcal{P}_0$, every $i \in [\ell]$, and all $t \in \mathbb{N}$,

$$d_{\text{TV}}\left(S^{\mathbf{p}_1, \dots, \mathbf{p}_\ell}(\mathbf{p}_i, t), \left\langle A^{\mathbf{p}_1, \dots, \mathbf{p}_\ell}, M_1(\mathbf{p}_1), \dots, M_\ell(\mathbf{p}_\ell) \right\rangle_{M_i}(t)\right) \leq \delta(t),$$

where S has the same sample complexity as Π and $\left\langle A^{\mathbf{p}_1, \dots, \mathbf{p}_\ell}, (M_1(\mathbf{p}_1), \dots, M_\ell(\mathbf{p}_\ell)) \right\rangle_{M_i}(t)$ denotes view of the interaction between all parties, from the i -th prover perspective. When $\delta = 0$, we say that Π is *perfect* HIPOZK.

Remark 1.8 (Generalizing HIPOZK to arbitrary subset of Merlins). Instead of simulating the interaction between all parties from the i -th prover's perspective as in Definition 1.7, we can also define zero knowledge for a fixed subset of provers, i.e., among that subset of provers, they do not learn more than the number of samples they could see from the protocol. This would capture the setting where a few provers (companies) can collude to learn secrets of other provers (companies).

2 Related Work

Property testing, introduced in [RS96; GGR98], is the study of ultra-efficient randomized algorithms for approximate decision-making on *unknown* massive objects. These algorithms, called testers, are given *local query access* to an oracle representing the object, e.g., as a function f . With only a number of queries sublinear in the size n of the object's representation, the tester must decide whether f is in the property $\mathcal{P}$ of interest, or ε -far (i.e., at distance at least ε) from every $g \in \mathcal{P}$, where the distance depends on the type of objects considered (e.g., Hamming distance for functions). See [Fis01; Fis04; Ron08; Gol10; Gol17] for introductions.

In the *distribution testing* setting [GGR98; Bat+00a], the “objects” are probability distributions over a discrete domain of size k , the type of queries is (typically) i.i.d. sample access to the distribution, and the measure of distance is (again, typically) the total variation distance. Uniformity [GR99; Pan08], identity [Bat+01; ADK15], and closeness [Bat+00b; Ach+12; Cha+13; Bat+13; DK16] testing were some of the earliest and fundamental tasks considered, and all require a sampling complexity polynomial in the domain size ($\Theta(\sqrt{k}/\varepsilon^2)$ for uniformity and identity, and $\Theta(\max(k^{2/3}/\varepsilon^{4/3}, \sqrt{k}/\varepsilon^2))$ for closeness, respectively).

A more challenging task in property testing is that of *tolerant* testing [PRR06], where one seeks to distinguish objects ε_1 -close to the property from those ε_2 -far from it. In the distribution testing setting, this is a much more challenging task, with tolerant uniformity testing requiring $\Theta(k/\log k)$ samples, for instance [VV11; Can+22]. Tolerant testing has been thoroughly investigated, especially for the broad class of *symmetric (label-invariant) properties* (that is, properties closed under relabeling of the domain), such as entropy and support size [Bat+05; GMV09; VV11; VV17; Ras+09; WY16].

The above are concerned with the non-interactive setting, in which the testing algorithm has access to the “object” – but no external help. Rothblum, Vadhan, and Wigderson introduce *Interactive Proofs of Proximity (IPP)* [RVW13] in the functional setting following work from [EKR04], where an unbounded and possibly adaptively cheating prover receives x , interacts with a sublinear-time verifier to decide whether $x \in L$ or is ε -far from L . They answer the question of whether a proof

system can improve complexity bounds for property testing by showing exponential improvements over the standard property testing model for specific properties.

[BRV18], motivated by the question of whether a prover convince a verifier of correct computation if the verifier only reads a few bits of input, without revealing additional non-local information about the input, defines the Zero-Knowledge Proof of Proximity (ZKPP) setting, where an unbounded and potentially malicious prover seeks to convince a sublinear-time verifier that the unknown object is ϵ -close to a property, without revealing *non-local information*. Unlike classical ZK [GMR89; GMW91] where a verifier sees the input x and learns whether $x \in L$ (without revealing anything else), ZKPP restricts the verifier to oracle access to the object (e.g., function f or graph G) and reveals no information beyond what the queries.

Combining the above two then-disjoint lines of inquiry, Chiesa and Gur [CG18] introduced the notion of proof systems for properties of distributions, *Proofs of Proximity for Distribution Testing*. The question there was to understand whether an untrusted prover with complete knowledge (or, alternatively, unbounded sample access) to a distribution can prove membership in a property to a verifier who only has sample access – using significantly fewer samples than required for testing the property. Chiesa and Gur initiated a systematic treatment of the question, introducing variants of the models (MA, AM, and IP distribution testers), and investigated the power and advantages conferred by public- and private-coin protocols, and interaction between prover and verifier. A line of follow-up works by Herman and Rothblum [HR22; HR23; HR24; HR25] improved efficiency first for symmetric properties, then for general properties; however, their results all require at least $\Omega(\sqrt{k})$ samples on the verifier’s side. [GR25] and [HR25] specifically investigate whether (and for which properties) both verifier and prover could have sample complexity significantly smaller than their “corresponding task” – namely, smaller than the testing sample complexity for the verifier, and smaller than the learning complexity for the prover.

Finally, the (functional) setting of *multiple* provers was introduced in [Ben+88], and investigated (among others) in [FRS94]. [CRR11; CRR13; GKR15; CLW25] then specifically consider this multiprover setting in the context of delegated computation.

We finally mention two other relevant works. First, that of Goldwasser, Rothblum, Shafer, and Yehudayoff [Gol+21], which introduces a setting analogous to that of Chiesa and Gur [CG18], but for verifying *PAC learning* instead of testing. Second, that of Tell [Tel16], which, in the (noninteractive, testing-only) setting, studies the *dual problems* of property testing in depth: namely, requiring that the tester accept when the object is ϵ -far from a property $\mathcal{P}$ and reject if it is ϵ -far from the set of objects that are ϵ -far from $\mathcal{P}$. As a special case of [Tel16, Theorem 1.9], it is shown that the dual of testing uniformity (testing fairness from uniformity) is the same as testing uniformity (when support size is sufficiently larger than $1/\epsilon$). Moreover, in general, when testing without provers, dual problems are at least as hard to test as the original problems [Tel16, Observation 1.4]).

2.1 Future Directions: the Holy Grail (Arthur without samples)

Our results leave open how far Arthur’s *sample* complexity can be pushed while still certifying interesting distributional properties. We highlight two complementary relaxations.

Hidden modes with game-theoretic incentives. Recent work shows that some tasks deemed impossible or unfair in the standard cryptographic model can be recovered once the adversary is associated with utility functions, and it wants to maximise its utility [Cle86; And+14; BK14; BB14;

BDD20; Chu+18; WAS22; TSW24]. Motivated by this, we could couple *hidden modes* with incentives: Arthur sometimes has *no* samples, but crafts challenges whose distribution is indistinguishable from those in rounds where he does have samples, so Merlin cannot tell the mode. Under a utility that rewards Arthur’s acceptance and imposes a loss or 0 utility for Arthur’s rejection, Merlin maximizes utility by responding exactly as if Arthur had samples in every round. This preserves soundness and lets a small fraction of with-sample rounds lift completeness, thereby lowering Arthur’s *expected* sample complexity. The main tasks are to (i) define and enforce mode-indistinguishability at the transcript level, (ii) calibrate the with-sample fraction and payoff parameters to meet a target δ -soundness and completeness, and (iii) prove robustness to collusion/adaptivity without leaking the mode.

Cryptographic assistance. If we move away from the information-theoretic setting to a computational one where we allow computational or setup assumptions, we can let Merlin (or Merlins) provide succinct, verifiable certificates (*e.g.*, commitments plus short arguments) that shift cost from Arthur’s sample complexity to communication. This points to protocols where Arthur, for instance, uses only a constant or amortized subconstant number of fresh samples to spot-audit a proof of “non-uniformity” computed over the full dataset. One has to characterize tradeoffs between Arthur’s samples, argument length and verification time, and whether reusable certificates can push per-instance sampling below current constants.

Along the way, we also state a few open problems of interest, starting with a very natural one.

Open Problem 1. Which other statistical tasks admit a dual version which can be verified with only a constant number of samples?

3 Preliminaries

We will use throughout the standard $O(\cdot)$ notation, and write $\Delta(k)$ to denote the set of discrete distributions supported on a domain $\mathcal{X}$ of size k (without loss of generality, assume that the domain is $[k] = \{1, 2, \dots, k\}$). For two discrete distributions $\mathbf{p} \in \Delta(k)$ and $\mathbf{q} \in \Delta(k)$, the total variation distance is defined as $d_{\text{TV}}(\mathbf{p}, \mathbf{q}) = \sup_{S \subseteq \mathcal{X}} \mathbf{p}(S) - \mathbf{q}(S) = \frac{1}{2} \sum_{x \in [k]} |\mathbf{p}(x) - \mathbf{q}(x)|$.

We now recall two results we will rely upon, the Polarization Lemma and the Pearson–Neyman Lemma. While originally written for distributions sampleable from (explicitly given) *circuits*, one can check by inspection of the proof that the former also applies to distributions provided *via* samples only. For completeness, we restate a proof sketch of the Polarization Lemma in Appendix A.

Lemma 3.1 (Polarization Lemma [Vad99, Lemma 3.1.12]). *Let $\alpha, \beta \in [0, 1]$ be any two constants such that $\alpha^2 > \beta$. There exists a polynomial-time computable mapping which, given sample access to two discrete probability distributions $\mathbf{p}, \mathbf{q}$ (over a domain of size k) and parameter $\delta \in (0, 1]$, provides sample access to two discrete probability distributions $\mathbf{p}', \mathbf{q}'$ (over a domain of size k') such that*

- If $d_{\text{TV}}(\mathbf{p}, \mathbf{q}) \geq \alpha$, then $d_{\text{TV}}(\mathbf{p}', \mathbf{q}') \geq 1 - \delta$;
- If $d_{\text{TV}}(\mathbf{p}, \mathbf{q}) \leq \beta$, then $d_{\text{TV}}(\mathbf{p}', \mathbf{q}') \leq \delta$;

and $k' = k^s$, where $s = (\log(1/\delta)) \frac{2 \log(1/\alpha)}{\log(\alpha^2/\beta)} + O(\log(1/\alpha))$. Moreover, one can efficiently sample from $\mathbf{p}'$ or $\mathbf{q}'$ given s samples from $\mathbf{p}$ and $\mathbf{q}$.

Lemma 3.2 (Pearson-Neyman Lemma (Folklore; see, e.g., [Can22, Lemma 1.4])). *Any (possibly randomized) statistical test which distinguishes between $\mathbf{p}$ and $\mathbf{q}$ from a single sample must have Type I (false positive) and Type-II (false negative) errors satisfying*

$$\text{Type I} + \text{Type II} \geq 1 - d_{\text{TV}}(\mathbf{p}, \mathbf{q}).$$

4 Verifying non-uniformity

Theorem 4.1. *There exists a one-round private-coin protocol $\Pi = \Pi(A, M)$ with the following guarantees: on input $k \geq 1$, and constants $\delta \in (0, 1]$, $\varepsilon_1, \varepsilon_2 \in [0, 1]$ with $\varepsilon_1 < \varepsilon_2^2$; and sample access to an unknown probability distribution $\mathbf{p} \in \Delta(k)$,*

- (Completeness) *If $d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k) \geq \varepsilon_2$, then $\Pr[\Pi(A, M) = \text{accept}] \geq 1 - \delta$;*
- (Soundness) *If $d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k) \leq \varepsilon_1$, then, for all M^* , $\Pr[\Pi(A, M^*) = \text{accept}] \leq \delta$.*

Moreover, A takes $(\log(1/\delta))^{O(1/\log(\varepsilon_1^2/\varepsilon_2))} = O(1)$ samples, and the protocol is (honest-verifier) Statistical Zero-Knowledge.

Algorithm 1 Basic protocol for verifying non-uniformity

Require: Sample access to $\mathbf{p}$ (Arthur) over domain $\mathcal{X}$, full knowledge of $\mathbf{p}$ (Merlin)

- ▷ **Arthur**
 - 1: Chooses $b \leftarrow_R \{0, 1\}$,
 - 2: Sample $x_0 \sim \mathbf{u}_k$ and $x_1 \sim \mathbf{p}$
 - 3: Send x_b to Merlin
 - ▷ **Merlin**
 - 4: Upon receiving x_b , must guess $\hat{b} \in \{0, 1\}$
 - 5: Send $\hat{b}$ to Arthur
 - ▷ **Arthur**
 - 6: **return** accept if $\hat{b} = b$.
-

Lemma 4.2. *There exists a one-round private-coin protocol (A, M) , Algorithm 1, with the following guarantees: on input $k \geq 1$, $\varepsilon_1, \varepsilon_2 \in [0, 1]$ with $\varepsilon_1 < \varepsilon_2$ and sample access to an unknown probability distribution $\mathbf{p} \in \Delta(k)$,*

- (Completeness) *If $d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k) \geq \varepsilon_2$, then $\Pr[\Pi(A, M) = \text{accept}] \geq \frac{1}{2} + \frac{\varepsilon_2}{2}$;*
- (Soundness) *If $d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k) \leq \varepsilon_1$, then, for all M^* , $\Pr[\Pi(A, M^*) = \text{accept}] \leq \frac{1}{2} + \frac{\varepsilon_1}{2}$.*

Moreover, A takes one sample, and the total communication is $O(\log k)$ bits.

Proof. The lemma will directly follow from the next two claims, which, respectively, establish completeness and soundness of the protocol. Its sample and communication complexity are immediate.

Claim 4.3 (Completeness). *If $d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k) \geq \varepsilon_2$, then there exists a strategy M for Merlin such that the protocol of Algorithm 1 returns **accept** with probability at least $\frac{1+\varepsilon_2}{2}$.*

Proof. Consider the so-called Scheffé set

$$S_{\mathbf{p}} := \left\{ x \in \mathcal{X} : \mathbf{p}(x) \geq \frac{1}{k} \right\}$$

which Merlin can compute explicitly, having full knowledge of the distribution $\mathbf{p}$. Merlin's strategy M is then simply to return the bit $\hat{b} := \mathbb{1}_{S_{\mathbf{p}}}(x_b)$, i.e., the indicator of whether the sample sent by Arthur falls in this Scheffé set.

By standard facts about total variation distance, $\mathbf{p}(S_{\mathbf{p}}) - \mathbf{u}_k(S_{\mathbf{p}}) = d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k) \geq \varepsilon_2$; as a result,

$$\begin{aligned} \Pr[\hat{b} = b] &= \Pr[x_0 \notin S_{\mathbf{p}} \mid b = 0] \Pr[b = 0] + \Pr[x_1 \in S_{\mathbf{p}} \mid b = 1] \Pr[b = 1] \\ &= \frac{1}{2} (\mathbf{u}_k(\overline{S_{\mathbf{p}}}) + \mathbf{p}(S_{\mathbf{p}})) \\ &= \frac{1}{2} (1 - \mathbf{u}_k(S_{\mathbf{p}}) + \mathbf{p}(S_{\mathbf{p}})) \\ &\geq \frac{1}{2} (1 + \varepsilon_2), \end{aligned}$$

as claimed. $\square$

Claim 4.4 (Soundness). *If $d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k) \leq \varepsilon_1$, then there is no strategy M for Merlin such that the protocol of Algorithm 1 returns **accept** with probability greater than $\frac{1+\varepsilon_1}{2}$.*

Proof. This follows from the characterization of total variation distance as best possible advantage for distinguishing two fixed distributions given a single observation. Namely, any possibly randomized strategy of Merlin to guess b can be seen as a distribution over deterministic strategies: each such strategy M^* corresponds to outputting the indicator of a subset $S_{M^*} \subseteq \mathcal{X}$: but the total variation distance is $d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k) = \sup_{S \subseteq \mathcal{X}} (\mathbf{p}(S) - \mathbf{u}_k(S)) \geq \mathbf{p}(S_{M^*}) - \mathbf{u}_k(S_{M^*})$; and so, as in the proof of the previous claim,

$$\Pr[\hat{b} = b] = \frac{1}{2} (1 - \mathbf{u}_k(S_{M^*}) + \mathbf{p}(S_{M^*})) \leq \frac{1}{2} (1 + \varepsilon_1),$$

proving the claim for deterministic strategies, and, by convexity, for randomized ones as well. $\square$

This concludes the proof of Lemma 4.2. $\square$

Given Lemma 4.2, it would be tempting to try and establish Theorem 4.1 by (parallel) repetition, invoking Theorem 6.1. While this would indeed lead to the desired sample and communication complexities (and only require $\varepsilon_1 < \varepsilon_2$ instead of the more stringent $\varepsilon_1 < \varepsilon_2^2$; though it would increase the number of rounds), it would unfortunately *not* yield the claimed SZK property: intuitively, the idea is that a simulator which does not know $d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k)$ *exactly* in the completeness case (but only has a lower bound ε_2 on it) can only simulate one instance of the protocol within total variation $\frac{1-\varepsilon_2}{2}$.⁵ After amplification over T repetitions, the total variation between distributions of simulated and true transcripts could then grow roughly as $\Omega(\min(T(1 - \varepsilon_2), 1))$.

⁵To see why, note that the simulator, without knowledge of $d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k)$, cannot “know” whether the probability to correctly guess b should be the lower bound $\frac{1}{2}(1 + \varepsilon_2)$, or arbitrarily close to 1.

Proof of Theorem 4.1. In order to obtain a meaningful Zero-Knowledge guarantee, we instead amplify the total variation distance in the completeness case (from ε_2 to $1 - \delta$) *before* running the basic protocol of Lemma 4.2. To do so, we let Arthur run the mapping of the Polarization Lemma (Lemma 3.1), so that with

$$(\log(1/\delta'))^{O(1/\log(\varepsilon_1^2/\varepsilon_2))}$$

samples from $\mathbf{p}$ and $\mathbf{u}_k$ (where δ' is a parameter to be determined later, and we rely on the assumption $\varepsilon_2^2 > \varepsilon_1$ in order to apply Lemma 3.1), Arthur and Merlin can run Algorithm 1 on a distribution $\mathbf{p}'$ over a new domain of size $k' = k'(k, \varepsilon_1, \varepsilon_2, \delta)$ such that (1) if $d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k) \geq \varepsilon_2$, then $d_{\text{TV}}(\mathbf{p}', \mathbf{u}_{k'}) \geq 1 - 2\delta'$, and (2) if $d_{\text{TV}}(\mathbf{p}, \mathbf{u}_k) \leq \varepsilon_1$, then $d_{\text{TV}}(\mathbf{p}', \mathbf{u}_{k'}) \leq 2\delta'$. Thereafter, we can apply $O(\log(1/\delta))$ rounds of sequential repetition to reduce the probability in the soundness case with Theorem 6.1. To show that it is SZK (with deviation parameter δ), we give its simulator:

Algorithm 2 Simulator for Algorithm 1 after sequential repetition

Require: Sample access to $\mathbf{p}$

- 1: $t \leftarrow O(\log(1/\delta))$
 - 2: **for all** $i \leftarrow 1$ to t **do**
 - ▷ Store b_i in memory to use at the end.
 - 3: Chooses $b_i \leftarrow_R \{0, 1\}$
 - 4: Sample $x_i \sim \mathbf{u}_k$ and $x'_i \sim \mathbf{p}$
 - 5: Set the view to be $(b_1, \dots, b_t, x_1, \dots, x_t, x'_1, \dots, x'_t, \hat{b}_1 = b_1, \dots, \hat{b}_t = b_t)$
-

The transcript's total variation distance from the true transcript is exactly the probability that $\{\hat{b}_1 \neq b_1\} \vee \dots \vee \{\hat{b}_t \neq b_t\}$ in the latter, which is (by completeness) at most $\delta' \cdot t = \delta' \log(1/\delta)$ by a union bound. Setting $\delta' = \delta / \log(1/\delta)$ suffices. This establishes Theorem 4.1. $\square$

4.1 An unexpected consequence of Theorem 4.1.

It was shown in [Vad99] (Theorem 3.4.8) that ENTROPY DIFFERENCE (ED) reduces to STATISTICAL DIFFERENCE (SD): the definition of the two computational problems assume that the input distributions are encoded as Boolean *circuits* (so that sampling from $\mathbf{p}$ amounts to evaluating the circuit $\mathcal{C}_{\mathbf{p}}$ on a string r of uniformly random input bits), and the reduction from ED to SD given in [Vad99, Theorem 3.4.8] relies on this assumption. To the best of our knowledge, whether such a reduction is possible even in the weaker *blackbox sampling* access model, was unknown.

A corollary of our results is that such a reduction given only sampling access to the distributions is impossible. To see why, recall that SD reduces to its complement: in more detail (see, e.g., [Vad99, Section 4.4]),

Theorem 4.5 (Reversing Statistical Difference [Vad99, Corollary 4.4.1]). *There is a polynomial-time computable function that maps pairs of distributions $(\mathbf{p}, \mathbf{q})$ (encoded as Boolean circuits) to pairs of distributions $(\mathbf{p}', \mathbf{q}')$ (encoded as Boolean circuits) such that*

- if $d_{\text{TV}}(\mathbf{p}, \mathbf{q}) \geq \frac{2}{3}$, then $d_{\text{TV}}(\mathbf{p}', \mathbf{q}') \leq \frac{1}{3}$;
- if $d_{\text{TV}}(\mathbf{p}, \mathbf{q}) \leq \frac{1}{3}$, then $d_{\text{TV}}(\mathbf{p}', \mathbf{q}') \geq \frac{2}{3}$.

The proof of this statement relies on combining three reductions: (1) reduction from SD to ED, (2) reduction from ED to its complement, and (3) reduction from ED to SD. Of these, inspection of the proof shows that (1) carries over to the blackbox sampling model, and so does (2) (which is immediate). Moreover, all three steps are *efficient*: for Boolean circuits with m input gates and n output gates (*i.e.*, inducing probability distributions over $\{0, 1\}^n$), (1) and (2) only require a *constant* number of circuit evaluations (or, equivalently, samples), while (3) only requires $\text{poly}(m, n)$ of them (a polylogarithmic dependence on the domain size $k = 2^n$).

From the above, it is apparent that *if* one could obtain a version of (3), the reduction from ED to SD, in the blackbox sampling access model, using t samples from $\mathbf{p}$ or $\mathbf{q}$, then the statistical difference reversal of Theorem 4.5 would be possible in the blackbox sampling access model as well, and require only $O(t)$ samples from $\mathbf{p}$ and $\mathbf{q}$ to sample from $\mathbf{p}'$ and $\mathbf{q}'$. But *then*, one could use this reduction from SD to its complement on $(\mathbf{p}, \mathbf{u}_k)$ to reduce *verifying uniformity* (a task known to require $\Omega(\sqrt{k})$ samples from Arthur) to *verifying non-uniformity* (which, by Theorem 4.1, can be done with $O(1)$ samples from Arthur)! This implies that any mapping reducing ED to SD in the blackbox sampling access model cannot be efficient, *i.e.*, must require $t = \Omega(\sqrt{k}) = \Omega(2^{n/2})$ samples from the input distributions:

Corollary 4.6. *Any reduction from ENTROPY DIFFERENCE to STATISTICAL DIFFERENCE in the blackbox sampling access model (for probability distributions over $\{0, 1\}^n$) requires $\Omega(2^{n/2})$ samples.*

4.2 Some non-unexpected consequences of Theorem 4.1.

Leveraging the (now-standard) reduction from testing identity to an arbitrary fixed distribution $\mathbf{q}$ over a domain of size k to that of testing uniformity over a domain of size $O(k)$ (see [DK16; Gol20], or [Can22, Section 2.2.3]), we can generalize to *non-identity*, where the goal is to verify that the unknown distribution $\mathbf{p}$ is far from a fixed, known reference distribution $\mathbf{p}^*$.

Lemma 4.7 (Verifying non-identity). *There exists a one-round private-coin protocol $\Pi = \Pi(A, M)$ with the following guarantees: on input $k \geq 1$, constants $\varepsilon, \delta \in (0, 1]$, full knowledge by all parties of a known reference distribution $\mathbf{p}^*$, and sample access to an unknown probability distribution $\mathbf{p} \in \Delta(k)$,*

- (Completeness) *If $d_{\text{TV}}(\mathbf{p}, \mathbf{p}^*) \geq \varepsilon$, then $\Pr[\Pi(A, M) = \text{accept}] \geq 1 - \delta$;*
- (Soundness) *If $d_{\text{TV}}(\mathbf{p}, \mathbf{p}^*) \leq \frac{\varepsilon^2}{2}$, then, for all M^* , $\Pr[\Pi(A, M^*) = \text{accept}] \leq \delta$.*

Moreover, A takes $(\log(1/\delta))^{O(1)} = O(1)$ samples, and the protocol is (honest-verifier) Statistical Zero-Knowledge.

The proof is immediate, relying on the fact that the mapping underlying the aforementioned reduction preserves total variation distances up to a constant factor. (This constant factor is the reason for the slightly weaker statement, which links the distance parameters in the completeness and soundness cases instead of allowing arbitrary $\varepsilon_1 < \varepsilon_2$.)

Another nearly-immediate consequence of Theorem 4.1 is for *non-independence verification*, whereby Merlin seeks to convince Arthur that the two variables (X, Y) of a joint distribution $\mathbf{p}_{XY}$ are *not* independent – an important task, *e.g.*, if one intends to fit a model to learn the statistical relation between two features of a dataset.

Lemma 4.8 (Verifying non-independence). *There exists a one-round private-coin protocol (A, M) , with the following guarantees: on input $k_X, k_Y \geq 1$, $\varepsilon_1, \varepsilon_2 \in [0, 1]$, $\varepsilon_1 < \varepsilon_2/3$ and sample access to an unknown probability distribution $\mathbf{p}_{XY} \in \Delta(k_X \cdot k_Y)$,*

- (Completeness) *If $d_{\text{TV}}(\mathbf{p}_{XY}, \mathbf{q}_X \otimes \mathbf{q}_Y) \geq \varepsilon_2$ for every $\mathbf{q}_X \in \Delta(k_X), \mathbf{q}_Y \in \Delta(k_Y)$, then*

$$\Pr[\Pi(A, M) = \text{accept}] \geq \frac{1}{2} + \frac{\varepsilon_2}{2};$$

- (Soundness) *If $d_{\text{TV}}(\mathbf{p}_{XY}, \mathbf{q}_X \otimes \mathbf{q}_Y) \leq \varepsilon_1$ for some $\mathbf{q}_X \in \Delta(k_X)$ and $\mathbf{q}_Y \in \Delta(k_Y)$, then, for all M^* ,*

$$\Pr[\Pi(A, M^*) = \text{accept}] \leq \frac{1}{2} + \frac{3\varepsilon_1}{2}.$$

Moreover, A takes two samples, and the total communication is $O(\log(k_X) + \log(k_Y))$ bits.

Proof. The proof follows from the fact that we can simulate $\mathbf{p}_X \otimes \mathbf{p}_Y$ given sample access to $\mathbf{p}_{XY}$: take two independent samples $(x_1, y_1), (x_2, y_2) \sim \mathbf{p}_{XY}$ and return (x_1, y_2) . This problem then reduces to Lemma 4.2 by setting $\varepsilon_2 = 0$.

For the completeness case, i.e., when $\mathbf{p}_{XY}$ is far from independent, we note that

$$d_{\text{TV}}(\mathbf{p}_{XY}, \mathbf{p}_X \otimes \mathbf{p}_Y) \geq d_{\text{TV}}(\mathbf{p}_{XY}, \mathbf{q}_X \otimes \mathbf{q}_Y) \geq \varepsilon_2.$$

For the soundness case, by [Bat+01, Proposition 1], we have that

$$d_{\text{TV}}(\mathbf{p}_{XY}, \mathbf{q}_X \otimes \mathbf{q}_Y) \leq \varepsilon_1 \Rightarrow d_{\text{TV}}(\mathbf{p}_{XY}, \mathbf{p}_X \otimes \mathbf{p}_Y) \leq 3\varepsilon_1$$

Therefore, checking non-identity for $\mathbf{p}_{XY}$ v.s. $\mathbf{p}_X \otimes \mathbf{p}_Y$ suffice. $\square$

5 Verifying fairness

In this section, we describe and analyze our two-prover protocols for verifying *fairness* of distributions: Merlin holds a distribution $\mathbf{p}$, Nimue a distribution $\mathbf{q}$, and they seek to convince Arthur that $d_{\text{TV}}(\mathbf{p}, \mathbf{q})$ is large. We begin by a very simple (but very much non-Zero-Knowledge) protocol in Section 5.1, before building up to more satisfying (but similarly efficient) protocols in Section 5.2: one for general distributions in Section 5.2.1, then one with slightly better communication for the specific case of “flat” distributions in Section 5.2.2, and finally a perfect-ZK one (but for a very stringent parameter regime) in Section 5.3.

As a baseline, we recall that *testing fairness* (without a prover) is as hard as testing closeness, and thus requires $\Omega(\sqrt{k}/\varepsilon^2 + k^{2/3}/\varepsilon^{4/3})$ even to distinguish $\mathbf{p} = \mathbf{q}$ from $d_{\text{TV}}(\mathbf{p}, \mathbf{q}) \geq \varepsilon$. In the IP for distribution testing world, verifying *closeness* also requires $\Omega(\sqrt{k}/\varepsilon^2)$ samples for Arthur, by the same simulation argument as for uniformity from [CG18]. As we will see, verifying *fairness* is again drastically easier – with Arthur only required to have a *constant* (with respect to k) number of samples.

Algorithm 3 Naive protocol for verifying farness

Require: Sample access to $\mathbf{p}, \mathbf{q}$ (Arthur) over domain $\mathcal{X}$, full knowledge of $\mathbf{p}$ (Merlin), full knowledge of $\mathbf{q}$ (Nimue)

▷ **Arthur**

1: Chooses $b \leftarrow_R \{0, 1\}$,

2: Samples $x_0 \sim \mathbf{p}$ and $x_1 \sim \mathbf{q}$

3: Sends x_b to Merlin

▷ **Merlin**

4: Upon receiving x_b , sends τ (real number $\mathbf{p}(x_b)$) to Arthur

▷ **Arthur**

5: Upon receiving τ , sends x_b and τ to Nimue

▷ **Nimue**

6: Receives τ , computes $\hat{b} \leftarrow \mathbb{1}_{\{\tau < \mathbf{q}(x_b)\}}$.

7: **return accept** if $\hat{b} = b$.

5.1 Verifying farness is easy

Theorem 5.1. *There exists a two-round private-coin protocol (Algorithm 3) $\Pi = \Pi(A, M, N)$ with the following guarantees: on input $k \geq 1$, $\varepsilon_1, \varepsilon_2 \in [0, 1]$ with $\varepsilon_1 < \varepsilon_2$ and sample access to two unknown probability distributions $\mathbf{p}, \mathbf{q} \in \Delta(k)$ (where A has oracle access to both, M full knowledge of and access to $\mathbf{p}$ only, and N full knowledge of and access to $\mathbf{q}$ only),*

- (Completeness) *If $d_{TV}(\mathbf{p}, \mathbf{q}) \geq \varepsilon_2$, then $\Pr[\Pi(A, M, N) = \text{accept}] \geq \frac{1+\varepsilon_2}{2}$;*
- (Soundness) *If $d_{TV}(\mathbf{p}, \mathbf{q}) \leq \varepsilon_1$, then, for all M^*, N^* , $\Pr[\Pi(A, M^*, N^*) = \text{accept}] \leq \frac{1+\varepsilon_1}{2}$;*

Moreover, A takes $O(1)$ samples.

Proof. Since Nimue will have both "full" information (since Arthur will be relaying $\mathbf{p}(x_b)$) from both $\mathbf{p}$ and $\mathbf{q}$, this proof follows from single prover non-identity verification as in Lemma 4.7. $\square$

While the above result appears promising, it is unsatisfying in two ways: first, its communication complexity is unbounded, as sending the threshold τ (a real number) may require an arbitrarily large number of bits. (As we will see later in our "better" protocol (Algorithm 4), this issue can be relatively easily handled.) Second, and much more critical, it is very much *not* zero-knowledge, and indeed blatantly fails to satisfy even the weaker notion of SK we introduced, Definition 1.7 in an inherent way: Merlin learns something about Nimue's distribution he could not have learned by himself, or even with the help of (sample-bounded) Arthur only.

5.2 Verifying farness with Zero-Knowledge

5.2.1 General distributions

In this section, we establish the following result:

Theorem 5.2. *There exists a $O(1)$ -round private-coin protocol $\Pi = (A, M, N)$ with the following guarantees: on input $k \geq 1$, constants $\delta \in (0, 1]$, $\varepsilon_1, \varepsilon_2 \in [0, 1]$ with $\varepsilon_1 < \varepsilon_2$ and sample access to unknown probability distributions $\mathbf{p}, \mathbf{q} \in \Delta(k)$,*

- (Completeness) If $d_{TV}(\mathbf{p}, \mathbf{q}) \geq \varepsilon_2$, then $\Pr[\Pi(A, M, N) = \text{accept}] \geq 1 - \delta$;
- (Soundness) If $d_{TV}(\mathbf{p}, \mathbf{q}) \leq \varepsilon_1$, then, for all (M^*, N^*) , $\Pr[\Pi(A, M^*, N^*) = \text{accept}] \leq \delta$.

Moreover, A takes $O(1)$ sample from either $\mathbf{p}$ or $\mathbf{q}$, and the total communication is $O(\log k)$ bits. In addition, this protocol achieves perfect HIPO Zero Knowledge.

The proof of this theorem will follow from the “basic” building block (Lemma 5.3) below, combined with the sequential repetition theorem (Theorem 6.1) (and the fact that $\delta, \varepsilon_1, \varepsilon_2$ are constant).

Algorithm 4 Basic protocol for verifying farness

Require: Sample access to $\mathbf{p}, \mathbf{q}$ (Arthur) over domain $\mathcal{X}$, full knowledge of $\mathbf{p}$ (Merlin), full knowledge of $\mathbf{q}$ (Nimue)

▷ **Arthur**

1: Chooses $b \leftarrow_R \{0, 1\}$,

2: Sample $x_0 \sim \mathbf{p}$ and $x_1 \sim \mathbf{q}$

3: Send x_b to both Merlin and Nimue

▷ $O(\log k)$ bits

▷ **Merlin**

4: Upon receiving x_b , sends τ (an approximation of $\mathbf{p}(x_b)$) to Arthur

▷ $O(\log \frac{\log k}{\varepsilon_2 - \varepsilon_1})$ bits

▷ **Nimue**

5: Upon receiving x_b , sends λ (an approximation of $\mathbf{q}(x_b)$) to Arthur

▷ $O(\log \frac{\log k}{\varepsilon_2 - \varepsilon_1})$ bits

▷ **Arthur**

6: Receives τ, λ , computes $\hat{b} \leftarrow \mathbb{1}_{\{\tau < \lambda\}}$.

7: **return** accept if $\hat{b} = b$.

Lemma 5.3. *There exists a one-round private-coin protocol $\Pi = (A, M, N)$, Algorithm 4, with the following guarantees: on input $k \geq 1, \varepsilon_1, \varepsilon_2 \in [0, 1]$ with $\varepsilon_1 < \varepsilon_2$ and sample access to unknown probability distributions $\mathbf{p}, \mathbf{q} \in \Delta(k)$,*

- (Completeness) If $d_{TV}(\mathbf{p}, \mathbf{q}) \geq \varepsilon_2$, then $\Pr[\Pi(A, M, N) = \text{accept}] \geq \frac{1}{2} + \frac{\varepsilon_1 + \varepsilon_2}{4}$;
- (Soundness) If $d_{TV}(\mathbf{p}, \mathbf{q}) \leq \varepsilon_1$, then, for all (M^*, N^*) , $\Pr[\Pi(A, M^*, N^*) = \text{accept}] \leq \frac{1}{2} + \frac{\varepsilon_1}{2}$.

Moreover, A takes one sample from either $\mathbf{p}$ or $\mathbf{q}$, and the total communication is $O(\log \frac{k}{\varepsilon_2 - \varepsilon_1})$ bits. In addition, this protocol achieves perfect HIPO Zero Knowledge.

Proof of Lemma 5.3. The lemma will directly follow from the next two claims, which, respectively, establish completeness and soundness of the protocol. Its sample complexity is immediate; the communication complexity will be addressed shortly.

Claim 5.4 (Completeness). *If $d_{TV}(\mathbf{p}, \mathbf{q}) \geq \varepsilon_2$, then there exists a strategy M, N for Merlin and Nimue such that the protocol of Algorithm 4 returns **accept** with probability at least $\frac{1 + \frac{1}{2}(\varepsilon_2 + \varepsilon_1)}{2}$.*

Proof. Let $\alpha := \frac{\varepsilon_2 - \varepsilon_1}{4\varepsilon_2}$, and denote by $\tilde{S}_{\mathbf{p}, \mathbf{q}}$ the “approximate Scheffé set” corresponding to $\mathbf{p}$ and $\mathbf{q}$, defined as

$$\tilde{S}_{\mathbf{p}, \mathbf{q}} := \left\{ x \in \mathcal{X} : \mathbf{p}(x) > \max \left((1 + \alpha\varepsilon_2)\mathbf{q}(x), \frac{\alpha\varepsilon_2}{k} \right) \right\}$$

and by $S_{\mathbf{p},\mathbf{q}}$ the actual Scheffé set (i.e., $S_{\mathbf{p},\mathbf{q}} = S_{\mathbf{p},\mathbf{q}}^0$). Then

$$\begin{aligned} \mathbf{p}(\tilde{S}_{\mathbf{p},\mathbf{q}}) - \mathbf{q}(\tilde{S}_{\mathbf{p},\mathbf{q}}) &= \mathbf{p}(S_{\mathbf{p},\mathbf{q}}) - \mathbf{q}(S_{\mathbf{p},\mathbf{q}}) - \left(\mathbf{p}(S_{\mathbf{p},\mathbf{q}} \setminus \tilde{S}_{\mathbf{p},\mathbf{q}}) - \mathbf{q}(S_{\mathbf{p},\mathbf{q}} \setminus \tilde{S}_{\mathbf{p},\mathbf{q}}) \right) \\ &\geq \varepsilon_2 - \alpha \varepsilon_2 \mathbf{q}(S_{\mathbf{p},\mathbf{q}} \setminus \tilde{S}_{\mathbf{p},\mathbf{q}}) - \frac{\alpha \varepsilon_2}{k} \cdot k \\ &\geq (1 - 2\alpha) \varepsilon_2 = \frac{1}{2}(\varepsilon_2 + \varepsilon_1). \end{aligned}$$

Considering the approximate Scheffé set instead of the “true” one allows Merlin and Nimue to be able to round the value τ, λ to send, as it suffices now to consider only $O(\log \frac{\log k}{\alpha \varepsilon_2}) = O(\log \frac{\log k}{\varepsilon_2 - \varepsilon_1})$ possible values.⁶ This enables the stated communication complexity. The probability of acceptance is then

$$\begin{aligned} \Pr[\hat{b} = b] &= \Pr[\tau < \lambda \mid b = 0] \Pr[b = 0] + \Pr[\tau \geq \lambda \mid b = 1] \Pr[b = 1] \\ &= \frac{1}{2} \left(\Pr[x_0 \in \tilde{S}_{\mathbf{p},\mathbf{q}}] + \Pr[x_1 \notin \tilde{S}_{\mathbf{p},\mathbf{q}}] \right) \\ &= \frac{1}{2} \left(\mathbf{p}(\tilde{S}_{\mathbf{p},\mathbf{q}}) + 1 - \mathbf{q}(\tilde{S}_{\mathbf{p},\mathbf{q}}) \right) \\ &\geq \frac{1}{2} \left(1 + \frac{1}{2}(\varepsilon_2 + \varepsilon_1) \right), \end{aligned}$$

as claimed. $\square$

Claim 5.5 (Soundness). *If $d_{\text{TV}}(\mathbf{p}, \mathbf{q}) \leq \varepsilon_1$, then there is no strategy M, N for Merlin and Nimue such that the protocol of Algorithm 4 returns **accept** with probability greater than $\frac{1+\varepsilon_1}{2}$.*

Proof. This again follows from the characterization of total variation distance as best possible advantage for distinguishing two fixed distributions given a single observation. Namely, one can consider a “friendly but forgetful” Arthur which does not remember b after Line 3, but gets the full knowledge of $\mathbf{p}, \mathbf{q}$ from Merlin and Nimue to help him figure it out: then, any strategy M, N yields a strategy for Arthur to guess b from a single sample from a distribution promised to be either $\mathbf{p}$ or $\mathbf{q}$. The best probability of success for this augmented Arthur is $\frac{1}{2}(1 + d_{\text{TV}}(\mathbf{p}, \mathbf{q}))$. $\square$

Claim 5.6 (Perfect HIPOZK). *There exists a simulator $S^{\mathbf{p},\mathbf{q}}(\mathbf{p}, t)$, when given sample access access to $\mathbf{p}, \mathbf{q}$ (sample complexity same as Π), and $\mathbf{p}_i$ ’s full description, can simulate M ’s view of Π in the completeness case*

$$d_{\text{TV}}\left(S^{\mathbf{p},\mathbf{q}}(\mathbf{p}, t), \left\langle A^{\mathbf{p},\mathbf{q}}, M(\mathbf{p}), N(\mathbf{q}) \right\rangle_M(t)\right) = 0,$$

Algorithm 5 Simulator for Algorithm 4

Require: Sample access to $\mathbf{p}, \mathbf{q}$ and $\mathbf{p}$ ’s full description.

- 1: Chooses $b \leftarrow_R \{0, 1\}$
 - 2: Sample $x_0 \sim \mathbf{p}$ and $x_1 \sim \mathbf{q}$
 - 3: Set the view to be $(x_b, \mathbf{p}(x_b))$
-

⁶The encoding will work by first dividing $[0, 1]$ into $O(\log k)$ buckets of a geometric series form: $\{[0, \frac{1}{k}], (\frac{1}{k}, \frac{2}{k}], \dots, (\frac{2^{\lfloor \log k \rfloor}}{k}, 1]\}$ – since we only need multiplicative approximation to encode $\tilde{S}_{\mathbf{p},\mathbf{q}}$. Then within each bucket, we will create uniform buckets of interval $\alpha \varepsilon_2$. Encoding the location of the final bucket takes $O(\log \frac{\log k}{\alpha \varepsilon_2})$ bits.

Proof. It is not hard to see that Algorithm 5 exactly simulates Merlin's view. $\square$

This concludes the proof of Lemma 5.3. $\square$

Note that the above protocol, while (perfect) HIPOZK, is *not* SZK. We leave as an open problem whether one can achieve SZK for fairness verification of general distributions:

Open Problem 2. Can one obtain similar guarantees as in Theorem 5.2 (fairness verification for general distributions), but with *Statistical ZK*?

As a step towards this, we provide in our next section a protocol for *flat* distributions which, while also not SZK, we conjecture can be more easily made so.

5.2.2 Flat Distributions

We here focus on the specific case of *flat* distributions, *i.e.*, distributions which are uniform on their support. While the resulting fairness verification protocol we provide under this promise, Lemma 5.7, will only provide HIPO SZK, it achieves better communication from the provers. Further, we believe – as discussed at the end of the previous subsection – that it can provide some valuable insight into how to achieve Statistical Zero Knowledge.

Algorithm 6 Protocol for verifying Flat Distributions Lemma 5.7

Require: Sample access to $\mathbf{p}, \mathbf{q}$ (Arthur) over domain $\mathcal{X}$, full knowledge of $\mathbf{p}$ (Merlin), full knowledge of $\mathbf{q}$ (Nimue)

▷ **Arthur**

- 1: Chooses $b_1 \leftarrow_R \{0, 1\}, b_2 \leftarrow_R \{0, 1\}$
- 2: **If** $b_1 = 0$ **Then** $x_1 \sim \mathbf{p}$; **Else** $x_1 \sim \mathbf{q}$.
- 3: **If** $b_2 = 0$ **Then** $x_2 \sim \mathbf{p}$; **Else** $x_2 \sim \mathbf{q}$.
- 4: Send x_1 to Merlin and x_2 to Nimue.

▷ **Merlin**

- 5: Upon receiving x_1 , sends $\hat{b}_1 \leftarrow \mathbb{1}_{\mathbf{p}(x_1) > 0}$ to Arthur

▷ **Nimue**

- 6: Upon receiving x_2 , sends $\hat{b}_2 \leftarrow \mathbb{1}_{\mathbf{q}(x_2) > 0}$ to Arthur

▷ **Arthur**

- 7: return **accept** if $\hat{b}_1 = b_1 \vee \hat{b}_2 = b_2$.
-

Lemma 5.7 (Verifying Flat Distributions). *There exists a one-round private-coin protocol $\Pi = \Pi(A, M, N)$ with the following guarantees: on input $k \geq 1, \varepsilon \in [0, 1)$ and sample access to two unknown probability distributions $\mathbf{p}, \mathbf{q} \in \Delta(k)$ promised to be uniform on their support (where A has oracle access to both, M full knowledge of and access to $\mathbf{p}$ only, and N full knowledge of and access to $\mathbf{q}$ only),*

- (Completeness) *If $d_{TV}(\mathbf{p}, \mathbf{q}) \geq \varepsilon$, then $\Pr[\Pi(A, M, N) = \text{accept}] \geq \frac{3+\varepsilon}{4}$;*
- (Soundness) *If $d_{TV}(\mathbf{p}, \mathbf{q}) = 0$, then, for all M^*, N^* , $\Pr[\Pi(A, M^*, N^*) = \text{accept}] \leq \frac{3}{4}$.*

Moreover, A takes 2 samples, and the communication complexity is $O(\log k)$ and via Algorithm 5 this protocol is perfect HIPO ZK.

Proof. A takes $b_1, b_2 \sim \text{Bern}(1/2)$ and samples $x_1 \sim \mathbf{p}$ ($x_2 \sim \mathbf{p}$) if $b_1 = 0$ ($b_2 = 0$ resp.) and $x_1 \sim \mathbf{q}$ ($x_2 \sim \mathbf{q}$) if $b_1 = 1$ ($b_2 = 1$ resp.). Then A sends x_1 to M and x_2 to N for them to guess the value of b_1 and b_2 respectively. Arthur will accept if Merlin or Nimue's response is correct.

(Completeness). We first look at the completeness case and give Π , i.e., what A , M and N will do when $d_{\text{TV}}(\mathbf{p}, \mathbf{q}) \geq 1 - \varepsilon$: M will return 0 if $\mathbf{p}(x_1) > 0$ and 1 otherwise; Likewise, N will return 0 if $\mathbf{q}(x_2) > 0$, and 1 otherwise. A will accept if one of the answer is correct return **accept**. Denote $r(M)$ as the response from M . Without loss of generality, assume that $|\text{supp}(\mathbf{p})| \leq |\text{supp}(\mathbf{q})|$: This means that $\mathbf{p}(x) \geq \mathbf{q}(x)$ for $x \in \text{supp}(\mathbf{p})$

$$d_{\text{TV}}(\mathbf{p}, \mathbf{q}) = \sup_S (\mathbf{p}(S) - \mathbf{q}(S)) = \sum_{x: \mathbf{p}(x) \geq \mathbf{q}(x)} \mathbf{p}(x) - \mathbf{q}(x) = \sum_{x \in \text{supp}(\mathbf{p})} \mathbf{p}(x) - \mathbf{q}(x).$$

Effectively, M knows the Sheffé set for “free” – without communication from N or A , and thus we have the following

$$\begin{aligned} \Pr[r(M) = b_1] &= \Pr[r(M) = 0 \wedge b_1 = 0] + \Pr[r(M) = 1 \wedge b_1 = 1] \\ &= \Pr[r(M) = 0 | b_1 = 0] \cdot \frac{1}{2} + \Pr[r(M) = 1 | b_1 = 1] \cdot \frac{1}{2} \\ &= \frac{1}{2} + \Pr_{x_1 \sim \mathbf{q}} [p(x_1) > 0] \cdot \frac{1}{2} \\ &= \frac{1}{2} + \frac{d_{\text{TV}}(\mathbf{p}, \mathbf{q})}{2} \end{aligned}$$

And

$$\begin{aligned} \Pr[r(N) = b_2] &= \Pr[r(N) = 0 \wedge b_2 = 0] + \Pr[r(N) = 1 \wedge b_2 = 1] \\ &= \Pr[r(N) = 0 | b_2 = 0] \cdot \frac{1}{2} + \frac{1}{2} \\ &= \Pr_{x_2 \sim \mathbf{p}} [q(x_2) > 0] \cdot \frac{1}{2} + \frac{1}{2} \geq \frac{1}{2}. \end{aligned}$$

Overall, the correct probability in the completeness case is:

$$\begin{aligned} \Pr[\Pi^*(A, M, N) = \text{accept}] &= 1 - \Pr[r(M) \neq b_1 \text{ and } r(N) \neq b_2] \\ &\geq 1 - \left(\frac{1}{2} - \frac{d_{\text{TV}}(\mathbf{p}, \mathbf{q})}{2} \right) \cdot \left(1 - \frac{1}{2} \right) \\ &= \frac{3}{4} + \frac{\varepsilon}{4}. \end{aligned}$$

(Soundness). When $d_{\text{TV}}(\mathbf{p}, \mathbf{q}) = 0$, then by Lemma 3.2: For any M^*, N^*

$$\Pr[r(M^*) = b] \leq \frac{1 + d_{\text{TV}}(\mathbf{p}, \mathbf{q})}{2} \leq \frac{1}{2} \text{ and } \Pr[r(N^*) = b] \leq \frac{1}{2}.$$

Therefore, for any strategy M^*, N^* employs, Arthur will accept if one of them guess b_1 (b_2 resp.) correctly, which happens with probability:

$$\Pr[\Pi(A, M^*, N^*) = \text{accept}] = 1 - \Pr[r(M) \neq b_1 \text{ and } r(N) \neq b_2] \leq 1 - \left(\frac{1}{2} \right)^2 = \frac{3}{4}.$$

This concludes the proof. $\square$

Remark 5.8. As mentioned in the previous section, while this protocol is not, as stated, SZK (but only (perfect) HIPO ZK), we conjecture it can be adapted to be made SZK, and see this as a stepping stone towards answering Open Problem 2.

5.3 Is Perfect Zero-Knowledge a dream?

To conclude this section, we provide here a protocol for fairness (in a very specific, yet non-trivial parameter regime) which does provide *perfect* ZK, showing that this requirement is indeed achievable.

Theorem 5.9. *There exists a one-round private-coin protocol $\Pi = \Pi(A, M, N)$ with the following guarantees: on input $k \geq 1$, $\varepsilon \in [0, 1)$ and sample access to two unknown probability distributions $\mathbf{p}, \mathbf{q} \in \Delta(k)$ (where A has oracle access to both, M full knowledge of and access to $\mathbf{p}$ only, and N full knowledge of and access to $\mathbf{q}$ only),*

- (Completeness) *If $d_{TV}(\mathbf{p}, \mathbf{q}) = 1$, then $\Pr[\Pi(A, M, N) = \text{accept}] = 1$;*
- (Soundness) *If $d_{TV}(\mathbf{p}, \mathbf{q}) \leq \varepsilon$, then, for all M^*, N^* , $\Pr[\Pi(A, M^*, N^*) = \text{accept}] \leq \frac{1}{2} + \frac{\varepsilon}{2}$.*

Moreover, A takes one sample, and the protocol is (honest-verifier) Perfect Zero-Knowledge.

Algorithm 7 Perfect Zero Knowledge protocol for verifying disjointness

Require: Sample access to $\mathbf{p}, \mathbf{q}$ (Arthur) over domain $\mathcal{X}$, full knowledge of $\mathbf{p}$ (Merlin), full knowledge of $\mathbf{q}$ (Nimue), distance parameter ε .

- ▷ **Arthur**
 - 1: Chooses $b \leftarrow_R \{0, 1\}$
 - 2: Sample $x_0 \sim \mathbf{p}$ and $x_1 \sim \mathbf{q}$
 - 3: Send x_b to both Merlin and Nimue ▷ $O(\log k)$ bits
 - ▷ **Merlin**
 - 4: Upon receiving x_b , sends $\hat{b}_1$ (a guess for b) to Arthur
 - ▷ **Nimue**
 - 5: Upon receiving x_b , sends $\hat{b}_2$ (a guess for b) to Arthur
 - ▷ **Arthur**
 - 6: **return accept** if $b = \hat{b}_1 = \hat{b}_2$.
-

Proof. We provide the completeness, soundness, and zero knowledge analysis below; the protocol is given in Algorithm 7.

Completeness. Merlin and Nimue's strategy is to send the indicator of whether the received sample belong to the support of their distribution: that is, $\hat{b}_1 = \mathbb{1}_{\text{supp}(\mathbf{p})}(x_b)$ and $\hat{b}_2 = \mathbb{1}_{\text{supp}(\mathbf{q})}(x_b)$. If $d_{TV}(\mathbf{p}, \mathbf{q}) = 1$, then $\text{supp}(\mathbf{p}) \cap \text{supp}(\mathbf{q}) = \emptyset$ and as a result $\Pr[b = \hat{b}_1 = \hat{b}_2] = 1$, as desired.

Soundness. Here again soundness directly follows from the hypothesis characterization of total variation distance.

Perfect Zero-Knowledge. There is an obvious simulator for Arthur's view of the interaction in the completeness case:

Algorithm 8 Simulator for Algorithm 7

Require: Sample access to $\mathbf{p}, \mathbf{q}$

- 1: Chooses $b \leftarrow_R \{0, 1\}$
 - 2: Sample $x_0 \sim \mathbf{p}$ and $x_1 \sim \mathbf{q}$
 - 3: Set the view to be $(b, x_0, x_1, x_b, \hat{b}_1 = b, \hat{b}_2 = b)$
-

The transcript (over the choice of b, x_0, x_1) has then exactly the same distribution as Arthur's view of a honest interaction with Merlin and Nimue in the completeness case. $\square$

6 Repetition theorems

Theorem 6.1 (Sequential repetition for MIP distribution verification). *Suppose an MIP proof system $\Pi(M_1(\mathbf{p}_1), \dots, M_\ell(\mathbf{p}_\ell), A^{\mathbf{p}_1, \dots, \mathbf{p}_\ell})$ for distribution verification of properties $\mathcal{P}_0, \mathcal{P}_1$ exists with completeness parameter $1/2 + \Delta_c$ and soundness parameter $1/2 + \Delta_s$, where $\Delta_c > \Delta_s$. Then, for any $t \in \mathbb{N}$, there exists an MIP proof system $\Pi'(M_1(\mathbf{p}_1), \dots, M_\ell(\mathbf{p}_\ell), A^{\mathbf{p}_1, \dots, \mathbf{p}_\ell})$ for verification of properties $\mathcal{P}_0, \mathcal{P}_1$ with completeness parameter $1 - 2^{-\Omega(t)}$ and soundness parameter $2^{-\Omega(t)}$, where Π' is obtained by sequentially repeating the original MIP protocol $O(t/(\Delta_c - \Delta_s)^2)$ times.*

Proof Sketch. The proof follows as for sequential repetition of IP protocols in the functional setting (see, in particular, Theorem 1.4 in [Vas21]). In our case, the sequential repetition with m repetitions works as follows: Run the protocol Π m times in sequence, and let ρ be the fraction of times the protocol Π accepts. Then the overall protocol Π' accepts if and only if $\rho \geq (1 + \Delta_c - \Delta_s)/2$. The required soundness and completeness guarantees then follow from the Chernoff bound as in [Vas21]. As a technical subtlety, we note that Theorem 1.4 in [Vas21] essentially relies on the observation that if the soundness error does not decay exponentially with sequential repetition as expected, then there exists a single round cheating prover that has soundness error more than what the original protocol Π claims (which is a contradiction). In the functional IP setting of [Vas21], this cheating prover can be realized in two ways: either by simulating interaction with verifier and conditioning on an event, or by hard-coding the desired simulated transcript in the prover (see proof of Theorem 1.2 and the remark thereafter in [Vas21]). In our MIP setting, however, a prover is oblivious of the inputs of the other provers, whereas the verifier has access to samples from distribution of every prover, which means that a cheating prover cannot simulate interaction with verifier. Thankfully, though, the hard-coding argument still goes through, which is why we can borrow arguments from the functional IP setting. $\square$

Open Problem 3 (Parallel repetition for MIP distribution verification). We conjecture that the thresholding protocol as in the sequential repetition above leads to an exponential decay in soundness parameter for parallel repetition in the single-prover following arguments in [Gol25].

References

- [Ach+12] Jayadev Acharya et al. “Competitive Classification and Closeness Testing”. In: *COLT 2012 - The 25th Annual Conference on Learning Theory, June 25-27, 2012, Edinburgh, Scotland*. Ed. by Shie Mannor, Nathan Srebro, and Robert C. Williamson. Vol. 23. JMLR

- Proceedings. JMLR.org, 2012, pp. 22.1–22.18. URL: <http://proceedings.mlr.press/v23/acharya12/acharya12.pdf> (cit. on p. 8).
- [ADK15] Jayadev Acharya, Constantinos Daskalakis, and Gautam Kamath. “Optimal Testing for Properties of Distributions”. In: *NIPS*. 2015, pp. 3591–3599 (cit. on pp. 3, 8).
- [And+14] Marcin Andrychowicz et al. “Secure Multiparty Computations on Bitcoin”. In: *2014 IEEE Symposium on Security and Privacy*. IEEE, 2014, pp. 443–458. DOI: [10.1109/SP.2014.35](https://doi.org/10.1109/SP.2014.35) (cit. on p. 9).
- [Bat+00a] Tugkan Batu et al. “Testing that distributions are close”. In: *Proceedings 41st Annual Symposium on Foundations of Computer Science*. IEEE. 2000, pp. 259–269 (cit. on pp. 1, 8).
- [Bat+00b] Tugkan Batu et al. “Testing that distributions are close”. In: *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, Redondo Beach, California, USA, November 12-14, 2000*. IEEE Computer Society, 2000, pp. 259–269. DOI: [10.1109/SFCS.2000.892113](https://doi.org/10.1109/SFCS.2000.892113). URL: <https://doi.org/10.1109/SFCS.2000.892113> (cit. on p. 8).
- [Bat+01] Tugkan Batu et al. “Testing Random Variables for Independence and Identity”. In: *FOCS*. IEEE Computer Society, 2001, pp. 442–451 (cit. on pp. 3, 8, 15).
- [Bat+05] Tugkan Batu et al. “The Complexity of Approximating the Entropy”. In: *SIAM J. Comput.* 35.1 (2005), pp. 132–150. DOI: [10.1137/S0097539702403645](https://doi.org/10.1137/S0097539702403645). URL: <https://doi.org/10.1137/S0097539702403645> (cit. on p. 8).
- [Bat+10] Tugkan Batu et al. “Testing Closeness of Discrete Distributions”. In: *CoRR abs/1009.5397* (2010) (cit. on pp. 3, 4).
- [Bat+13] Tugkan Batu et al. “Testing Closeness of Discrete Distributions”. In: *J. ACM* 60.1 (2013), 4:1–4:25. DOI: [10.1145/2432622.2432626](https://doi.org/10.1145/2432622.2432626). URL: <https://doi.org/10.1145/2432622.2432626> (cit. on p. 8).
- [BB14] Adam Back and Iddo Bentov. “Note on fair coin toss via Bitcoin”. In: *CoRR abs/1402.3698* (2014). arXiv: [1402.3698](https://arxiv.org/abs/1402.3698). URL: <http://arxiv.org/abs/1402.3698> (cit. on p. 9).
- [BDD20] Carsten Baum, Bernardo David, and Rafael Dowsley. “Insured MPC: Efficient Secure Computation with Financial Penalties”. In: *Financial Cryptography and Data Security (FC 2020)*. Ed. by Joseph Bonneau and Nadia Heninger. Vol. 12059. Lecture Notes in Computer Science. Springer, 2020, pp. 404–420. DOI: [10.1007/978-3-030-51280-4_22](https://doi.org/10.1007/978-3-030-51280-4_22) (cit. on p. 10).
- [Ben+88] Michael Ben-Or et al. “Multi-Prover Interactive Proofs: How to Remove Intractability Assumptions”. In: *STOC*. ACM, 1988, pp. 113–131 (cit. on pp. 3, 5, 9).
- [BK14] Iddo Bentov and Ranjit Kumaresan. “How to Use Bitcoin to Design Fair Protocols”. In: *Advances in Cryptology – CRYPTO 2014*. Ed. by Juan A. Garay and Rosario Gennaro. Vol. 8617. Lecture Notes in Computer Science. Springer, 2014, pp. 421–439. DOI: [10.1007/978-3-662-44381-1_24](https://doi.org/10.1007/978-3-662-44381-1_24) (cit. on p. 9).
- [BKR04] Tugkan Batu, Ravi Kumar, and Ronitt Rubinfeld. “Sublinear algorithms for testing monotone and unimodal distributions”. In: *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*. Ed. by László Babai. ACM, 2004, pp. 381–390. DOI: [10.1145/1007352.1007414](https://doi.org/10.1145/1007352.1007414). URL: <https://doi.org/10.1145/1007352.1007414> (cit. on p. 2).

- [BRV18] Itay Berman, Ron D. Rothblum, and Vinod Vaikuntanathan. “Zero-Knowledge Proofs of Proximity”. In: *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*. Ed. by Anna R. Karlin. Vol. 94. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, 19:1–19:20. DOI: [10.4230/LIPICS.ITCS.2018.19](https://doi.org/10.4230/LIPICS.ITCS.2018.19) (cit. on p. 9).
- [Can+18] Clément L. Canonne et al. “Testing Shape Restrictions of Discrete Distributions”. In: *Theory Comput. Syst.* 62.1 (2018), pp. 4–62 (cit. on p. 2).
- [Can+22] Clément L. Canonne et al. “The Price of Tolerance in Distribution Testing”. In: *COLT*. Vol. 178. Proceedings of Machine Learning Research. PMLR, 2022, pp. 573–624 (cit. on p. 8).
- [Can22] Clément L. Canonne. “Topics and Techniques in Distribution Testing: A Biased but Representative Sample”. In: *Found. Trends Commun. Inf. Theory* 19.6 (2022), pp. 1032–1198 (cit. on pp. 11, 14).
- [CG18] Alessandro Chiesa and Tom Gur. “Proofs of Proximity for Distribution Testing”. In: *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*. Ed. by Anna R. Karlin. Vol. 94. LIPIcs. 2018, 53:1–53:14. DOI: [10.4230/LIPICS.ITCS.2018.53](https://doi.org/10.4230/LIPICS.ITCS.2018.53). URL: <https://doi.org/10.4230/LIPICS.ITCS.2018.53> (cit. on pp. 1–5, 9, 15).
- [Cha+13] Siu On Chan et al. “Optimal Algorithms for Testing Closeness of Discrete Distributions”. In: *CoRR abs/1308.3946* (2013). arXiv: [1308.3946](https://arxiv.org/abs/1308.3946). URL: <http://arxiv.org/abs/1308.3946> (cit. on pp. 3, 4, 8).
- [Chu+18] Kai-Min Chung et al. “Game theoretic notions of fairness in multi-party coin toss”. In: *Theory of Cryptography Conference*. Springer. 2018, pp. 563–596 (cit. on p. 10).
- [Cle86] Richard Cleve. “Limits on the Security of Coin Flips when Half the Processors Are Faulty”. In: *Proceedings of the 18th Annual ACM Symposium on Theory of Computing (STOC '86)*. ACM, 1986, pp. 364–369. ISBN: 0-89791-193-8. DOI: [10.1145/12130.12168](https://doi.org/10.1145/12130.12168) (cit. on p. 9).
- [CLW25] Ran Canetti, Ephraim Linder, and Connor Wagaman. “Refereed Learning”. en. In: (2025). DOI: [10.48550/arXiv.2510.05440](https://arxiv.org/abs/2510.05440). URL: <http://arxiv.org/abs/2510.05440> (visited on 10/22/2025) (cit. on p. 9).
- [CRR11] Ran Canetti, Ben Riva, and Guy N. Rothblum. “Practical delegation of computation using multiple servers”. In: *Proceedings of the 18th ACM Conference on Computer and Communications Security, CCS 2011, Chicago, Illinois, USA, October 17-21, 2011*. Ed. by Yan Chen, George Danezis, and Vitaly Shmatikov. ACM, 2011, pp. 445–454. DOI: [10.1145/2046707.2046759](https://doi.org/10.1145/2046707.2046759) (cit. on p. 9).
- [CRR13] Ran Canetti, Ben Riva, and Guy N. Rothblum. “Refereed delegation of computation”. In: *Inf. Comput.* 226 (2013), pp. 16–36. DOI: [10.1016/J.IC.2013.03.003](https://doi.org/10.1016/J.IC.2013.03.003) (cit. on p. 9).
- [DK16] Ilias Diakonikolas and Daniel M. Kane. “A New Approach for Testing Properties of Discrete Distributions”. In: *FOCS*. IEEE Computer Society, 2016, pp. 685–694 (cit. on pp. 3, 4, 8, 14).

- [EKR04] Funda Ergün, Ravi Kumar, and Ronitt Rubinfeld. “Fast approximate probabilistically checkable proofs”. In: *Inf. Comput.* 189.2 (2004), pp. 135–159. DOI: [10.1016/J.IC.2003.09.005](https://doi.org/10.1016/J.IC.2003.09.005) (cit. on p. 8).
- [Fis01] Eldar Fischer. “The Art of Uninformed Decisions: A Primer to Property Testing”. In: *Bull. EATCS* 75 (2001), p. 97. DOI: [10.1142/9789812562494_0014](https://doi.org/10.1142/9789812562494_0014) (cit. on p. 8).
- [Fis04] Eldar Fischer. “The art of uninformed decisions: A primer to property testing”. In: *Current Trends in Theoretical Computer Science: The Challenge of the New Century Vol 1: Algorithms and Complexity Vol 2: Formal Models and Semantics*. World Scientific, 2004, pp. 229–263 (cit. on p. 8).
- [FRS94] Lance Fortnow, John Rompel, and Michael Sipser. “On the Power of Multi-Prover Interactive Protocols”. In: *Theor. Comput. Sci.* 134.2 (1994), pp. 545–557 (cit. on pp. 6, 9).
- [GGR98] Oded Goldreich, Shafi Goldwasser, and Dana Ron. “Property Testing and its Connection to Learning and Approximation”. In: *J. ACM* 45.4 (1998), pp. 653–750. DOI: [10.1145/285055.285060](https://doi.org/10.1145/285055.285060) (cit. on pp. 1, 8).
- [GKR15] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. “Delegating Computation: Interactive Proofs for Muggles”. In: *J. ACM* 62.4 (2015), 27:1–27:64. DOI: [10.1145/2699436](https://doi.org/10.1145/2699436) (cit. on p. 9).
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. “The Knowledge Complexity of Interactive Proof Systems”. In: *SIAM J. Comput.* 18.1 (1989), pp. 186–208. DOI: [10.1137/0218012](https://doi.org/10.1137/0218012) (cit. on p. 9).
- [GMV09] Sudipto Guha, Andrew McGregor, and Suresh Venkatasubramanian. “Sublinear estimation of entropy and information distances”. In: *ACM Trans. Algorithms* 5.4 (2009), 35:1–35:16. DOI: [10.1145/1597036.1597038](https://doi.org/10.1145/1597036.1597038). URL: <https://doi.org/10.1145/1597036.1597038> (cit. on p. 8).
- [GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. “Proofs that Yield Nothing But Their Validity for All Languages in NP Have Zero-Knowledge Proof Systems”. In: *J. ACM* 38.3 (1991), pp. 691–729. DOI: [10.1145/116825.116852](https://doi.org/10.1145/116825.116852) (cit. on p. 9).
- [Gol+21] Shafi Goldwasser et al. “Interactive Proofs for Verifying Machine Learning”. In: *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*. Ed. by James R. Lee. Vol. 185. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 41:1–41:19. DOI: [10.4230/LIPICS.ITCS.2021.41](https://doi.org/10.4230/LIPICS.ITCS.2021.41) (cit. on p. 9).
- [Gol10] Oded Goldreich, ed. *Property Testing - Current Research and Surveys*. Vol. 6390. Lecture Notes in Computer Science. Springer, 2010. ISBN: 978-3-642-16366-1. DOI: [10.1007/978-3-642-16367-8](https://doi.org/10.1007/978-3-642-16367-8). URL: <https://doi.org/10.1007/978-3-642-16367-8> (cit. on p. 8).
- [Gol17] Oded Goldreich. *Introduction to Property Testing*. Cambridge University Press, 2017. DOI: [10.1017/9781108135252](https://doi.org/10.1017/9781108135252). URL: <http://www.cambridge.org/us/catalogue/catalogue.asp?isbn=9781107194052> (cit. on p. 8).
- [Gol20] Oded Goldreich. “The Uniform Distribution Is Complete with Respect to Testing Identity to a Fixed Distribution”. In: *Computational Complexity and Property Testing*. Vol. 12050. Lecture Notes in Computer Science. Springer, 2020, pp. 152–172 (cit. on pp. 2, 14).

- [Gol25] Oded Goldreich. “On Parallel Repetition of Interactive Proof Systems”. In: *Computational Complexity and Local Algorithms*. Vol. 15700. Lecture Notes in Computer Science. Springer, 2025, pp. 119–126 (cit. on pp. 5, 22).
- [Gol98] Oded Goldreich. *Modern Cryptography, Probabilistic Proofs and Pseudorandomness*. Vol. 17. Algorithms and Combinatorics. Springer, 1998 (cit. on p. 5).
- [GR25] Oded Goldreich and Guy N. Rothblum. “Location-Invariant Properties of Functions versus Properties of Distributions: United in Testing but Separated in Verification”. In: *Electron. Colloquium Comput. Complex.* TR25-105 (2025) (cit. on pp. 1, 9).
- [GR99] Oded Goldreich and Dana Ron. “A Sublinear Bipartiteness Tester for Bounded Degree Graphs”. In: *Comb.* 19.3 (1999), pp. 335–373 (cit. on p. 8).
- [HR22] Tal Herman and Guy N. Rothblum. “Verifying the unseen: interactive proofs for label-invariant distribution properties”. In: *STOC ’22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*. Ed. by Stefano Leonardi and Anupam Gupta. ACM, 2022, pp. 1208–1219. DOI: [10.1145/3519935.3519987](https://doi.org/10.1145/3519935.3519987) (cit. on pp. 1, 3, 9).
- [HR23] Tal Herman and Guy N. Rothblum. “Doubly-Efficient Interactive Proofs for Distribution Properties”. In: *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023*. IEEE, 2023, pp. 743–751. DOI: [10.1109/FOCS57990.2023.00049](https://doi.org/10.1109/FOCS57990.2023.00049) (cit. on pp. 1, 9).
- [HR24] Tal Herman and Guy N. Rothblum. “Interactive Proofs for General Distribution Properties”. In: *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*. IEEE, 2024, pp. 528–538. DOI: [10.1109/FOCS61266.2024.00041](https://doi.org/10.1109/FOCS61266.2024.00041) (cit. on pp. 1, 9).
- [HR25] Tal Herman and Guy N. Rothblum. “How to Verify Any (Reasonable) Distribution Property: Computationally Sound Argument Systems for Distributions”. In: *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025 (cit. on pp. 1, 9).
- [LRR12] Reut Levi, Dana Ron, and Ronitt Rubinfeld. “Testing Similar Means”. In: *ICALP (1)*. Vol. 7391. Lecture Notes in Computer Science. Springer, 2012, pp. 629–640 (cit. on p. 3).
- [LRR13] Reut Levi, Dana Ron, and Ronitt Rubinfeld. “Testing Properties of Collections of Distributions”. In: *Theory Comput.* 9 (2013), pp. 295–347 (cit. on p. 3).
- [Pan08] Liam Paninski. “A Coincidence-Based Test for Uniformity Given Very Sparsely Sampled Discrete Data”. In: *IEEE Trans. Inf. Theory* 54.10 (2008), pp. 4750–4755. DOI: [10.1109/TIT.2008.928987](https://doi.org/10.1109/TIT.2008.928987). URL: <https://doi.org/10.1109/TIT.2008.928987> (cit. on p. 8).
- [PRR06] Michal Parnas, Dana Ron, and Ronitt Rubinfeld. “Tolerant property testing and distance approximation”. In: *J. Comput. Syst. Sci.* 72.6 (2006), pp. 1012–1042 (cit. on p. 8).
- [Ras+09] Sofya Raskhodnikova et al. “Strong Lower Bounds for Approximating Distribution Support Size and the Distinct Elements Problem”. In: *SIAM J. Comput.* 39.3 (2009), pp. 813–842. DOI: [10.1137/070701649](https://doi.org/10.1137/070701649). URL: <https://doi.org/10.1137/070701649> (cit. on p. 8).

- [Ron08] Dana Ron. “Property Testing: A Learning Theory Perspective”. In: *Found. Trends Mach. Learn.* 1.3 (2008), pp. 307–402. DOI: [10.1561/2200000004](https://doi.org/10.1561/2200000004). URL: <https://doi.org/10.1561/2200000004> (cit. on p. 8).
- [RS96] Ronitt Rubinfeld and Madhu Sudan. “Robust Characterizations of Polynomials with Applications to Program Testing”. In: *SIAM J. Comput.* 25.2 (1996), pp. 252–271. DOI: [10.1137/S0097539793255151](https://doi.org/10.1137/S0097539793255151) (cit. on p. 8).
- [RVW13] Guy N. Rothblum, Salil P. Vadhan, and Avi Wigderson. “Interactive proofs of proximity: delegating computation in sublinear time”. In: *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*. Ed. by Dan Boneh, Tim Roughgarden, and Joan Feigenbaum. ACM, 2013, pp. 793–802. DOI: [10.1145/2488608.2488709](https://doi.org/10.1145/2488608.2488709) (cit. on p. 8).
- [Tel16] Roei Tell. “On Being Far from Far and on Dual Problems in Property Testing: [Extended Abstract]”. In: *ITCS*. ACM, 2016, pp. 103–110 (cit. on p. 9).
- [TSW24] Sri AravindaKrishnan Thyagarajan, Pratik Soni, and Ke Wu. “Game-Theoretically Fair Distributed Sampling”. In: *Annual International Cryptology Conference*. Springer. 2024, pp. 207–239 (cit. on p. 10).
- [Vad99] Salil Pravin Vadhan. *A study of statistical zero-knowledge proofs*. Thesis (Ph.D.)–Massachusetts Institute of Technology. ProQuest LLC, Ann Arbor, MI, 1999 (cit. on pp. 10, 13, 27, 28).
- [Val11] Paul Valiant. “Testing Symmetric Properties of Distributions”. In: *SIAM J. Comput.* 40.6 (2011), pp. 1927–1968. DOI: [10.1137/080734066](https://doi.org/10.1137/080734066). URL: <https://doi.org/10.1137/080734066> (cit. on p. 4).
- [Vas21] Prashant Nalini Vasudevan. *Lecture 03: Resources in Interactive Proofs*. 2021 (cit. on p. 22).
- [VV11] Gregory Valiant and Paul Valiant. “Estimating the unseen: an $n/\log(n)$ -sample estimator for entropy and support size, shown optimal via new CLTs”. In: *STOC*. ACM, 2011, pp. 685–694 (cit. on pp. 3, 8).
- [VV17] Gregory Valiant and Paul Valiant. “Estimating the Unseen: Improved Estimators for Entropy and Other Properties”. In: *J. ACM* 64.6 (2017), 37:1–37:41. DOI: [10.1145/3125643](https://doi.org/10.1145/3125643). URL: <https://doi.org/10.1145/3125643> (cit. on p. 8).
- [WAS22] Ke Wu, Gilad Asharov, and Elaine Shi. “A complete characterization of game-theoretically fair, multi-party coin toss”. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2022, pp. 120–149 (cit. on p. 10).
- [WY16] Yihong Wu and Pengkun Yang. “Minimax rates of entropy estimation on large alphabets via best polynomial approximation”. In: *IEEE Trans. Inform. Theory* 62.6 (2016), pp. 3702–3720. ISSN: 0018-9448. DOI: [10.1109/TIT.2016.2548468](https://doi.org/10.1109/TIT.2016.2548468). URL: <https://doi.org/10.1109/TIT.2016.2548468> (cit. on p. 8).

A Recap: Proof sketch of the Polarization Lemma (Lemma 3.1)

We recapitulate the proof of polarization lemma in [Vad99]. In what follows, $\ell, m, r \geq 1$ are parameters to be optimized. Note that one can simulate a distribution resulting from the XOR Lemma ([Vad99, Lemma 3.1.16], with parameter ℓ) and the Direct Product Lemma (Lemma 3.1.15

in [Vad99], with parameter m) using ℓ and m samples respectively. In what follows, we assume without loss of generality that $\delta \in (0, c]$ for a sufficiently small absolute constant $c > 0$ ($c = 1/8$ suffices), as one can replace δ by $\min(\delta, c)$ at the cost of only constant factors in the claimed result.

- Apply the XOR lemma on $\mathbf{p}, \mathbf{q}$ ℓ times to get $\mathbf{p}''', \mathbf{q}'''$. Then

$$\begin{aligned} d_{\text{TV}}(\mathbf{p}, \mathbf{q}) \geq \alpha &\implies d_{\text{TV}}(\mathbf{p}''', \mathbf{q}''') \geq \alpha^\ell \\ d_{\text{TV}}(\mathbf{p}, \mathbf{q}) \leq \beta &\implies d_{\text{TV}}(\mathbf{p}''', \mathbf{q}''') \leq \beta^\ell \end{aligned}$$

- Apply the Direct Product lemma on $\mathbf{p}''', \mathbf{q}'''$ m times to get $\mathbf{p}'', \mathbf{q}''$. Then

$$\begin{aligned} d_{\text{TV}}(\mathbf{p}''', \mathbf{q}''') \geq \alpha^\ell &\implies d_{\text{TV}}(\mathbf{p}'', \mathbf{q}'') \geq 1 - 2e^{-\frac{m\alpha^{2\ell}}{2}} \\ d_{\text{TV}}(\mathbf{p}''', \mathbf{q}''') \leq \beta^\ell &\implies d_{\text{TV}}(\mathbf{p}'', \mathbf{q}'') \leq m\beta^\ell \end{aligned}$$

- Finally, apply the XOR lemma r times on $\mathbf{p}'', \mathbf{q}''$ to get $\mathbf{p}', \mathbf{q}'$. Then

$$\begin{aligned} d_{\text{TV}}(\mathbf{p}'', \mathbf{q}'') \geq 1 - 2e^{-\frac{m\alpha^{2\ell}}{2}} &\implies d_{\text{TV}}(\mathbf{p}', \mathbf{q}') \geq \left(1 - 2e^{-\frac{m\alpha^{2\ell}}{2}}\right)^r \geq 1 - 2re^{-\frac{m\alpha^{2\ell}}{2}} \\ d_{\text{TV}}(\mathbf{p}'', \mathbf{q}'') \leq m\beta^\ell &\implies d_{\text{TV}}(\mathbf{p}', \mathbf{q}') \leq (m\beta^\ell)^r \end{aligned}$$

Overall, generating a sample from $\mathbf{p}'$ or $\mathbf{q}'$ takes $s := \ell mr$ samples from $\mathbf{p}, \mathbf{q}$. To get a concrete bound, set

$$t = 2 \ln \frac{1}{\delta}, \lambda := \min \left\{ \frac{\alpha^2}{\beta}, 2 \right\} \in (1, 2], \ell := \lceil \log_\lambda 4t \rceil, m := \frac{\lambda^\ell}{2\alpha^{2\ell}} \leq \frac{1}{2\beta^\ell}, r := \left\lceil \log \frac{1}{\delta} \right\rceil;$$

one can verify that $(m\beta^\ell)^r \leq \left(\frac{1}{2}\right)^r \leq \delta$, and that

$$2re^{-\frac{m\alpha^{2\ell}}{2}} = 2r \exp\left(-\frac{\lambda^\ell}{4}\right) \leq 2r \exp\left(-\frac{\lambda^{\log_\lambda 4t}}{4}\right) = 2r \exp(-t) = 2 \left\lceil \log \frac{1}{\delta} \right\rceil \cdot \delta^2 \leq \delta;$$

(the last inequality as $\delta \leq c$); and since

$$\begin{aligned} m = \frac{\lambda^\ell}{2\alpha^{2\ell}} &\leq \frac{1}{2} \frac{\lambda^{\log_\lambda(4t)+1}}{\alpha^{2(\log_\lambda(4t)+1)}} \\ &= \frac{1}{2} \frac{4t\lambda}{\alpha^2 \cdot \alpha^{\log_\alpha 4t \cdot 2 \frac{\log \alpha}{\log \lambda}}} \\ &\leq \frac{1}{2} \frac{4t\lambda}{\alpha^2 \cdot (4t)^{\frac{2 \log \alpha}{\log \lambda}}} \\ &= \frac{\lambda}{2\alpha^2} (4t)^{1 - \frac{2 \log \alpha}{\log \lambda}} \\ &\leq \frac{\lambda}{2\alpha^2} (4t)^{1 + 2 \log(1/\alpha) \cdot \max\left(\frac{1}{\log(\alpha^2/\beta)}, 1\right)} \\ &\leq \min\left(\frac{1}{2\beta}, \frac{1}{\alpha^2}\right) (4 \log(1/\delta))^{1 + \frac{2 \log(1/\alpha)}{\log(\alpha^2/\beta)} + 2 \log(1/\alpha)}. \end{aligned}$$

and so, in particular,

$$s = \ell mr \leq (\log(1/\delta))^{\frac{2 \log(1/\alpha)}{\log(\alpha^2/\beta)} + O(\log(1/\alpha))}.$$

This concludes the proof. □