

On the Complexity of Statistical and Multiplicative Zero-Knowledge

Daniel Alabi ¹ and Marshall Ball ²

¹ University of Illinois at Urbana-Champaign

² New York University

Abstract

Multiplicative zero-knowledge (MZK) replaces the negligible statistical distance required by statistical zero-knowledge (SZK) with symmetric (ε, δ) approximate max-divergence between real and simulated verifier views. Prior work established efficiency gains from this relaxation; we study its computational power.

With negligible completeness, soundness, and additive simulation errors,

$$\text{HVMZK}[\varepsilon] = \text{MZK}[\varepsilon] = \text{SZK} \quad \text{for } \varepsilon(n) = O(\log n),$$

where n is the input length. Taking the union over polynomial leakage bounds gives

$$\text{MZK}[\text{poly}] = \text{HVMZK}[\text{poly}] = \text{PSPACE},$$

even with zero additive simulation error. For every fixed $c > 0$, a polynomially padded TQBF language remains PSPACE-complete and belongs to $\text{MZK}[O(n^c)]$.

A linear-leakage protocol for SAT separates MZK from SZK unless the polynomial hierarchy collapses to its second level. Padding yields NP-complete separators at every fixed positive polynomial exponent. Under a subexponential weakening of the nonuniform nondeterministic Strong Exponential Time Hypothesis, every polylogarithmic level $\text{MZK}[O((\log n)^j)]$ with fixed $j \geq 2$ strictly contains SZK. The logarithmic threshold is tight relative to oracles: every polynomial-time computable, polynomially bounded superlogarithmic leakage budget admits an oracle separating classical MZK from quantum statistical zero-knowledge. Finally, fraction statistical knowledge complexity exactly characterizes one-sided real-to-simulator multiplicative simulation, with the sharp knowledge parameter and negligible statistical error.

Contents

1	Introduction	4
1.1	Our Results	4
1.2	Additional Related Work	8
1.3	Organization of Contents	9
2	Technical Overview	9
2.1	Collapse to Statistical Zero-Knowledge	9
2.2	A Conditional Separation at Polynomial Leakage	11
2.3	A Tight Relativized Threshold	12
2.4	Complete Problems Across Leakage Regimes	13
2.5	Fraction Knowledge and One-Sided Multiplicative Simulation	15
2.6	Why a Distribution-Free Symmetric Converse Fails	16
2.7	Repetition-Based Symmetrization	17
3	Preliminaries and Notation	19
3.1	Notation	19
3.2	Basic Complexity Background	20
3.3	Statistical and Multiplicative Zero-Knowledge	21
3.4	Fraction Statistical Knowledge Complexity	23
3.5	Arthur–Merlin Complexity Classes	23
4	Collapse to Statistical Zero-Knowledge	24
4.1	Setup	24
4.2	Views, Validity, and Fixed-Length Padding	24
4.3	Bounded Conditioning on Yes-Instances	25
4.4	Support Regularization and the Yes-Instance KL Bound	27
4.5	Projection onto a Genuine Prover	28
4.6	No-Instances: Low and High Sampler Success	29
4.6.1	Low success and five independent bits	30
4.6.2	High success, soundness, and data processing	30
4.7	Reductions to Entropy Difference and Statistical Difference	31
4.8	Combining the Queries by an NC^1 Truth-Table Reduction	32
5	A uniform-branch simulation lemma	33
6	A Conditional Unrelativized Separation	34
6.1	The randomized-witness protocol for SAT	35
6.2	Padding to obtain every fixed polynomial exponent	35
7	Polylogarithmic Leakage and Limited Nondeterminism	36
8	Fully Relativized Separation	38
8.1	Unambiguous computation and the oracle language	38
8.1.1	The classes UP and coUP	38
8.1.2	The classical oracle model	38
8.1.3	The unary witness language	39
8.2	Relativized quantum state distinguishability	39

8.2.1	States, trace distance, and fidelity	39
8.2.2	Oracle query gates and uniform circuit families	39
8.2.3	The relativized QSZK characterization	40
8.2.4	The imported black-box bounds	40
8.3	Finite-extension oracle diagonalization	41
8.4	The relativized multiplicative zero-knowledge protocol	41
9	Complete Problems and a Canonical Protocol Characterization	43
9.1	Polynomial leakage: PSPACE and TQBF	43
9.1.1	Fixed polynomial budgets and padding	44
9.2	A canonical reduction for fixed protocol profiles	44
9.2.1	Views generated by a prover	45
9.2.2	Hockey-stick divergence	45
9.2.3	Definition of the promise problem	45
9.3	The scope of the canonical reduction	46
9.4	Why an ordinary two-sampler problem does not yet follow	47
10	Implications for Fraction Knowledge Complexity	47
10.1	Implications for AM and coAM	49
10.2	One-Sided Characterization and the Symmetric Converse	49
10.3	The class-level converse	50
11	Towards a Symmetric Converse for Knowledge Complexity	50
11.1	What fraction knowledge does imply	51
11.2	There is no reverse KL bound	51
11.3	Bounded forward KL does not yield approximate max-divergence	51
11.4	Consequence for the proposed converse	52
12	When a Prover Can Generate the Simulator's Views	52
12.1	The prover and the product protocol	52
12.2	Smoothing the good conditional distribution	53
12.3	Random coordinate embedding	54
12.4	Negligibility, completeness, and soundness	55
12.5	What the argument assumes	55
13	Conclusion	55
14	Acknowledgements	55
A	A Communication Separation for Distributed Simulation	59

1 Introduction

Zero-knowledge proofs let a prover convince a verifier that a statement is true without revealing more than its validity [25, 26]. In statistical zero-knowledge (SZK), the verifier’s view must be reproducible by a polynomial-time simulator up to negligible statistical distance. The resulting class is well understood [33, 42]: honest-verifier and general-verifier statistical zero-knowledge coincide, STATISTICAL DIFFERENCE and ENTROPY DIFFERENCE are complete problems, and $\text{SZK} \subseteq \text{AM} \cap \text{coAM}$ [1, 22, 24, 34, 37]. In particular, unlike computational zero-knowledge under standard cryptographic assumptions [18, 19], statistical zero-knowledge is not expected to contain all of NP.

Talwar introduced differential zero-knowledge for distributed verification [41], motivated by differential privacy [14, 15]. It allows real and simulated views to differ by a multiplicative factor e^ε and an additive slack δ . Writing R for the real verifier-view distribution and Q for the simulator output, the requirement is that, for every event E ,

$$R(E) \leq e^\varepsilon Q(E) + \delta \quad \text{and} \quad Q(E) \leq e^\varepsilon R(E) + \delta.$$

We call this notion *multiplicative zero-knowledge* (MZK).¹ The parameter ε measures multiplicative leakage, while δ measures additive simulation error. We write $\text{HVMZK}[\varepsilon, \delta]$ for the honest-verifier class and $\text{MZK}[\varepsilon, \delta]$ when simulation is required against every probabilistic polynomial-time verifier. (We omit δ when it is negligible.)

Multiplicative simulation can reduce communication and client computation, as Talwar’s distributed verification protocols demonstrate [41]. Beimel, Nissim, and Omri showed a related separation: a distributed differentially private computation can communicate less than secure evaluation of a centralized private analysis [7]. Appendix A strengthens the comparison by holding the randomized functionality fixed. For m parties, with m a power of two, constant multiplicative simulation uses $O(m \log m)$ bits, whereas statistical simulation with error at most $1/(100m^2)$ requires $\Omega(m^2)$. The model uses private channels, independent private randomness, no correlated setup, and a fixed public communication schedule; security is against every static passive coalition of fewer than $m/3$ parties.

These efficiency gains leave open the effect on language classes. A constant multiplicative factor permits constant statistical distance, so it need not give a statistical simulator for the same protocol. We ask:

How does the computational power of multiplicative zero-knowledge depend on the allowed leakage?

Logarithmic leakage preserves the language class: honest-verifier and general-verifier MZK both coincide with SZK. This threshold is tight relative to oracles, even against quantum statistical zero-knowledge. We prove conditional unrelativized separations and $\text{MZK}[\text{poly}] = \text{PSPACE}$, give a canonical reduction preserving fixed leakage and proof-error profiles, and characterize one-sided multiplicative simulation by fraction statistical knowledge complexity. This characterization identifies an obstruction to symmetrization and a structural condition under which it succeeds.

1.1 Our Results

Unless stated otherwise, completeness, soundness, and additive simulation errors are negligible. The security parameter is the input length n . In the canonical circuit encoding below, profiles are instead

¹Talwar uses the term “differential zero-knowledge.” We use “multiplicative zero-knowledge” to emphasize that the simulation requirement compares real and simulated views without invoking a neighboring input relation.

evaluated at a distinguished unary tag; that representation is not identified with the ordinary class at the total encoding length. Our class characterizations, like those of statistical zero-knowledge, are formulated for promise problems.

Logarithmic leakage does not enlarge statistical zero-knowledge. Our main collapse theorem states that, for every nonnegative leakage profile $\varepsilon(n) = O(\log n)$,

$$\text{HVMZK}[\varepsilon] = \text{MZK}[\varepsilon] = \text{SZK}.$$

Constant or logarithmic leakage can therefore change a proof system's efficiency without enlarging its language class. The upper bound only requires acceptance probability bounded below by a positive constant, with negligible soundness and additive simulation errors. We do not assume that standard error amplification preserves a fixed leakage budget.

The proof reduces the language to SZK-complete problems. It first conditions the simulator on valid accepting conversations. The real-to-simulator inequality guarantees inverse-polynomial accepting mass when $\varepsilon = O(\log n)$, so this conditioning is efficient. After support regularization, the reverse inequality bounds relative entropy against a genuine prover's view. A simulation-based prover and the entropy chain rule express this bound as instances of ENTROPY DIFFERENCE, with STATISTICAL DIFFERENCE handling low accepting mass. This adapts the classical simulation-based prover and prefix-entropy method [24, 34]. The resulting SZK proof need not simulate the original protocol. With negligible completeness error, the reverse inequality alone gives overwhelming accepting mass at logarithmic leakage. Beyond that regime, the forward lower bound can be superpolynomially small, and the reverse bound can amplify negligible errors by a superpolynomial factor; the collapse argument does not extend to leakage $(\log n)^2$.

Conditional separations at every fixed polynomial exponent. We give a one-message protocol establishing $\text{SAT} \in \text{MZK}[n \ln 2, 0]$. The protocol has perfect soundness and negligible completeness error. Its prover sends a satisfying assignment except on an exponentially rare branch that sends a uniform assignment; the simulator always sends a uniform assignment. The rare branch supplies the support needed for both multiplicative inequalities, including against cheating verifiers. If the polynomial hierarchy does not collapse to its second level, then $\text{SAT} \notin \text{SZK}$ [10].

Polynomial padding shows that, for every fixed $\alpha > 0$, there is an NP-complete language L_α such that $L_\alpha \in \text{MZK}[O(N^\alpha), 0]$, where N is the padded input length. Under the same noncollapse assumption, $L_\alpha \notin \text{SZK}$. Choosing the rare-branch probability at the unpadded length keeps completeness error negligible in N .

A polylogarithmic hierarchy above the collapse. For fixed integers $j \geq 1$, let $\mathcal{M}_j = \text{MZK}[O((\log n)^j)]$. Then

$$\text{SZK} = \mathcal{M}_1 \subseteq \mathcal{M}_2 \subseteq \cdots.$$

Every $\mathcal{M}_j$ with $j \geq 2$ strictly contains SZK, assuming that some fixed-width DNF tautology problem has no subexponential-size nondeterministic circuits. This assumption follows from the nonuniform nondeterministic Strong Exponential Time Hypothesis of Carmosino et al. [11]. Section 7 states the assumption and proves the separation by padding limited-nondeterminism languages. Strictness between consecutive levels remains open.

The logarithmic threshold is tight relative to oracles. Fix a polynomial-time computable, polynomially bounded integer profile $\ell(n) \geq 1$ with $\ell(n) = \omega(\log n)$. There exist an oracle A and a

unary language $L_\ell(A)$ such that

$$L_\ell(A) \in \text{MZK}^A[(\ell(n) - 1) \ln 2, 0] \setminus \text{QSZK}^A.$$

Moreover, $L_\ell(A) \in \text{UP}^A \cap \text{coUP}^A$. The protocol sends $\ell(n) - 1$ prover bits, has perfect soundness, and has completeness error at most $2^{-(\ell(n)-1)}$. We reparameterize the singleton-oracle argument of Menda and Watrous [30], using their quantum query bound and a diagonalization that preserves earlier failures. A rare uniform branch then gives simulation against every classical PPT oracle verifier. The collapse at $O(\log n)$ holds relative to every oracle; above it, each such profile admits an oracle separation even from QSZK, which contains relativized SZK.

Polynomial leakage captures all interactive proofs. Taking the union over polynomial leakage bounds recovers interactive proofs:

$$\text{MZK}[\text{poly}] = \text{HVMZK}[\text{poly}] = \text{PSPACE}.$$

The equality holds even with zero additive simulation error.

The construction extends the rare-uniform-branch idea to arbitrary interactive proofs. If the prover sends at most $m(n)$ bits, mixing its strategy with uniform replies with probability $2^{-h(n)}$ gives pure symmetric simulation with leakage

$$\varepsilon(n) = \ln 2 \cdot \max\{m(n), h(n)\}.$$

The simulator runs the verifier and supplies uniform replies. For a polynomial-time computable, polynomially bounded integer-valued $h(n) = \omega(\log n)$, the modification adds only negligible completeness error and leaves soundness unchanged. Applying it to amplified interactive proofs and using $\text{IP} = \text{PSPACE}$ [39] proves the characterization. More generally, an interactive proof with negligible proof errors and at most $b(n) = \omega(\log n)$ prover bits lies in $\text{MZK}[b(n) \ln 2, 0]$ under the usual polynomial bounds (Corollary 3); this connects multiplicative leakage to laconic provers [17, 23].

Consequently, TQBF is complete for both polynomial-leakage classes. For every fixed $c > 0$, an appropriately polynomially padded version of TQBF is PSPACE-complete and belongs to $\text{MZK}[O(N^c)]$. It is therefore complete for that fixed-budget class under ordinary polynomial-time many-one reductions. This does not identify $\text{MZK}[O(N^c)]$ with PSPACE: ordinary reductions can increase the input length and hence change the leakage budget.

A canonical reduction preserving fixed profiles. For fixed leakage and additive-error profiles ε, δ , and fixed proof-error profiles $c_{\text{ip}}, s_{\text{ip}}$ satisfying $1 - c_{\text{ip}}(n) > s_{\text{ip}}(n)$, we introduce *Interactive Max-Divergence Acceptance* (IMDA). Every honest-verifier protocol meeting these four bounds compiles in polynomial time to instances

$$x \mapsto (1^{|x|}, V_x, S_x)$$

of the corresponding fixed-profile promise problem, with one fixed polynomial encoding bound B chosen large enough for that source protocol. The verifier and simulator circuits explicitly encode coins, message slots, and outputs. A yes-instance admits an online prover whose view meets the acceptance and two simulation bounds. On a no-instance, every online prover accepts with probability at most $s_{\text{ip}}(n)$.

The unary tag preserves the original parameter n . A guarantee $\delta(n)$ need not meet the prescribed bound $\delta(N)$ at the larger output length N . Thus this is a canonical reduction for fixed protocol

profiles, not completeness for $\text{HVMZK}[\varepsilon, \delta]$ at total input length. It applies, in particular, at leakage $(\log n)^2$.

Each prover message depends only on the input, transcript so far, and private randomness. A distribution supported on valid transcripts need not be generated by such a prover, and a view generated by an unbounded prover need not be efficiently samplable. The representation therefore retains interaction. It does not yield general-verifier completeness or an interaction-free problem analogous to `STATISTICAL DIFFERENCE`.

Fraction knowledge exactly captures one-sided simulation. We characterize one-sided simulation by fraction statistical knowledge complexity [20,21,34]. Suppose only the real-to-simulator inequality $R(E) \leq e^\varepsilon Q(E) + \delta$ holds for every event E . Under our polynomial bounds on parameters and sampler lengths, for every inverse-exponential negligible ξ , there exists a set G of random tapes of a suitably extended simulator such that

$$\Pr[G] \geq e^{-\varepsilon}, \quad \text{SD}(Q_G, R) \leq \delta + 2\xi.$$

Here Q_G is the simulator output conditioned on its random tape belonging to G . Thus, for negligible δ , the same protocol has fraction statistical knowledge complexity at most

$$\frac{\varepsilon}{\ln 2}.$$

The leading term is tight in the exact case. The set G need only exist; efficient recognition of good tapes is not required.

Conversely, conditioning on a good-tape set of density at least 2^{-k} gives the real-to-simulator inequality with exponent $k \ln 2$. Writing $\text{HVoMZK}^\rightarrow$ for the honest-verifier notion requiring only this direction, we obtain the exact indexed equality under our strict-PPT conventions,

$$\text{SKC}_{\text{frac}}[k] = \text{HVoMZK}^\rightarrow[k \ln 2].$$


Obstructions to a symmetric converse. Fraction knowledge does not control the simulator's behavior on bad tapes. That part of its output may have substantial mass on events that are impossible, or extremely unlikely, in the real interaction. Our examples show that fraction knowledge yields neither a reverse multiplicative bound nor a reverse relative-entropy bound for the given unconditioned simulator. We also show that even bounded forward relative entropy does not, in general, imply a constant multiplicative exponent with negligible additive slack. These are separate obstructions to a conversion based only on the two distributions.

At the class level, a general symmetric converse would settle a longstanding knowledge-complexity collapse question. Our logarithmic collapse theorem and the identity $\text{SZK} = \text{SKC}_{\text{frac}}[0]$ imply that

$$\text{SKC}_{\text{frac}}[O(1)] \subseteq \text{MZK}[O(1)] \iff \text{SKC}_{\text{frac}}[O(1)] = \text{SKC}_{\text{frac}}[0].$$

Symmetrization when a prover can generate the simulator's views. Suppose one uniform, possibly unbounded prover generates a fraction simulator's unconditioned output distribution exactly on every yes-instance, choosing each message from the input, transcript so far, and private randomness. If the fraction knowledge is $k = O(\log n)$ and the proof and statistical simulation errors are negligible, a polynomial sequential repetition yields an honest-verifier symmetric multiplicative zero-knowledge protocol with leakage $O(k + 1)$ and negligible additive error.

The verifier accepts if any repeated copy accepts. The prescribed prover uses the original strategy in one privately selected uniform coordinate and the strategy that generates the simulator's views



If PH does not collapse to its second level, these padded complete languages lie outside SZK.

Explicit linear-leakage witness: $\text{SAT} \in \text{MZK}[n \ln 2]$.

Figure 1: The complexity of multiplicative zero-knowledge as the leakage bound grows. Completeness, soundness, and additive simulation errors are negligible. Here $\alpha > 0$ is any fixed constant, and poly denotes the union over polynomial leakage bounds. The first inclusion is strict under Definition 10; the other displayed inclusions need not be strict. Relative to oracles, the logarithmic threshold is tight for every polynomial-time computable, polynomially bounded superlogarithmic profile (Theorem 12).

in all others; the simulator outputs independent copies of its original distribution. A concentration bound supplies the reverse simulation inequality. Together with the collapse theorem, this places the underlying promise problem in SZK under this hypothesis on the simulator’s output.

The emerging complexity landscape. Figure 1 illustrates the dependence on the leakage budget. Logarithmic leakage preserves SZK; every fixed positive polynomial exponent already allows padded PSPACE-complete languages; and the union over polynomial bounds equals PSPACE, even with pure simulation. Under the stated nonuniform assumption, every fixed polylogarithmic degree $j \geq 2$ already exceeds SZK. The logarithmic threshold is tight relative to oracles. Unrelativized unconditional separations, strictness between consecutive polylogarithmic levels, and honest-verifier/general-verifier equivalence at fixed superlogarithmic profiles remain open. So does an interaction-free complete problem analogous to STATISTICAL DIFFERENCE.

1.2 Additional Related Work

Multiplicative simulation. Backes, Kate, Meiser, and Ruffing [5] defined computational differential zero-knowledge for bounded weak randomness. Talwar [41] formulated its information-theoretic counterpart for approximate norm verification and robust vector summation. PINE [36] adopts Talwar’s criterion, explicitly requires an efficiently samplable reference distribution, and gives exact norm verification; its differential variant reduces communication in some parameter regimes. Both Talwar and PINE require symmetric (ε, δ) -closeness. Their reference distribution depends only on the protected parties’ permitted information. We study language classes in the conventional interactive-proof model, with a public statement, an unbounded prover, and simulation of the verifier’s full view, for both honest and arbitrary efficient verifiers.

Knowledge complexity and entropy methods. Fraction statistical knowledge complexity measures the density of simulator tapes on which simulation succeeds [20, 21]. Petrank and Tardos [34] proved $\text{SKC}_{\text{frac}}[O(\log n)] \subseteq \text{AM} \cap \text{coAM}$. We characterize fraction knowledge by real-to-simulator

domination. The reverse inequality bounds additional simulator mass and supports our SZK bound. The unrestricted fraction-knowledge collapse remains open; logarithmic *hint* knowledge already collapses to SZK [37, Theorem 4.15]. Our proof uses the classical simulation-based prover and prefix-entropy identity [34], entropy completeness, Boolean closure, and honest-verifier/general-verifier equivalence [22, 24, 37]. The resulting SZK proof need not simulate the original protocol.

Our minimum-KL formulation includes the Jensen–Shannon divergence problem of Berman, Degwekar, Rothblum, and Vasudevan [9] as a special case (Remark 2).

1.3 Organization of Contents

Section 2 gives the full technical overview, and Section 3 fixes the definitions and conventions. Sections 4 and 5 prove the logarithmic collapse and uniform-branch lemma. Section 6 gives the SAT and padding separations. Section 7 defines the polylogarithmic hierarchy and proves its conditional separation from SZK. The remaining technical sections develop fraction knowledge and its symmetric-converse obstructions (Sections 10–12), the oracle separation (Section 8), and the polynomial endpoint and canonical fixed-profile encoding (Section 9). These appear in that order in the submission appendices. Appendix A proves the strengthened distributed communication separation. Section 13 discusses the remaining hierarchy questions.

2 Technical Overview

We explain the main constructions and the parameter choices that make them work. The collapse uses both directions of simulation to bound the minimum KL divergence from a sampled view to a view generated by a prover. The separations exploit the cost of finding a witness, and the knowledge-complexity results isolate the role of simulator mass outside the real distribution.

A recurring idea: rare-branch support regularization. Several constructions control likelihood ratios by changing only a negligible amount of probability mass. A view that has positive simulator probability but zero real probability precludes pure symmetric simulation. We supply this missing support by having the prover privately choose a uniform-message strategy with probability θ ; the simulator always uses that strategy. The modification increases completeness error by at most θ , preserves soundness, and bounds the simulator-to-real ratio by $1/\theta$. A forward bound in terms of prover communication completes the simulation analysis. This gives our witness protocols and extends to arbitrary interactive proofs at polynomial leakage. The logarithmic collapse proof uses a related regularization of auxiliary distributions to bound relative entropy.

2.1 Collapse to Statistical Zero-Knowledge

Theorem 1 (Informal form of Theorem 8). *Symmetric multiplicative simulation with leakage $\varepsilon(n) = O(\log n)$ and negligible additive, completeness, and soundness errors does not enlarge statistical zero-knowledge: every promise problem with such a protocol lies in SZK. The upper bound also holds when yes-input acceptance is bounded below by a positive constant and soundness and additive simulation errors are negligible.*

The proof reduces the source problem to Statistical Difference and Entropy Difference, adapting the simulation-based prover and prefix-entropy framework of Petrank–Tardos and Goldreich–Vadhan [24, 34]. The full argument appears in Section 4.

Choose C such that $\varepsilon(n) \leq C \ln n$ for all sufficiently large n , and fix a protocol (P, V) satisfying the theorem. For an input x , write $R_x = \text{View}_{P,V}(x)$ for the real view, $Q_x = S(x)$ for the simulator output, and A_x for the event of a valid accepting view. On a yes-input, completeness and the real-to-simulator inequality give

$$q_x := Q_x(A_x) \geq e^{-\varepsilon(n)}(R_x(A_x) - \delta(n)) \geq \frac{1}{\text{poly}(n)}, \quad (1)$$

for negligible $\delta(n)$. A bounded rejection sampler therefore finds an accepting simulator view with negligible failure probability in strict polynomial time. To see how the two simulation directions combine, put $p_x = R_x(A_x)$ and $Q_{A,x} = Q_x \mid A_x$. The reverse inequality gives

$$Q_{A,x}(E) \leq \frac{e^{2\varepsilon(n)}}{p_x - \delta(n)} R_x(E) + \frac{e^{\varepsilon(n)} \delta(n)}{p_x - \delta(n)}.$$

One factor e^ε comes from reverse simulation and the other from the lower bound on the conditioning probability. The additive term remains negligible because p_x is bounded below and e^ε is polynomial.

Additive domination alone does not bound KL divergence: even negligible mass outside the support of R_x can make it infinite. Let $R_{0,x}$ be the full-support view of a prover sending uniform messages. The bounded sampler uses $R_{0,x}$ on failure; denote its output by X_x . For $\gamma = 2^{-n}$, define

$$Y_x = (1 - \gamma)X_x + \gamma R_{0,x}, \quad W_x = (1 - \gamma)R_x + \gamma R_{0,x}.$$

Now Y_x is efficiently samplable, W_x is a genuine prover view, and W_x gives every valid view probability at least $2^{-\text{poly}(n)}$. Clipping the additive excess in the analysis costs only a negligible amount of relative entropy: its mass is negligible and the logarithm of the inverse minimum probability is polynomial. This gives

$$D_{\text{KL},2}(Y_x \parallel W_x) \leq 2C \log_2 n + \text{negl}(n) \quad (x \in \Pi_Y), \quad (2)$$

With acceptance probability only bounded below by $p_0 > 0$, the bound becomes $2C \log_2 n + \log_2(1/p_0) + \text{negl}(n)$ and remains logarithmic. For negligible completeness error c , the reverse inequality also gives $Q_x(A_x^c) \leq e^\varepsilon c + \delta = \text{negl}(n)$ directly. In this case it already supplies overwhelming accepting mass; the argument above also covers constant positive completeness.

Although Y_x is close to a genuine view in relative entropy, it may correlate an early prover message with verifier coins revealed only later. Soundness therefore does not apply directly to Y_x . Define the simulation-based prover P_{Y_x} by sampling, at each prover turn, from the next-message conditional distribution induced by Y_x given the public transcript prefix, and let $T_{Y_x} = \text{View}_{P_{Y_x},V}(x)$. This is a genuine, possibly unbounded prover view. The reduction never runs P_{Y_x} ; it uses this prover only in the soundness analysis. The KL chain rule shows that

$$D_{\text{KL}}(Y_x \parallel T_{Y_x}) \leq D_{\text{KL}}(Y_x \parallel W_x). \quad (3)$$

Indeed, the difference is a sum, over prover turns and Y_x -distributed prefixes, of nonnegative conditional KL divergences. Equations (2)–(3) give a constant K for which $D_{\text{KL},2}(Y_x \parallel T_{Y_x}) \leq K \log_2 n$ on yes-instances.

The bounded sampler also runs on no-inputs, where no simulation guarantee is assumed. Let b_x be its probability of finding an accepting simulator view. If $b_x > 3/4$, then Y_x accepts with probability at least $3/4 - o(1)$, while soundness gives $T_{Y_x}(A_x) \leq s(n)$ for negligible $s(n)$. High

success forces $s(n) > 0$, because T_{Y_x} is positive on every Y_x -positive view. Data processing under the acceptance bit therefore gives

$$D_{\text{KL},2}(Y_x \| T_{Y_x}) \geq \left(\frac{3}{4} - o(1) \right) \log_2 \frac{1}{s(n)} - 1 = \omega(\log n). \quad (4)$$

A second chain-rule calculation expresses this gap through prefix entropies. Let $Y_{x,i}$ be the i -message prefix distribution, d the final position containing the verifier tape, $\mathcal{I}_P$ the set of prover turns, and $r(n)$ the verifier random-tape length. Then

$$D_{\text{KL},2}(Y_x \| T_{Y_x}) = r(n) - H(Y_{x,d}) + \sum_{i \in \mathcal{I}_P} (H(Y_{x,i}) - H(Y_{x,i-1})). \quad (5)$$

Independent products of the efficiently samplable prefix distributions produce circuits U_x, W_x satisfying $H(U_x) - H(W_x) = D_{\text{KL},2}(Y_x \| T_{Y_x})$. Adding an intermediate number of uniform bits to W_x gives an *Entropy Difference* instance that has opposite promised answers on yes-instances and high-success no-instances [24]. A Statistical Difference query handles low-success no-inputs. Both problems are SZK-complete, and SZK is closed under the resulting truth-table reduction. The full proof verifies its three-valued promise semantics, including cases where one query falls outside its promise.

The same entropy identity connects the proof to polarization. For any efficiently sampled valid view Q , it expresses the minimum KL divergence from Q to a view generated by a prover as an entropy difference. Thus inverse-polynomial gaps in this minimum divergence reduce to statistical-distance gaps that can be polarized to 2^{-k} versus $1 - 2^{-k}$. Jensen–Shannon divergence is a special case: the verifier receives Y before revealing a fresh fair bit B , while the sampler chooses B first and draws $Y \sim C_B$. The minimum is then $I(Y; B) = \text{JS}(C_0, C_1)$; Remark 2 gives the reduction.

2.2 A Conditional Separation at Polynomial Leakage

Assuming the polynomial hierarchy does not collapse, the logarithmic collapse cannot extend to linear leakage [1, 2, 10, 37]. Section 6 proves this using a randomized witness protocol for SAT.

Theorem 2 (Informal form of Theorem 9). *If PH does not collapse to its second level, then $\text{SAT} \notin \text{SZK}$ but $\text{SAT} \in \text{MZK}[N \ln 2]$.*

Let φ be a satisfiable formula of encoding length N on $m \leq N$ variables, fix a satisfying assignment w , and set $\eta = 2^{-N}$. The prover sends w with probability $1 - \eta$ and an independent $U_m \leftarrow \text{Unif}(\{0, 1\}^m)$ with probability η ; the simulator always outputs U_m . If R_w and Q denote the corresponding message distributions, then

$$R_w(y) = (1 - \eta)\mathbf{1}[y = w] + \eta 2^{-m}, \quad Q(y) = 2^{-m}. \quad (6)$$

At $y = w$, $R_w(w)/Q(w) \leq 2^m \leq 2^N$ and $Q(w)/R_w(w) \leq 1$; at every $y \neq w$, the two ratios are respectively η and $\eta^{-1} = 2^N$. Hence $R_w \approx_{N \ln 2, 0} Q$. For any PPT verifier V^* , its complete view is a randomized postprocessing K of the sole prover message. Thus, for every event E ,

$$\begin{aligned} (KR_w)(E) &\leq 2^N (KQ)(E), \\ (KQ)(E) &\leq 2^N (KR_w)(E). \end{aligned} \quad (7)$$

The protocol has perfect soundness and completeness error at most 2^{-N} . Hence $\text{SAT} \in \text{MZK}[N \ln 2]$. If PH does not collapse to its second level, then $\text{SAT} \notin \text{SZK}$: otherwise NP-completeness and $\text{SZK} \subseteq \text{AM} \cap \text{coAM}$ would imply $\text{coNP} \subseteq \text{AM}$, which collapses PH to its second level.

The same phenomenon occurs at every fixed polynomial leakage exponent.

Corollary 1 (Informal form of Corollary 4). *Given any $\alpha > 0$, there exists an NP-complete language L_α in $\text{MZK}[O(N^\alpha)]$ which is outside SZK, assuming the polynomial hierarchy does not collapse.*

For $k \geq 2$, let $\text{PadSAT}_k = \{\varphi \# 1^{|\varphi|^k} : \varphi \in \text{SAT}\}$. If $s = |\varphi|$, then the padded length is $N = \Theta(s^k)$. Run the preceding protocol with the choice $\eta = 2^{-s}$, rather than 2^{-N} . The leakage is

$$\begin{aligned} s \ln 2 &= O(N^{1/k}), \\ 2^{-s} &= 2^{-\Theta(N^{1/k})} = \text{negl}(N). \end{aligned} \tag{8}$$

Given any $\alpha > 0$, choosing $k > 1/\alpha$ yields an NP-complete language L_α in $\text{MZK}[O(N^\alpha)]$ and, under the same assumption that PH does not collapse, outside SZK.

Reaching polylogarithmic leakage. Section 7 uses a larger padding scale. For a fixed integer $j \geq 2$, pad an m -variable fixed-width DNF formula to length $N = 2^{\lceil m^{1/j} \rceil}$. A falsifying assignment then has only $m = O((\log N)^j)$ bits, so the uniform-branch construction gives a pure MZK proof for non-tautology at this leakage. If the padded language were in SZK, its complement would lie in $\text{AM} \subseteq \text{NP/poly}$. Hardwiring the padding into those circuits would decide the original tautology problem with nondeterministic circuits of size $N^{O(1)} = 2^{O(m^{1/j})} = 2^{o(m)}$. The nonuniformity comes from fixing Arthur's random coins: after amplifying the AM error below $2^{-N^{-1}}$, some coin string works for all inputs of length N simultaneously. That string becomes circuit advice. Definition 10 rules out the resulting subexponential-size nondeterministic circuits. Thus each fixed polylogarithmic degree above one conditionally exceeds SZK; the argument does not separate consecutive degrees.

2.3 A Tight Relativized Threshold

The logarithmic collapse holds relative to every oracle. We show it is tight: every polynomial-time computable, polynomially bounded integer profile $\ell(n) \geq 1$ with $\ell(n) = \omega(\log n)$ admits an oracle separation, even from quantum statistical zero-knowledge. The proof reparameterizes the singleton-oracle argument of Menda and Watrous [30]. Section 8 gives the quantum background and full proof.

Theorem 3 (Informal version of Theorem 11). *Fix such a profile ℓ . Let $\mathcal{A}$ contain the oracles with exactly one string a_r of each positive length r , and set $L_\ell(A) = \{0^n : n \geq 1 \text{ and } a_{\ell(n)} \text{ begins with } 1\}$. There exists $A \in \mathcal{A}$ such that $L_\ell(A) \in \text{UP}^A \cap \text{coUP}^A$ and $L_\ell(A) \notin \text{QSZK}^A$.*

Uniqueness gives the two unambiguous machines: on 0^n , guess the $\ell(n) - 1$ suffix bits and query the candidate with prefix 1 or 0. The complement machine also accepts empty and nonunary strings. Classical machines use membership queries; quantum circuits may query coherently. By the relativized Quantum State Distinguishability characterization, excluding QSZK^A amounts to defeating every polynomial-time uniform circuit-pair family.

Fix one such family and a polynomial $p(n)$ bounding its query count and address lengths. The superlogarithmic hypothesis gives, for all sufficiently large n ,

$$p(n) \leq \frac{2^{\ell(n)/2}}{20736}. \tag{9}$$

Consider its pair on input 0^n , put $r = \ell(n)$, and fix all oracle slices other than length r . Write d_B for the pair's output-state trace distance when the remaining black box is $B \subseteq \{0, 1\}^r$. The Menda–Watrous fidelity bound gives at least $(2/3)^{2^r}$ strings z with $|d_{\{z\}} - d_\emptyset| < 1/6$. If $d_\emptyset \leq 1/2$, at least $2^r/6$ of these strings begin with 1 and force a yes-input with distance below $2/3$. If $d_\emptyset \geq 1/2$, at least $2^r/6$ begin with 0 and force a no-input with distance above $1/3$. Either case supplies a singleton that defeats the pair.

We choose one oracle by diagonalizing over the countably many clocked uniform families. At stage i , choose n_i so that $r_i = \ell(n_i)$ exceeds every previously fixed slice length and satisfies (9) for the i th family's bound p_i . Fix every other still-free slice through $\max\{r_i, p_i(n_i)\}$, then choose the length- r_i singleton to defeat the pair on 0^{n_i} . Every slice that pair can query is now fixed; later stages preserve its failure. Both choices are possible because $\ell(n) = \omega(\log n)$. This constructs the oracle for the chosen profile directly; padding a previously fixed oracle would not supply the required lower bound.

Theorem 4 (Informal version of Theorem 12). *For every such ℓ , there exist an oracle A and a unary language $L_\ell(A)$ such that*

$$L_\ell(A) \in \text{MZK}^A[(\ell(n) - 1) \ln 2, 0] \setminus \text{QSZK}^A.$$

The protocol communicates $\ell(n) - 1$ prover bits, has perfect soundness, and has completeness error at most $2^{-(\ell(n)-1)}$.

Fix this oracle. Empty and nonunary inputs are rejected before interaction. On a yes-input 0^n , put $q = \ell(n) - 1$, let y^* be the unique string with $1y^* \in A$, and set $\eta = 2^{-q}$. The prover sends y^* with probability $1 - \eta$ and $U_q \leftarrow \text{Unif}(\{0, 1\}^q)$ with probability η ; the simulator always outputs U_q . Thus

$$R_{y^*}(y) = (1 - \eta)\mathbf{1}[y = y^*] + \eta 2^{-q}, \quad Q(y) = 2^{-q}, \quad (10)$$

and both pointwise likelihood ratios are at most 2^q , with equality in the reverse direction away from y^* . Hence $R_{y^*} \approx_{q \ln 2, 0} Q$. The verifier accepts iff $1Y \in A$, so soundness is perfect and completeness error is at most 2^{-q} , which is negligible since $q = \omega(\log n)$. For $q = 0$, the verifier simply queries 1 and completeness is perfect. Randomized postprocessing gives the same simulation bound for every cheating oracle verifier's complete view. Since $\text{SZK}^A \subseteq \text{QSZK}^A$, we obtain

$$L_\ell(A) \in \text{MZK}^A[(\ell(n) - 1) \ln 2, 0] \setminus \text{SZK}^A. \quad (11)$$

The choice $\ell(n) = n$ recovers linear leakage; profiles of order $\log n \log \log n$ or $(\log n)^2$ are also covered. When $\ell(n) = O(\log n)$, the verifier can enumerate the whole oracle slice in polynomial time.

2.4 Complete Problems Across Leakage Regimes

At an unspecified polynomial budget, even pure symmetric multiplicative simulation captures all of PSPACE. At a fixed intermediate budget, such as $\varepsilon(n) = (\ln n)^2$, we give a canonical reduction for protocols with fixed error profiles. The reduction retains the verifier and the online constraints on prover messages, and evaluates profiles at a unary tag rather than the full encoding length. Section 9 proves both results.

Theorem 5 (Informal form of Theorem 13). *With polynomially bounded leakage,*

$$\text{MZK}[\text{poly}] = \text{HVMZK}[\text{poly}] = \text{PSPACE}, \quad (12)$$

even when the additive simulation error is zero.

Consequently, TQBF is complete for both the general-verifier and honest-verifier classes under ordinary polynomial-time many-one reductions. The proof uses the uniform-branch lemma. Consider an interactive proof (P, V) with at most $m(n)$ prover bits under the fixed interface of Lemma 3, and let V^* be any PPT verifier. The strict-PPT simulator U^* runs V^* within that interface and supplies fresh uniform bits in each prover slot. Both prover branches enforce the same interface and termination rules. Modify the prescribed prover so that, before the interaction, it follows P with probability $1 - \theta$ and the same uniform-message strategy with probability $\theta = 2^{-h(n)}$. If R^* is the view generated by the original prover and R_θ^* is the modified view, then

$$R_\theta^* = (1 - \theta)R^* + \theta U^*. \quad (13)$$

Fixing a full view z fixes the verifier coins and all prover replies on the observed path. If that path contains $c(z) \leq m(n)$ prover bits, the uniform strategy sends those replies with probability $2^{-c(z)}$, while any prover does so with probability at most one. Thus $R^*(z) \leq 2^{c(z)}U^*(z) \leq 2^{m(n)}U^*(z)$, including views that terminate early. Combining this relation with (13) gives the two pointwise bounds

$$R_\theta^*(z) \leq 2^{m(n)}U^*(z), \quad U^*(z) \leq 2^{h(n)}R_\theta^*(z). \quad (14)$$

Thus the simulation is pure and symmetric with leakage $\ln 2 \cdot \max\{m(n), h(n)\}$; soundness is unchanged, while completeness decreases by at most θ . Applying this transformation with $h(n) = n$ to an amplified polynomial-communication proof for a PSPACE language gives the nontrivial inclusion $\text{PSPACE} \subseteq \text{MZK}[\text{poly}]$. The reverse inclusions follow by forgetting simulation and using $\text{IP} = \text{PSPACE}$.

Corollary 2 (Informal form of Corollary 5). *For every fixed $c > 0$, a polynomially padded version of TQBF is complete for $\text{MZK}[O(N^c)]$.*

For $c < 1$, tune the mixing probability together with the padding. If an amplified TQBF proof on length- ℓ formulas uses at most ℓ^D prover bits, choose $0 < b < c$ and $k > D/b$, and pad to $N = \Theta(\ell^k)$. The probability $2^{-\lceil N^b \rceil}$ is negligible in N and yields

$$\varepsilon(N) \leq \ln 2 \cdot \max\{\ell^D, \lceil N^b \rceil\} = O\left(N^{\max\{D/k, b\}}\right) = O(N^c). \quad (15)$$

Using probability 2^{-N} instead would already incur $\Theta(N)$ leakage and would not establish the claimed sublinear budgets.

Theorem 6 (Informal form of Theorem 14). *Fix leakage and additive-error profiles ε, δ and proof-error profiles $c_{\text{ip}}, s_{\text{ip}}$ with $1 - c_{\text{ip}}(n) > s_{\text{ip}}(n)$. Every honest-verifier protocol meeting these four bounds has a polynomial-time reduction to $\text{IMDA}_{\varepsilon, \delta}^{c_{\text{ip}}, s_{\text{ip}}}[B]$ for a fixed polynomial encoding bound B large enough for that protocol. The reduction preserves the original input length as a unary tag.*

An instance is a tagged triple $(1^n, \mathcal{V}, \mathcal{S})$, with an explicit Boolean circuit encoding, random-coin inputs, and message and output slots, all within the one fixed bound $B(n)$. Writing Q for the output law of $\mathcal{S}$, its yes-condition asks for an accepting distribution R induced by some online prover for $\mathcal{V}$ such that

$$\begin{aligned} R(A_{\mathcal{V}}) &\geq 1 - c_{\text{ip}}(n), \\ \text{HS}_{\varepsilon(n)}(R\|Q) &\leq \delta(n), \quad \text{HS}_{\varepsilon(n)}(Q\|R) \leq \delta(n), \end{aligned} \quad (16)$$

where $\text{HS}_\alpha(P\|Q) = \sup_E (P(E) - \alpha Q(E))$. The no-condition bounds every online prover's acceptance probability by $s_{\text{ip}}(n)$. Views generated by a prover have the following factorization, where $r \in \{0, 1\}^q$

is the verifier tape and a_i is the i th prover message:

$$R(r, a_1, \dots, a_t) = 2^{-q} \prod_{i=1}^t K_i(a_i \mid h_{i-1}), \quad (17)$$

Here K_i depends on the public history h_{i-1} , without access to unrevealed verifier coins. The reduction compiles the verifier and simulator of a protocol meeting the four profiles on $x \in \{0, 1\}^n$ into circuits $\mathcal{V}_x, \mathcal{S}_x$ and maps $x \mapsto (1^n, \mathcal{V}_x, \mathcal{S}_x)$. On a source yes-input the honest prover generates the required view; on a source no-input soundness bounds every prover strategy. Disjointness of the target promises follows from $1 - c_{\text{ip}}(n) > s_{\text{ip}}(n)$.

The reduction evaluates every profile at the source length n . Ordinary class membership would require the corresponding bounds at the full encoding length N , which the tag does not supply. For example, $\delta(n) = 2^{-n}$ does not imply $\delta(N) = 2^{-N}$ when $N = n^2$. Moreover, an arbitrary negligible-error protocol need not meet a particular fixed pair $c_{\text{ip}}, s_{\text{ip}}$. The theorem is therefore a reduction for protocols satisfying those bounds, not an ordinary honest-verifier completeness theorem or a general-verifier membership claim.

The encoding retains interaction because ordinary efficient samplers do not capture all the required constraints. At $\varepsilon(n) = (\ln n)^2$, the forward simulation inequality alone guarantees only $e^{-\varepsilon(n)}(1 - c_{\text{ip}}(n) - \delta(n)) = n^{-\ln n}(1 - c_{\text{ip}}(n) - \delta(n))$ accepting simulator mass, too little for polynomial-time rejection sampling. At this leakage, factors of e^ε can also turn negligible error terms into nonnegligible ones. A real view may require an unbounded prover, and an arbitrary distribution on valid transcripts can correlate early messages with unrevealed verifier randomness. Thus a complete problem involving only ordinary efficient samplers does not follow from this encoding. Section 7 nevertheless separates the ordinary class at $O((\ln n)^2)$ from SZK under its explicit nonuniform assumption; an unconditional unrelativized separation remains open.

2.5 Fraction Knowledge and One-Sided Multiplicative Simulation

Section 10 converts one-sided multiplicative simulation into fraction statistical knowledge without changing the protocol. Fix a yes-input, let $R = \text{View}_{P,V}(x)$ be the real honest-verifier view, and let $Q = S(x)$ be the output of a strict finite-coin PPT simulator. Only the real-to-simulator inequality $R \preceq_{\varepsilon, \delta} Q$ is required.

Theorem 7 (Informal form of Theorem 15). *If $R \preceq_{\varepsilon, \delta} Q$ and the parameter and sampler lengths have the polynomial bounds fixed in the paper, then, for every $\xi = 2^{-p(n)}$ with a positive integer-valued polynomial p of degree at least one, appending polynomially many unused random bits gives the simulator a set G of tapes such that*

$$\Pr[G] \geq e^{-\varepsilon}, \quad \text{SD}(Q_G, R) \leq \delta + 2\xi, \quad (18)$$

where Q_G is the simulator output conditioned on G . The appended tape length depends only on n , while $G = G_x$ may depend on the input.

Consequently, when δ is negligible, the same protocol has fraction knowledge complexity

$$k \leq \frac{\varepsilon}{\ln 2}. \quad (19)$$

Write the simulation inequality in hockey-stick form:

$$\sum_z (R(z) - e^\varepsilon Q(z))_+ \leq \delta. \quad (20)$$

The common subdistribution $\mu(z) = \min\{R(z), e^\varepsilon Q(z)\}$ has mass $m \geq 1 - \delta$. The subdistribution $\tau = e^{-\varepsilon} \mu \leq Q$ therefore has mass ρm , where $\rho = e^{-\varepsilon}$. Selecting only this mass would lose the factor m in the good-tape density. To preserve the exact knowledge parameter, fill τ inside Q to a subdistribution σ of total mass exactly ρ . This is possible because Q has total mass one. The normalized law σ/ρ retains at least the common mass m with R , so it is within statistical distance $1 - m \leq \delta$ of R . For a genuine finite-coin simulator, append polynomially many unused bits, round down each desired output multiplicity, and then add unused tapes until the good set has exactly $\lceil \rho T \rceil$ elements, where T is the number of extended tapes. The final density is at least ρ . Choosing the extra tape length so the discarded and added conditional mass is at most 2ξ proves (18). Density completion and rounding only establish the existence of the good-tape set; recognizing it efficiently is unnecessary. The extended simulator ignores its extra bits, so the finite-coin loss affects only statistical error and preserves the exact knowledge index. For negligible δ and $\varepsilon = O(\log n)$, the bound also recovers $\text{MZK}[O(\log n)] \subseteq \text{SKC}_{\text{frac}}[O(\log n)] \subseteq \text{AM} \cap \text{coAM}$ from the classical fraction-knowledge upper bound [34].

Proposition 1 (Informal form of Proposition 3). *The following (protocol-preserving) implications hold.*

1. *If a simulator has a good-tape event of density at least 2^{-k} and its conditional output is within statistical distance ν of the real view, then*

$$R \preceq_{k \ln 2, \nu} Q. \quad (21)$$

2. *If $R \preceq_{\kappa, \delta} Q$ with polynomially bounded $\kappa \geq 0$, then for every inverse-exponential negligible $\xi = 2^{-p(n)}$, appending polynomially many unused random bits gives the simulator a good-tape event of density at least*

$$e^{-\kappa} \quad (22)$$

whose conditional output is within $\delta + 2\xi$ of R .

For the other implication, expand the simulator as a mixture. If $\Pr[G] = \rho \geq 2^{-k}$ and $\text{SD}(Q_G, R) \leq \nu$, then

$$Q = \rho Q_G + (1 - \rho) Q_{\bar{G}} \implies R(E) \leq 2^k Q(E) + \nu \quad (23)$$

for every event E . Together with Theorem 15, this gives the following equality for every nonnegative polynomially bounded k , with the exact knowledge parameter, negligible statistical error, and the same uniform strict-PPT convention on both sides:

$$\text{SKC}_{\text{frac}}[k] = \text{HVoMZK}^{\rightarrow}[k \ln 2]. \quad (24)$$

Fraction knowledge thus characterizes real-to-simulator domination. It does not control excess simulator mass, the obstacle to a symmetric converse examined next.

2.6 Why a Distribution-Free Symmetric Converse Fails

Section 11 identifies two obstacles to a symmetric converse. First consider exact fraction simulation, with good-tape density $\rho \geq 2^{-k}$ and $Q_G = R$. Then

$$Q = \rho R + (1 - \rho) Q_{\bar{G}}, \quad D_{\infty}(R \| Q) \leq \ln \frac{1}{\rho} \leq k \ln 2, \quad (25)$$

and hence $D_{\text{KL}}(R\|Q) \leq k \ln 2$. With $\text{SD}(R, Q_G) \leq \nu < 1$, the support of R need not be contained in that of Q , so the correct statement is smoothed: for

$$\begin{aligned} \mu(z) &= \min\{R(z), Q_G(z)\}, \\ m &= \sum_z \mu(z) \geq 1 - \nu, \quad R' = \mu/m, \end{aligned} \tag{26}$$

one has $\text{SD}(R, R') \leq \nu$ and

$$D_\infty(R'\|Q), D_{\text{KL}}(R'\|Q) \leq k \ln 2 + \ln \frac{1}{1 - \nu}. \tag{27}$$

These remain forward, real-to-simulator estimates.

The first obstruction is that the bad-tape component $Q_{\overline{G}}$ is unrestricted. For $R = \delta_a$ and $Q = \rho\delta_a + (1 - \rho)\delta_b$, with $0 < \rho < 1$, the forward divergence is $D_{\text{KL}}(R\|Q) = \ln(1/\rho)$, whereas $D_{\text{KL}}(Q\|R) = +\infty$. Giving b real probability $0 < \gamma < 1$ makes the reverse divergence finite but does not bound it uniformly: for $R_\gamma = (1 - \gamma)\delta_a + \gamma\delta_b$,

$$D_{\text{KL}}(Q\|R_\gamma) = \rho \ln \frac{\rho}{1 - \gamma} + (1 - \rho) \ln \frac{1 - \rho}{\gamma} \rightarrow \infty \tag{28}$$

as $\gamma \rightarrow 0$. Thus no reverse-KL bound follows from fraction knowledge.

The second obstruction is independent of direction: even bounded forward KL does not imply approximate max-divergence with constant exponent and negligible slack. Fix $K > 0$, take $n \geq 2$, put $\tau = 1/n$, and define

$$\begin{aligned} X(a) &= \tau, & Y(a) &= \tau e^{-K/\tau}, \\ X(b) &= 1 - \tau, & Y(b) &= 1 - \tau e^{-K/\tau}. \end{aligned} \tag{29}$$

An explicit finite-coin variant takes $K = \ln 2$, $\tau = 2^{-\lceil \log_2 n \rceil}$, and $Y(a) = \tau 2^{-1/\tau}$. Both laws then use $O(n + \log n)$ fair bits, with $\tau = \Theta(1/n)$ and the same asymptotic estimates. One has $D_{\text{KL}}(X\|Y) = K + O(\tau)$, but $X \preceq_{\varepsilon, \delta} Y$ applied to $\{a\}$ forces

$$\delta \geq \tau(1 - e^{\varepsilon - K/\tau}). \tag{30}$$

For the original choice $\tau = 1/n$ and constant ε , the right-hand side is eventually at least $1/(2n)$; the finite-coin variant has the same $\Omega(1/n)$ obstruction. Negligible δ instead forces $\varepsilon = \Omega(n)$. Equivalently, every X' within distance $< \tau/2$ of X satisfies $D_\infty(X'\|Y) \geq K/\tau - \ln 2 = \Omega(n)$.

A symmetric converse must therefore use interaction structure beyond these distributional estimates. At the class level, it would resolve a longstanding knowledge-complexity question: using $\text{MZK}[O(1)] = \text{SZK} = \text{SKC}_{\text{frac}}[0]$, an inclusion $\text{SKC}_{\text{frac}}[O(1)] \subseteq \text{MZK}[O(1)]$ is equivalent to $\text{SKC}_{\text{frac}}[O(1)] = \text{SKC}_{\text{frac}}[0]$, a collapse not currently known [34, 37].

2.7 Repetition-Based Symmetrization

Section 12 proves a converse when one uniform, possibly unbounded online prover $\widehat{P}$ generates the simulator's *unconditioned* output distribution on every yes-input: writing $Q_x = S(x)$, we require

$$Q_x = \text{View}_{\widehat{P}, V}(x). \tag{31}$$

Each prover message may depend only on the common input, observed prefix, and private randomness. Support on valid transcripts alone does not ensure this: an offline sampler may correlate early messages with unrevealed verifier randomness.

Proposition 2 (Informal form of Proposition 4). *Suppose (P, V) has completeness error c , soundness error s , and fraction knowledge k through a simulator S with statistical error ν . If one online prover $\hat{P}$ generates Q_x exactly on every yes-input, $k = O(\log n)$, and c, s, ν are negligible, then a polynomial sequential repetition of the protocol is honest-verifier symmetric multiplicative zero knowledge with leakage $O(k + 1)$ and negligible additive error.*

The repeated verifier runs all t independent copies of V to completion and accepts their OR. The prescribed prover privately chooses $I \leftarrow [t]$, uses P in coordinate I , and uses $\hat{P}$ in every other coordinate, with fresh independent randomness in each copy. The full view is the tuple of the copy views, and the simulator outputs $Q^{\otimes t}$. If G is the fraction simulator's good event, then $\rho = \Pr[G] \geq 2^{-k}$ and $\text{SD}(R, Q_G) \leq \nu$. Smoothing as in (26) gives a distribution R' with $\text{SD}(R, R') \leq \nu$ and

$$R'(z) \leq e^\kappa Q(z), \quad \kappa := \ln \frac{1}{\rho m} \leq k \ln 2 + \ln \frac{1}{1 - \nu}. \quad (32)$$

Writing $L(z) = R'(z)/Q(z)$ on the support of Q , one has $0 \leq L \leq e^\kappa$ and $\mathbb{E}_Q[L] = 1$. The smoothed repeated real view is the uniform-coordinate mixture $T'_t = t^{-1} \sum_i T'_{t,i}$, where

$$\begin{aligned} T'_{t,i} &= Q^{\otimes(i-1)} \otimes R' \\ &\quad \otimes Q^{\otimes(t-i)}, \\ \frac{T'_t(z_1, \dots, z_t)}{Q^{\otimes t}(z_1, \dots, z_t)} &= \frac{1}{t} \sum_{i=1}^t L(z_i). \end{aligned} \quad (33)$$

The ratio is at most e^κ , giving the forward inequality. Under $Q^{\otimes t}$ the summands are independent, bounded, and have mean one. Hoeffding's inequality bounds the probability that their average falls below $e^{-\eta}$ by

$$\beta_t(\eta, \kappa) = \exp\left(-\frac{2t(1 - e^{-\eta})^2}{e^{2\kappa}}\right), \quad (34)$$

which supplies the reverse multiplicative inequality. Only one coordinate uses R' in each mixture component, so replacing it by R changes the whole repeated view by at most ν , rather than $t\nu$. Transferring the inequalities therefore costs at most ν in the forward direction and $e^\eta \nu$ in the reverse direction.

Concretely, for sufficiently large n , take

$$\bar{\kappa} = k \ln 2 + \ln 2, \quad \eta = \bar{\kappa} + 1,$$

and any polynomial-time computable, polynomially bounded integer $t \geq n^2 e^{2\bar{\kappa}}$. Then the actual repeated view T_t satisfies

$$\begin{aligned} T_t &\approx_{\eta, \Delta} Q^{\otimes t}, \\ \Delta &= e^\eta \nu + \exp\left(-\frac{2t(1 - e^{-\eta})^2}{e^{2\bar{\kappa}}}\right) \\ &= \text{negl}(n). \end{aligned} \quad (35)$$

Because $k = O(\log n)$, both e^η and the required t are polynomial. The selected genuine coordinate preserves completeness error at most c ; on a no-instance, each marginal copy is a valid single-copy cheating strategy, so a union bound gives soundness at most ts . This gives an honest-verifier converse for the modified product protocol when one online prover generates the simulator's full view distribution; the general symmetric converse remains open.

3 Preliminaries and Notation

3.1 Notation

Basics. The security parameter is the input length $n = |x|$. For a nonnegative integer m , write $[m] = \{1, \dots, m\}$, with $[0] = \emptyset$. All sample spaces are finite. Protocol messages, verifier tapes, and efficient-sampler outputs have length polynomial in n . Uppercase letters denote random variables and lowercase letters their realizations. We write $z \leftarrow Z$ or $z \leftarrow P$ for sampling, identifying a random variable with its distribution when convenient.

For a distribution P on a finite set Ω and an event $E \subseteq \Omega$, write

$$P(E) = \sum_{z \in E} P(z), \quad \text{supp}(P) = \{z \in \Omega : P(z) > 0\}.$$

If $P(E) > 0$, the conditional distribution $P \mid E$ is

$$(P \mid E)(z) = \frac{P(z)\mathbf{1}_E(z)}{P(E)}.$$

Write $P_1 \otimes \dots \otimes P_k$ for an independent product and $P^{\otimes k}$ for k independent copies of P . Let Unif_m denote the uniform distribution on $\{0, 1\}^m$, including the point mass on the empty string when $m = 0$, and $\text{Bern}(p)$ the Bernoulli law with success probability p . For a message sequence $M = (M_1, \dots, M_d)$, write $M_{\leq i} = (M_1, \dots, M_i)$. If Y is a distribution on full message sequences, Y_i denotes the law of the prefix $M_{\leq i}$, and Y_0 is the point mass on the empty prefix.

Logarithms and entropy. We use $\ln$ for natural logarithms and $\log_2$ for base-two logarithms. Unsubscripted $\log$ appears only in asymptotic expressions, where the base is immaterial. Shannon entropy is measured in bits:

$$H(P) := - \sum_{z \in \Omega} P(z) \log_2 P(z), \tag{36}$$

with $0 \log_2 0 = 0$. We also write $H(Z)$ for the entropy of Z . Conditional entropy and the chain rule are

$$\begin{aligned} H(X \mid Y) &:= \sum_y \Pr[Y = y] H(X \mid Y = y), \\ H(X, Y) &= H(Y) + H(X \mid Y). \end{aligned}$$

For independent X, Y , this gives $H(X, Y) = H(X) + H(Y)$. Binary entropy in nats and bits, respectively, is

$$\begin{aligned} h(t) &= -t \ln t - (1 - t) \ln(1 - t), \\ h_2(t) &= -t \log_2 t - (1 - t) \log_2(1 - t). \end{aligned}$$

Relative entropy. For distributions P, Q on a common finite space, define

$$D_{\text{KL}}(P \parallel Q) := \sum_z P(z) \ln \frac{P(z)}{Q(z)}, \tag{37}$$

$$D_{\text{KL},2}(P \parallel Q) := \sum_z P(z) \log_2 \frac{P(z)}{Q(z)} = \frac{D_{\text{KL}}(P \parallel Q)}{\ln 2}. \tag{38}$$

The units are nats and bits, respectively. A summand with $P(z) = 0$ is zero; if $P(z) > 0 = Q(z)$ for some z , both divergences are infinite. Relative entropy is directional and is not a metric. For $a, t \in [0, 1]$, binary relative entropy in bits is

$$d_2(a||t) := a \log_2 \frac{a}{t} + (1-a) \log_2 \frac{1-a}{1-t},$$

with the same zero-mass conventions. Abbreviating $P_{Y|X=x}$ by $P_{Y|x}$, the KL chain rule is

$$D_{\text{KL},2}(P_{XY}||Q_{XY}) = D_{\text{KL},2}(P_X||Q_X) + \mathbb{E}_{x \leftarrow P_X} [D_{\text{KL},2}(P_{Y|x}||Q_{Y|x})], \quad (39)$$

whenever the conditional laws are defined on P_X -positive outcomes. For a deterministic map f , data processing gives

$$D_{\text{KL},2}(P||Q) \geq D_{\text{KL},2}(f(P)||f(Q)), \quad (40)$$

where $f(P)$ is the law of $f(Z)$ for $Z \leftarrow P$.

Asymptotics and protocol views. A function $\nu: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ is negligible if, for every constant $c > 0$, $\nu(n) < n^{-c}$ for all sufficiently large n . We write $\text{negl}(n)$ for an unspecified negligible function, which may change between occurrences. The notation $o(1)$ means convergence to zero, which need not be negligible.

Samplers and protocol machines are uniform unless stated otherwise. Throughout, PPT means polynomial worst-case running time (*strict* PPT), including in the SZK, MZK, and fraction-SKC definitions. Such machines use polynomially many fair random bits; padding unused bits preserves their output law. Provers are uniform algorithms with unrestricted running time and private randomness, inducing behavioral strategies; they may use an unbounded number of random bits but must terminate each message computation almost surely. Soundness holds against every cheating strategy.

For a prover P and verifier V on common input x , $\text{View}_{P,V}(x)$ denotes the verifier's complete view: its random tape and all exchanged messages. Write $R_x = \text{View}_{P,V}(x)$ for the real view and $Q_x = S(x)$ for the simulator output. Let A_x be the set of valid accepting views, with $p_x = R_x(A_x)$ and $q_x = Q_x(A_x)$. We omit subscripts when the input is fixed.

General-verifier notions require auxiliary-input simulation: for every PPT verifier V^* there is one PPT simulator S_{V^*} whose guarantee holds for every yes-input x and every auxiliary string z of polynomially bounded length. Both receive (x, z) and run in time polynomial in $|x| + |z|$; leakage and error profiles are evaluated at $|x|$. For each fixed polynomial bound on $|z|$, negligible error is uniform over all permitted (x, z) . Write $\text{View}_{P,V^*}(x, z)$ for the view. The auxiliary input and verifier output may be included explicitly, since they are determined by the inputs, coins, and transcript. Honest-verifier notions use the prescribed V without auxiliary input.

3.2 Basic Complexity Background

A *promise problem* is a pair $\Pi = (\Pi_Y, \Pi_N)$ of disjoint subsets of $\{0, 1\}^*$, called the yes- and no-instances [2]. Inputs outside $\Pi_Y \cup \Pi_N$ impose no requirements. A language is the case $\Pi_N = \{0, 1\}^* \setminus \Pi_Y$. We formulate the classes below for promise problems.

An *interactive proof system* for Π is a pair (P, V) of an unbounded prover and a PPT verifier on common input x [3, 4, 25, 26]. It has completeness error $c(n)$ and soundness error $s(n)$ if

$$\begin{aligned} x \in \Pi_Y &\implies \Pr[(P, V)(x) \text{ accepts}] \geq 1 - c(|x|), \\ x \in \Pi_N &\implies \sup_{P^*} \Pr[(P^*, V)(x) \text{ accepts}] \leq s(|x|), \end{aligned}$$

where the supremum ranges over all cheating provers P^* . Completeness and soundness errors are negligible unless stated otherwise.

For sequential arguments, we use the standard full-view normalization [20, 34]: the verifier starts the conversation and appends its random tape as a final message. A conversation is valid when every verifier message is consistent with that tape. This reveals no coins before the original protocol ends and gives an equivalent encoding of the view.

3.3 Statistical and Multiplicative Zero-Knowledge

For distributions P, Q on a common finite space, statistical distance is

$$\text{SD}(P, Q) := \frac{1}{2} \sum_z |P(z) - Q(z)| = \max_E |P(E) - Q(E)|.$$

For $\varepsilon \geq 0$ and $0 \leq \delta < 1$, write

$$P \preceq_{\varepsilon, \delta} Q \iff P(E) \leq e^\varepsilon Q(E) + \delta \text{ for every event } E, \quad (41)$$

and $P \approx_{\varepsilon, \delta} Q$ when both $P \preceq_{\varepsilon, \delta} Q$ and $Q \preceq_{\varepsilon, \delta} P$ hold. This symmetric closeness relation is used in PINE [36, Definition 4.1]. For a common randomized channel K , with fresh randomness independent of its input, both statistical distance and the one-sided relation obey postprocessing:

$$\begin{aligned} \text{SD}(K(P), K(Q)) &\leq \text{SD}(P, Q), \\ P \preceq_{\varepsilon, \delta} Q &\implies K(P) \preceq_{\varepsilon, \delta} K(Q). \end{aligned}$$

The second implication follows by averaging the event inequality over the channel's random coins; applying it in both directions preserves symmetric closeness. The forward inequality bounds real mass on events rare under the simulator; the reverse inequality bounds simulator mass on events rare in the interaction. The collapse proof uses the reverse bound after conditioning on acceptance.

Definition 1 (Honest-verifier statistical zero-knowledge). An interactive proof (P, V) for Π is honest-verifier statistical zero-knowledge if a PPT simulator S satisfies

$$\text{SD}(\text{View}_{P,V}(x), S(x)) \leq \nu(|x|) \quad (x \in \Pi_Y)$$

for some negligible ν . The corresponding class is HVSZK.

Definition 2 (Statistical zero-knowledge). A promise problem Π belongs to SZK if it has an interactive proof (P, V) such that, for every PPT verifier V^* , there is one PPT simulator S_{V^*} with the following property: for each fixed polynomial bound p on auxiliary-input length, some negligible ν satisfies

$$\text{SD}(\text{View}_{P,V^*}(x, z), S_{V^*}(x, z)) \leq \nu(|x|)$$

for every yes-input x and every z with $|z| \leq p(|x|)$. The error function may depend on V^* and p , but the simulator is the same for all auxiliary inputs and length bounds.

Goldreich, Sahai, and Vadhan proved $\text{HVSZK} = \text{SZK}$ [22]. Thus an honest-verifier protocol suffices to establish SZK membership.

Multiplicative zero-knowledge uses symmetric approximate max-divergence in place of statistical distance.

Definition 3 (Multiplicative zero-knowledge). An interactive proof (P, V) is honest-verifier (ε, δ) -multiplicative zero-knowledge if a PPT simulator S satisfies, for every yes-instance x ,

$$\text{View}_{P,V}(x) \approx_{\varepsilon(|x|), \delta(|x|)} S(x).$$

The corresponding class is $\text{HVMZK}[\varepsilon, \delta]$. Requiring that, for every PPT verifier V^* , one PPT simulator S_{V^*} satisfies

$$\text{View}_{P,V^*}(x, z) \approx_{\varepsilon(|x|), \delta(|x|)} S_{V^*}(x, z)$$

for every yes-input x and polynomially bounded auxiliary input z gives $\text{MZK}[\varepsilon, \delta]$. An unspecified δ is negligible and may depend on V^* and the auxiliary-length bound, but not on x or z .

An asymptotic leakage index, such as $\text{MZK}[O(f(n))]$, denotes the union over nonnegative profiles $\varepsilon(n) = O(f(n))$, with a fixed implicit constant for each profile.

Definition 4 (Honest-verifier one-sided multiplicative zero-knowledge). Let $\kappa, \delta: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$, with $\delta(n) < 1$, and let (P, V) be an interactive proof for Π . It is *honest-verifier one-sided* (κ, δ) -multiplicative zero-knowledge if a PPT sampler S satisfies, for every $x \in \Pi_Y$ of length n , with $R_x = \text{View}_{P,V}(x)$ and $Q_x = S(x)$,

$$\begin{aligned} R_x &\preceq_{\kappa(n), \delta(n)} Q_x, \\ \text{i.e., } R_x(E) &\leq e^{\kappa(n)} Q_x(E) + \delta(n) \quad \text{for every event } E. \end{aligned} \tag{42}$$

The class $\text{HVomZK}^{\rightarrow}[\kappa, \delta]$ consists of promise problems with such a proof and negligible completeness and soundness errors. Omitting the additive error denotes the union

$$\text{HVomZK}^{\rightarrow}[\kappa] := \bigcup_{\substack{\delta: \mathbb{N} \rightarrow [0,1] \\ \delta \text{ negligible}}} \text{HVomZK}^{\rightarrow}[\kappa, \delta].$$

Remark 1 (Direction of the inequality). *The arrow denotes real-to-simulator domination, with no reverse bound. For $\delta = 0$, Definition 4 is equivalent to*

$$R_x(z) \leq e^{\kappa(n)} Q_x(z) \quad \text{for every view } z.$$

For $\delta > 0$, the event formulation (42) is essential: the additive slack must not be applied separately to each outcome and then summed.

At zero multiplicative leakage,

$$P(E) \leq Q(E) + \delta \quad \forall E \iff \text{SD}(P, Q) \leq \delta,$$

because both one-sided event suprema equal statistical distance:

$$\sup_E (P(E) - Q(E)) = \sup_E (Q(E) - P(E)) = \text{SD}(P, Q).$$

Hence $\text{HVMZK}[0] = \text{HVSZK}$, and honest-verifier/general-verifier equivalence [22, 37] gives

$$\text{HVMZK}[0] = \text{MZK}[0] = \text{SZK}. \tag{43}$$

3.4 Fraction Statistical Knowledge Complexity

Fraction statistical knowledge complexity measures the density of simulator tapes on which simulation is statistically accurate [20, 21]. Write $S(x; r)$ for a PPT simulator using a uniform tape $r \in \{0, 1\}^{m(|x|)}$, where m is polynomially bounded.

Definition 5 (Fraction statistical knowledge complexity). Let $k: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ be polynomially bounded. An interactive proof (P, V) has fraction statistical knowledge complexity at most k if there exist a PPT simulator S and a negligible ν such that, for every $x \in \Pi_Y$, there is a set of good tapes

$$G_x \subseteq \{0, 1\}^{m(|x|)}, \quad \Pr_r[r \in G_x] \geq 2^{-k(|x|)},$$

with

$$\text{SD}(S(x; r) \mid r \in G_x, \text{View}_{P,V}(x)) \leq \nu(|x|).$$

The corresponding class of promise problems is $\text{SKC}_{\text{frac}}[k]$.

The set G_x need not be efficiently recognizable. At $k = 0$, all tapes are good, so $\text{SKC}_{\text{frac}}[0] = \text{HVSZK} = \text{SZK}$. At $k = 1$, at least half are good; at $k = O(\log n)$, an inverse-polynomial fraction suffices. If conditioning on an event of probability ρ yields the exact real-view law, then the unconditioned simulator assigns at least $\rho R(E)$ to every event E . Taking $\rho = 2^{-k}$ gives the multiplicative factor $2^k = e^{k \ln 2}$.

3.5 Arthur–Merlin Complexity Classes

Definition 6 (The class AM). A promise problem Π belongs to AM if there exist a polynomial-time predicate $V(x, r, m)$ and a polynomial p such that, with m ranging over $\{0, 1\}^{p(|x|)}$,

$$\begin{aligned} x \in \Pi_Y &\implies \Pr_{r \leftarrow \{0, 1\}^{p(|x|)}}[\exists m : V(x, r, m) = 1] \geq \frac{2}{3}, \\ x \in \Pi_N &\implies \Pr_{r \leftarrow \{0, 1\}^{p(|x|)}}[\exists m : V(x, r, m) = 1] \leq \frac{1}{3}. \end{aligned}$$

Equivalently, AM consists of promise problems with constant-round public-coin proofs: Arthur sends uniform coins, and Merlin responds after seeing them. The error constants $1/3$ can be amplified.

Definition 7 (The class coAM). For $\Pi = (\Pi_Y, \Pi_N)$, let $\bar{\Pi} = (\Pi_N, \Pi_Y)$ and define $\text{coAM} = \{\Pi : \bar{\Pi} \in \text{AM}\}$. Thus $\text{AM} \cap \text{coAM}$ consists of problems whose yes- and no-sides both admit Arthur–Merlin verification.

Petrack and Tardos [34] proved

$$\text{SKC}_{\text{frac}}[O(\log n)] \subseteq \text{AM} \cap \text{coAM}.$$

Together with Theorem 15, this places logarithmic one-sided multiplicative zero-knowledge in $\text{AM} \cap \text{coAM}$. Our collapse theorem strengthens the bound to SZK for symmetric simulation.

Complete problems for SZK. We use the following standard circuit problems [24, 37].

Definition 8 (Statistical Difference). An instance consists of circuits C_0, C_1 sampling distributions D_0, D_1 . For fixed constants $0 \leq a < b \leq 1$, the promise problem $\text{SD}_{a,b}$ has the promises

$$\text{SD}(D_0, D_1) \leq a \quad (\text{no}), \quad \text{SD}(D_0, D_1) \geq b \quad (\text{yes}).$$

For fixed $0 < a < b < 1$ with $a < b^2$, it is SZK-complete.

Definition 9 (Entropy Difference). An instance consists of circuits sampling distributions X, Y , with promises

$$H(X) \geq H(Y) + 1 \quad (\text{yes}), \quad H(Y) \geq H(X) + 1 \quad (\text{no}).$$

Entropy Difference is SZK-complete.

4 Collapse to Statistical Zero-Knowledge

Theorem 8 (Main collapse theorem). *Suppose that $0 \leq \varepsilon(n) \leq C \ln n$ for a constant C and all sufficiently large n , and that the completeness error, soundness error, and additive simulation error are negligible. Then*

$$\text{HVMZK}[\varepsilon] = \text{SZK}.$$

Consequently,

$$\text{MZK}[\varepsilon] = \text{HVMZK}[\varepsilon] = \text{SZK}.$$

More generally, the upper bound holds for any such protocol whose yes-instance acceptance probability is at least a fixed constant $p_0 > 0$, provided soundness and additive simulation error remain negligible.

4.1 Setup

Only the containment

$$\text{HVMZK}[\varepsilon] \subseteq \text{SZK} \tag{44}$$

requires proof. Indeed, statistical distance at most $\nu(n)$ implies symmetric $(0, \nu(n))$ approximate-max-divergence [12]. Hence $\text{SZK} = \text{MZK}[0] \subseteq \text{MZK}[\varepsilon] \subseteq \text{HVMZK}[\varepsilon]$.

Fix a promise problem $\Pi = (\Pi_Y, \Pi_N)$ witnessed by an interactive proof (P, V) and a strict-PPT simulator S . We prove the upper bound under the weaker completeness requirement $p_x \geq p_0 > 0$ on yes-inputs. The proof reduces Π to Statistical Difference and Entropy Difference, both SZK-complete [24, 37], and uses closure of SZK under NC^1 truth-table reductions [37, Corollary 4.14].

Fix one global cutoff $n_0 \geq 4$, large enough for every eventual inequality below to hold uniformly over inputs. All constructions below concern $n \geq n_0$. For shorter inputs, a fixed finite table maps promised yes- and no-inputs to fixed instances with the corresponding answer; entries outside the source promise are arbitrary. Hardwiring this table gives a single uniform reduction at all lengths.

4.2 Views, Validity, and Fixed-Length Padding

Fix a public input x of length n . The verifier's full view is

$$Z = (M_1, \dots, M_{d-1}, \rho),$$

where $M_1, \dots, M_{d-1}$ are *all* messages exchanged in both directions and $\rho \in \{0, 1\}^{r(n)}$ is the verifier's complete random tape. Normalize the protocol so that the verifier sends the first message and

appends ρ after the original interaction ends. This records the full view without revealing coins during the original interaction.

Each message position has a polynomially bounded length depending only on n . Add length information and fixed-width dummy slots to normalize variable-length or early-terminating protocols. Dummy verifier messages are deterministic; unused prover slots contain arbitrary bits that the verifier ignores. All bit patterns are allowed in a prover slot, although the verifier rejects malformed payloads. These conventions give every sampling circuit a fixed output length.

A string z is a *valid view* if replaying V on (x, ρ) and the recorded prior messages reproduces every verifier message. Validity thus permits malformed prover payloads that lead to rejection. Let A_x be the set of valid accepting views; replaying V decides membership in A_x in polynomial time.

Write

$$\begin{aligned} R_x &= \text{View}_{P,V}(x), & Q_x &= S(x), \\ p_x &= R_x(A_x), & q_x &= Q_x(A_x). \end{aligned}$$

The machines and distributions are defined on every input, but the simulation guarantee

$$R_x \approx_{\varepsilon(n), \delta(n)} Q_x \tag{45}$$

is required only for $x \in \Pi_Y$. We use completeness only on yes-inputs and soundness only on no-inputs.

We also need an efficiently samplable genuine view with full support. Let P_0 independently send a uniform string of the prescribed padded length at every prover position, and let

$$R_{0,x} = \text{View}_{P_0,V}(x).$$

Because the verifier follows its prescribed algorithm with a uniform tape, $R_{0,x}$ is supported on valid views. Let $\mathcal{I}_P$ be the set of prover positions and $\ell_i(n)$ the padded length at position i . Writing

$$\Lambda(n) = r(n) + \sum_{i \in \mathcal{I}_P} \ell_i(n),$$

every valid view z has

$$R_{0,x}(z) = 2^{-\Lambda(n)}. \tag{46}$$

The uniform choices include ignored prover slots. Later, to equalize sampling-circuit output lengths, we append deterministic zeros, which do not change entropy.

4.3 Bounded Conditioning on Yes-Instances

Fix $x \in \Pi_Y$ and suppress the subscript x . Since $p \geq p_0$, the real-to-simulator half of (45), applied to $A = A_x$, gives

$$p \leq e^\varepsilon q + \delta,$$

and hence

$$q \geq e^{-\varepsilon}(p - \delta). \tag{47}$$

For all sufficiently large n , $p - \delta \geq p_0/2$. Since $\varepsilon(n) \leq C \ln n$,

$$q \geq \frac{p_0}{2} n^{-C}. \tag{48}$$

Choose the integer $D = \lceil C \rceil + 3$ and set $N(n) = n^D$. The bounded sampler draws $Z_1, \dots, Z_N$ independently from Q_x and sets

$$B = \mathbf{1}\{Z_j \in A_x \text{ for at least one } j\}. \quad (49)$$

The *success bit* B records whether any trial produced a valid accepting view. On success, output the first such Z_j ; on failure, output a fresh sample from $R_{0,x}$. Write X_x for the output law and $b_x = \Pr[B = 1]$. This fallback ensures valid support and a uniform strict-PPT sampler on every input, without requiring a known rejecting view. The sampler only evaluates V and S and uses the fixed trial bound n^D . The probabilities p_x, q_x, b_x and the profiles ε, δ are used in the analysis, not computed by the sampler.

Independence gives

$$1 - b_x = \Pr[B = 0] = (1 - q_x)^{N(n)} \leq e^{-N(n)q_x}. \quad (50)$$

On a yes-instance, (48) therefore implies

$$\Pr[B = 0] \leq \exp\left(-\frac{p_0}{2} n^{D-C}\right) \leq \exp(-p_0 n^3/2), \quad (51)$$

which is negligible.

Let $Q_A = Q \mid A$. Conditioned on $B = 1$, the first successful sample has distribution exactly Q_A . Indeed, for $z \in A$,

$$\begin{aligned} \Pr[X = z, B = 1] &= \sum_{j=1}^N (1 - q)^{j-1} Q(z) \\ &= \frac{1 - (1 - q)^N}{q} Q(z) = b Q_A(z). \end{aligned}$$

Consequently

$$\begin{aligned} X &= bQ_A + (1 - b)R_0, \\ \text{SD}(X, Q_A) &\leq 1 - b = \text{negl}(n) \quad (x \in \Pi_Y). \end{aligned} \quad (52)$$

The simulator-to-real half of (45) is used after conditioning. For every event E ,

$$\begin{aligned} Q_A(E) &= \frac{Q(E \cap A)}{q} \\ &\leq \frac{e^\varepsilon R(E \cap A) + \delta}{q} \\ &\leq \frac{e^{2\varepsilon}}{p - \delta} R(E) + \frac{e^\varepsilon \delta}{p - \delta}. \end{aligned} \quad (53)$$

The forward inequality supplies the inverse-polynomial lower bound on q ; the reverse inequality supplies (53). With negligible completeness error, the reverse inequality also gives the stronger direct estimate $Q(A^c) \leq e^\varepsilon c + \delta = \text{negl}(n)$ at logarithmic leakage. Both directions are used here to cover constant positive completeness; they need not both be necessary in the negligible-completeness-error case.

4.4 Support Regularization and the Yes-Instance KL Bound

Set $\gamma = 2^{-n}$ and define, on every input,

$$Y_x = (1 - \gamma)X_x + \gamma R_{0,x}. \quad (54)$$

On a yes-instance also define

$$W_x = (1 - \gamma)R_x + \gamma R_{0,x}. \quad (55)$$

The distribution W_x is a genuine view: the prover privately chooses P with probability $1 - \gamma$ and P_0 with probability γ , then follows that strategy throughout.

Let

$$\kappa = \frac{e^{2\varepsilon}}{p - \delta}, \quad \alpha = \ln \kappa = 2\varepsilon + \ln \frac{1}{p - \delta},$$

and let

$$\eta = (1 - b) + \frac{e^\varepsilon \delta}{p - \delta}.$$

Equations (52) and (53) give $X(E) \leq \kappa R(E) + \eta$ for every E . Since $\kappa \geq 1$,

$$\begin{aligned} Y(E) &\leq (1 - \gamma)\kappa R(E) + \gamma R_0(E) + \eta \\ &\leq \kappa W(E) + \eta. \end{aligned}$$

Thus

$$Y \preceq_{\alpha, \eta} W. \quad (56)$$

Moreover, (46) implies

$$W(z) \geq \gamma R_0(z) = 2^{-(n + \Lambda(n))} \quad (z \in \text{supp}(Y)). \quad (57)$$

The following estimate bounds relative entropy, measured in nats.

Lemma 1 (Approximate multiplicative simulation to relative entropy). *Let $\alpha \geq 0$ and $U \preceq_{\alpha, \eta} W$, suppose $\min_{z \in \text{supp}(U)} W(z) \geq \mu > 0$, and let $0 \leq \eta \leq 1/2$. Then*

$$D_{\text{KL}}(U \| W) \leq \alpha + h(\eta) + \eta \ln \frac{1}{\mu}, \quad (58)$$

where $h(t) = -t \ln t - (1 - t) \ln(1 - t)$.

Proof. The event formulation of $U \preceq_{\alpha, \eta} W$ is equivalent to

$$\sum_z (U(z) - e^\alpha W(z))_+ \leq \eta.$$

Let $Z(z) = \min\{U(z), e^\alpha W(z)\}$ and $\tau = 1 - \sum_z Z(z) \leq \eta$. If $\tau > 0$, write

$$U = (1 - \tau)U_0 + \tau U_1, \quad U_0 = \frac{Z}{1 - \tau}, \quad U_1 = \frac{U - Z}{\tau};$$

the case $\tau = 0$ is the same with the second component omitted. Convexity of relative entropy in its first argument gives

$$D_{\text{KL}}(U \| W) \leq (1 - \tau)D_{\text{KL}}(U_0 \| W) + \tau D_{\text{KL}}(U_1 \| W).$$

Because $U_0(z) \leq e^\alpha W(z)/(1 - \tau)$,

$$D_{\text{KL}}(U_0 \| W) \leq \alpha + \ln \frac{1}{1 - \tau}.$$

The minimum-mass hypothesis gives $D_{\text{KL}}(U_1 \| W) \leq \ln(1/\mu)$. Hence

$$D_{\text{KL}}(U \| W) \leq \alpha + (1 - \tau) \ln \frac{1}{1 - \tau} + \tau \ln \frac{1}{\mu}.$$

Since $(1 - \tau) \ln(1/(1 - \tau)) \leq h(\tau)$ and h is increasing on $[0, 1/2]$, the result follows from $\tau \leq \eta \leq 1/2$. $\square$

On yes-instances, $e^\varepsilon \leq n^C$, while δ and $1 - b$ are negligible and $p - \delta \geq p_0/2$ eventually. Thus η is negligible. Also

$$\ln \frac{1}{p - \delta} \leq \ln \frac{1}{p_0} + \text{negl}(n),$$

and $\Lambda(n)$ is polynomially bounded. Applying Lemma 1 with $\mu = 2^{-(n + \Lambda(n))}$ yields

$$D_{\text{KL}}(Y_x \| W_x) \leq 2\varepsilon(n) + \ln \frac{1}{p_0} + \text{negl}(n) \quad (x \in \Pi_Y). \quad (59)$$

The estimate $h(t) = O(\sqrt{t})$ as $t \rightarrow 0$ shows that $h(\eta)$ is negligible; so is the product $\eta(n + \Lambda(n))$. If completeness error c is negligible, so is $\ln(1/(p - \delta))$, giving the sharper bound $2\varepsilon(n) + \text{negl}(n)$. Constant positive completeness costs only an additive constant.

4.5 Projection onto a Genuine Prover

An offline sampler may correlate early prover messages with verifier coins revealed only later. Its output need not be a genuine prover view, so soundness cannot be applied directly to Y_x . We instead construct a prover whose messages depend only on the public history.

Let Y be any distribution supported on valid views, and let Y_i be the distribution of the prefix $(M_1, \dots, M_i)$, where the final position contains the verifier tape. Define the simulation-based prover P_Y as follows. At a prover position $i \in \mathcal{I}_P$ and a public history h with $Y_{i-1}(h) > 0$, it sends m with probability

$$\pi_i^Y(m \mid h) = \frac{Y_i(h, m)}{Y_{i-1}(h)}. \quad (60)$$

At a Y -null history, it sends a uniform string of the prescribed padded length, which does not affect expectations under Y .

For our ensemble Y_x , one uniform unbounded algorithm implements these kernels. Given (x, h) , enumerate all tapes of the strict-PPT sampler for Y_x and count outputs extending h and (h, m) . The resulting rational probabilities in (60) can be sampled exactly with fair coins by rejection from a finite integer range. This algorithm terminates almost surely; its running time is unrestricted. Soundness also applies directly to its behavioral kernels, since at each fixed input a deterministic strategy attains at least their acceptance probability. Set

$$T_Y = \text{View}_{P_Y, V}(x).$$

Every $y \in \text{supp}(Y)$ has positive probability under this interaction: its verifier tape has positive probability, and each of its prover transitions has positive conditional probability under Y . Thus $D_{\text{KL}}(Y \| T_Y)$ is finite.

Lemma 2 (Simulation-based projection). *Let W be the view of any genuine prover interacting with V on x . If $W(z) > 0$ for every $z \in \text{supp}(Y)$, then*

$$D_{\text{KL}}(Y \| T_Y) \leq D_{\text{KL}}(Y \| W). \quad (61)$$

Proof. For the prover inducing W , write $\pi_i^W(m | h)$ for its next-message distribution given public history h . A valid full view y fixes the verifier tape ρ and all messages. Since ρ and prior messages determine each verifier message,

$$T_Y(y) = 2^{-r(n)} \prod_{i \in \mathcal{I}_P} \pi_i^Y(y_i | y_{<i}), \quad (62)$$

$$W(y) = 2^{-r(n)} \prod_{i \in \mathcal{I}_P} \pi_i^W(y_i | y_{<i}). \quad (63)$$

Private verifier coins contribute the common factor $2^{-r(n)}$; validity enforces consistency with those coins.

Using (60)–(63),

$$\begin{aligned} & D_{\text{KL}}(Y \| W) - D_{\text{KL}}(Y \| T_Y) \\ &= \mathbb{E}_{y \leftarrow Y} \left[\ln \frac{T_Y(y)}{W(y)} \right] \\ &= \sum_{i \in \mathcal{I}_P} \mathbb{E}_{h \leftarrow Y_{i-1}} [D_{\text{KL}}(\pi_i^Y(\cdot | h) \| \pi_i^W(\cdot | h))] \\ &\geq 0. \end{aligned}$$

The support hypothesis ensures that every denominator used on a Y -positive transition is nonzero. This proves (61). $\square$

If a genuine view W vanishes on a Y -positive outcome, then $D_{\text{KL}}(Y \| W) = \infty$ and the same comparison is immediate. For our efficiently samplable ensemble, one uniform unbounded prover generates T_{Y_x} on every input. Its view attains the minimum even when competing views have smaller support.

Apply the lemma with the efficiently samplable Y_x from (54) and, on a yes-instance, the genuine view W_x from (55). By (57), the support hypothesis holds. Hence

$$D_{\text{KL}}(Y_x \| T_{Y_x}) \leq 2\varepsilon(n) + \ln \frac{1}{p_0} + \text{negl}(n) \quad (x \in \Pi_Y). \quad (64)$$

Writing $D_{\text{KL},2} = D_{\text{KL}} / \ln 2$, we obtain

$$D_{\text{KL},2}(Y_x \| T_{Y_x}) \leq 2C \log_2 n + O(1) \quad (x \in \Pi_Y). \quad (65)$$

The reduction uses P_{Y_x} only to invoke soundness in the analysis; it never runs this prover.

4.6 No-Instances: Low and High Sampler Success

The uniform circuits for Q_x , B , X_x , and Y_x run on every input. On $x \in \Pi_N$, only soundness is available. We divide these inputs into two cases:

$$b_x \leq \frac{3}{4} \quad \text{or} \quad b_x > \frac{3}{4}.$$

The reduction constructs both queries without evaluating b_x ; the case split is used only to verify their promises. Low success is detected through the efficiently sampled bit B . High success forces a relative-entropy gap because no genuine prover can produce that much accepting mass.

4.6.1 Low success and five independent bits

Take five independent executions of the bounded sampler, with success bits $B_1, \dots, B_5$, and output

$$C_x = B_1 \wedge \dots \wedge B_5.$$

Compare the distribution of C_x with the deterministic bit 1. Since $\Pr[C_x = 1] = b_x^5$,

$$\text{SD}(C_x, 1) = \Pr[C_x = 0] = 1 - b_x^5. \quad (66)$$

On yes-instances, $1 - b_x$ is negligible by (51), so $1 - b_x^5 \leq 5(1 - b_x) = \text{negl}(n)$. On a low-success no-instance,

$$\text{SD}(C_x, 1) \geq 1 - \left(\frac{3}{4}\right)^5 = \frac{781}{1024} > \frac{2}{3}. \quad (67)$$

Thus the query is promised close on yes-inputs and far on low-success no-inputs. Four copies would also exceed the $2/3$ threshold; five give a larger constant margin at constant cost.

4.6.2 High success, soundness, and data processing

Suppose now that $x \in \Pi_N$ and $b_x > 3/4$. Whenever $B = 1$, X_x is accepting. Therefore, with

$$a_x = Y_x(A_x), \quad t_x = T_{Y_x}(A_x),$$

we have

$$a_x \geq (1 - \gamma)b_x > \frac{3}{4}(1 - 2^{-n}). \quad (68)$$

The simulation-based prover P_{Y_x} is a genuine, possibly unbounded prover. Soundness therefore gives

$$t_x \leq s(n). \quad (69)$$

Apply $z \mapsto \mathbf{1}_{A_x}(z)$ to Y_x and T_{Y_x} . Data processing, or the log-sum inequality grouping accepting and rejecting outcomes, gives

$$D_{\text{KL},2}(Y_x \| T_{Y_x}) \geq d_2(a_x \| t_x), \quad (70)$$

where

$$d_2(a \| t) = a \log_2 \frac{a}{t} + (1 - a) \log_2 \frac{1 - a}{1 - t}$$

is binary relative entropy. Since $a_x > 0$ and T_{Y_x} is positive on the support of Y_x , we have $t_x > 0$. In particular, the high-success case cannot occur at a length with $s(n) = 0$. For the remaining inputs,

$$\begin{aligned} d_2(a_x \| t_x) &= a_x \log_2 \frac{1}{t_x} - h_2(a_x) + (1 - a_x) \log_2 \frac{1}{1 - t_x} \\ &\geq a_x \log_2 \frac{1}{t_x} - 1 \\ &\geq \frac{3}{4}(1 - 2^{-n}) \log_2 \frac{1}{s(n)} - 1. \end{aligned} \quad (71)$$

Because $s(n)$ is negligible, $\log_2(1/s(n)) = \omega(\log n)$: for every constant k , eventually $s(n) < n^{-k}$. Combining (70) and (71) gives

$$D_{\text{KL},2}(Y_x \| T_{Y_x}) = \omega(\log n) \quad (x \in \Pi_N, b_x > 3/4). \quad (72)$$

4.7 Reductions to Entropy Difference and Statistical Difference

The Statistical Difference query is $(C_x, 1)$ from (66), which is close on yes-inputs and far on low-success no-inputs. We now construct the entropy query for high-success no-inputs.

For the entropy query, let $Y = Y_x$, let Y_i be its first- i -message prefix distribution, and let d be the final position, which contains the verifier tape. From (62), every $y \in \text{supp}(Y)$ satisfies

$$T_Y(y) = 2^{-r(n)} \prod_{i \in \mathcal{I}_P} \frac{Y_i(y_{\leq i})}{Y_{i-1}(y_{< i})}.$$

Taking the Y -expectation of $\log_2(Y(y)/T_Y(y))$ gives the exact identity

$$D_{\text{KL},2}(Y \| T_Y) = r(n) - H(Y_d) + \sum_{i \in \mathcal{I}_P} (H(Y_i) - H(Y_{i-1})). \quad (73)$$

This is the entropy form of the simulation-based-prover method [34].

Each Y_i is efficiently samplable: run the strict-PPT sampler for Y and truncate its output after message i . Using independent randomness for each factor, define

$$U_x = \text{Unif}_{r(n)} \otimes \bigotimes_{i \in \mathcal{I}_P} Y_i, \quad (74)$$

$$W_x = Y_d \otimes \bigotimes_{i \in \mathcal{I}_P} Y_{i-1}. \quad (75)$$

Polynomially many rounds and polynomial output lengths give polynomial-size sampling circuits. A deterministic polynomial-time compiler constructs them from x by unrolling the fixed running-time bounds of V and S , the n^D trials, validity tests, and prefix projections, with randomness as explicit circuit input. Neither P nor P_{Y_x} is needed. Appending deterministic zeros equalizes output lengths without changing entropy. Additivity for independent products and (73) give

$$H(U_x) - H(W_x) = D_{\text{KL},2}(Y_x \| T_{Y_x}). \quad (76)$$

Choose a fixed integer $K > 2C$. Equation (65) gives

$$H(U_x) - H(W_x) \leq K \log_2 n \quad (x \in \Pi_Y) \quad (77)$$

for all sufficiently large n . By (72), for every high-success no-instance and all sufficiently large n ,

$$H(U_x) - H(W_x) > (K + 3) \log_2 n. \quad (78)$$

Let

$$k(n) = \lceil (K + 2) \log_2 n \rceil, \quad W_x^+ = W_x \otimes \text{Unif}_{k(n)}.$$

Compute $k(n)$ exactly as the least integer k for which $2^k \geq n^{K+2}$. Again, deterministically zero-pad the two final circuit outputs to a common length if necessary. On a yes-instance,

$$\begin{aligned} H(W_x^+) - H(U_x) &= k(n) - (H(U_x) - H(W_x)) \\ &\geq 2 \log_2 n \geq 1. \end{aligned}$$

On a high-success no-instance, for all sufficiently large n ,

$$\begin{aligned} H(U_x) - H(W_x^+) &= H(U_x) - H(W_x) - k(n) \\ &> \log_2 n - 1 \geq 1. \end{aligned}$$

Thus the Entropy Difference query (W_x^+, U_x) is a yes-instance on every original yes-input and a no-instance on every high-success original no-input. The finitely many exceptional input lengths are covered by the single cutoff and finite table fixed in Section 4.1.

4.8 Combining the Queries by an NC^1 Truth-Table Reduction

Let z_{SD} be the Statistical Difference answer, with 1 for far and 0 for close. Let z_{ED} be the Entropy Difference answer, with 1 when the first distribution's entropy exceeds the second's by at least one bit. Combine them by

$$g(z_{\text{SD}}, z_{\text{ED}}) = (\neg z_{\text{SD}}) \wedge z_{\text{ED}}. \quad (79)$$

It has constant depth, and hence is an NC^1 function.

We use the three-valued promise semantics of Sahai–Vadhan [37, Definitions 4.6 and 4.12–4.13]. A query has value 1 on the yes-side, 0 on the no-side, and ? outside the promise. Gates are evaluated recursively: $\neg ? = ?$, $0 \wedge a = 0$, $1 \wedge 1 = 1$, and all remaining AND cases involving ? evaluate to ? (with the analogous OR rules). The reduction must evaluate to the correct unambiguous value on every promised source input. This is stronger than correctness under every Boolean completion of the unknown inputs: $? \vee \neg ? = ?$ in these semantics, although every Boolean completion gives 1.

For $x \in \Pi_Y$, both queries are promised: $z_{\text{SD}} = 0$ because (66) is negligible, and $z_{\text{ED}} = 1$ by (77). Hence $g = 1$. If $x \in \Pi_N$ and $b_x \leq 3/4$, the Statistical Difference query is promised far by (67), so $z_{\text{SD}} = 1$ and $g = 0$ even if the entropy-query value is ?. If $x \in \Pi_N$ and $b_x > 3/4$, the Entropy Difference query is promised negative by (78), so $z_{\text{ED}} = 0$ and $g = 0$ even if the statistical-difference value is ?. Thus the actual gates satisfy the required three-valued semantics.

To use a single oracle promise problem, compose each query with a Karp reduction to a fixed SZK-complete problem, such as Statistical Difference. This preserves promised answers; arbitrary images of unpromised queries are tolerated by (79). Closure under NC^1 truth-table reductions therefore proves $\text{HVMZK}[\varepsilon] \subseteq \text{SZK}$.

Finally,

$$\text{SZK} = \text{MZK}[0] \subseteq \text{MZK}[\varepsilon] \subseteq \text{HVMZK}[\varepsilon] \subseteq \text{SZK},$$

which proves Theorem 8.

Remark 2 (Polarizing the minimum KL divergence to a prover view). *Consider verifier and sampler circuits V, S whose random-input bits, message slots, and output bits are explicitly listed in their encoding. Let Q be the output law of S , supported on valid full views. Let $\mathcal{R}(V)$ consist of views generated by interacting with V : each prover message may depend only on the input, private randomness, and transcript available when it is sent. Lemma 2 and (73) give*

$$\begin{aligned} d_V(Q) &:= \min_{R \in \mathcal{R}(V)} D_{\text{KL},2}(Q \| R) = D_{\text{KL},2}(Q \| T_Q) \\ &= r - H(Q) + \sum_{i \in \mathcal{I}_P} (H(Q_i) - H(Q_{i-1})), \end{aligned} \quad (80)$$

Here r is the verifier's coin length and Q_i are transcript prefixes. The simulation-based prover attains the minimum by sampling each next message from Q conditioned on the observed transcript prefix. These conditional distributions are exactly samplable by a uniform unbounded prover.

The products in (74)–(75) give samplers U, W with $H(U) - H(W) = d_V(Q)$. Include rational thresholds $0 \leq a < b$ in binary in the input, and let N be the total encoding length of V, S, a, b . Fix polynomial bounds on b and $1/(b - a)$ as part of the promise. Malformed encodings and thresholds violating $0 \leq a < b$ or the polynomial bounds are mapped to a fixed circuit pair, so the reduction has polynomial running time on every input. Take $m = \lceil 4/(b - a) \rceil$ independent copies and add $\lfloor m(a + b)/2 \rfloor$ fair bits to the W side. The promises $d_V(Q) \leq a$ and $d_V(Q) \geq b$ then give opposite entropy gaps of at least one bit. Deterministic zero-padding equalizes output lengths. Entropy

completeness and standard polarization [24, 37] uniformly produce, in time polynomial in $N + k$, circuits D_0, D_1 with

$$\begin{aligned} d_V(Q) \leq a &\implies \text{SD}(D_0, D_1) \leq 2^{-k}, \\ d_V(Q) \geq b &\implies \text{SD}(D_0, D_1) \geq 1 - 2^{-k}. \end{aligned}$$

This direct circuit reduction does not use IMDA membership. It applies standard statistical-distance polarization after the entropy reduction; the statistical-distance threshold remains tight for oblivious transformations [28].

For the Jensen–Shannon special case, let V receive Y before revealing a fresh fair bit B , and let Q sample B first, then $Y \sim C_B$. Ignoring deterministic view records, every genuine view is $\nu \otimes \text{Unif}_1$. With $M = (C_0 + C_1)/2$,

$$D_{\text{KL},2}(Q\|\nu \otimes \text{Unif}_1) = I_Q(Y; B) + D_{\text{KL},2}(M\|\nu).$$

Thus $d_V(Q) = \text{JS}(C_0, C_1)$, attained at $\nu = M$. For fixed rational $0 < a < b < 1$, SZK-completeness follows from the Jensen–Shannon characterization [9] and closure under complement.

5 A uniform-branch simulation lemma

A uniform component turns a bound on prover communication into a multiplicative simulation bound.

Lemma 3. *Let $\Pi = (P, V)$ be an interactive proof with a fixed, polynomial-time computable, polynomially bounded interface and at most $m(n)$ prover bits. The prescribed prover enforces this interface against cheating verifiers. Early termination follows a deterministic public rule shared with the uniform strategy; any other abort is encoded in a fixed-width prover slot, counted in $m(n)$ and sampled by the uniform strategy. Both strategies thus terminate on the same invalid verifier messages. Let $h(n)$ be a nonnegative integer-valued, polynomial-time computable, polynomially bounded function with $h(n) = \omega(\log n)$, and let $\theta(n) = 2^{-h(n)}$. There is a modified prescribed prover P_θ whose view R_θ^* against every PPT verifier V^* has a strict-PPT simulator U^* satisfying*

$$R_\theta^* \approx_{\varepsilon_\theta, 0} U^*, \quad \varepsilon_\theta(n) = \ln 2 \cdot \max\{m(n), h(n)\}. \quad (81)$$

The verifier and its soundness error are unchanged, while the completeness error increases by at most $\theta(n)$.

Proof. The simulator U^* runs V^* , supplying independent uniform strings in the prover slots under the shared interface and abort rules. Before interacting, P_θ privately chooses to run P with probability $1 - \theta$ or this uniform strategy with probability θ .

Fix an input of length n and any auxiliary input, and write $m = m(n)$. Pad the verifier tape to its polynomial bound q^* . For each completed view z consistent with V^* and the shared interface rules, let $c(z) \leq m$ be the number of prover bits sent, counting an encoded abort’s full slot. The shared public termination rule determines $c(z)$ from the view. Under the uniform strategy,

$$U^*(z) = 2^{-q^* - c(z)}. \quad (82)$$

For the verifier tape in z , the probability that P produces its prover messages is at most one. Thus the original real view satisfies

$$R^*(z) \leq 2^{-q^*} = 2^{c(z)} U^*(z) \leq 2^m U^*(z). \quad (83)$$

This bound includes early termination by the verifier or the shared interface rule. Views inconsistent with V^* or these rules have zero mass under both strategies.

The modified view is

$$R_\theta^* = (1 - \theta)R^* + \theta U^*. \quad (84)$$

Pointwise,

$$R_\theta^*(z) \leq ((1 - \theta)2^m + \theta)U^*(z) \leq 2^m U^*(z), \quad (85)$$

$$U^*(z) \leq \theta^{-1} R_\theta^*(z) = 2^h R_\theta^*(z). \quad (86)$$

These bounds prove (81). The uniform branch reduces completeness by at most its selection probability θ . The unchanged verifier retains its soundness against arbitrary provers. $\square$

Laconic provers and limited nondeterminism. Protocols with limited prover communication are studied as proofs with *laconic provers* [17, 23]. Let $\text{IP}_{P \rightarrow V}[b]$ denote the promise problems having interactive proofs with negligible completeness and soundness errors and a fixed message interface containing at most $b(n)$ prover bits in total. This definition includes all interface hypotheses of Lemma 3: enforcement against cheating verifiers, shared public invalid-message and termination rules, and abort responses counted in $b(n)$. Let $\text{GC}(b, P)$ denote languages verified in polynomial time using at most $b(n)$ nondeterministically guessed bits.

Corollary 3 (From prover communication to multiplicative leakage). *Let b be a nonnegative integer-valued, polynomial-time computable, polynomially bounded function satisfying $b(n) = \omega(\log n)$. Then*

$$\text{IP}_{P \rightarrow V}[b] \subseteq \text{MZK}[b \ln 2, 0], \quad \text{GC}(b, P) \subseteq \text{MZK}[b \ln 2, 0].$$

The second inclusion has a one-message protocol with perfect soundness and completeness error at most $2^{-b(n)}$.

Proof. Apply Lemma 3 with $h = b$ and $m \leq b$, giving leakage $b \ln 2$, unchanged soundness, and negligible additional completeness error 2^{-b} . For the second inclusion, the single message supplies $b(n)$ guessed bits to a deterministic checker; unused bits are ignored. Before mixing, completeness and soundness are perfect; an unbounded prescribed prover finds a witness by exhaustive search on yes-inputs. $\square$

The budget refers to the original negligible-error proof. Amplifying a constant-error laconic proof may increase it and is not implicit here.

6 A Conditional Unrelativized Separation

We use natural logarithms in multiplicative parameters. For distributions P and Q on the same finite set, write $P \approx_{\varepsilon, 0} Q$ when, for every event E ,

$$P(E) \leq e^\varepsilon Q(E) \quad \text{and} \quad Q(E) \leq e^\varepsilon P(E). \quad (87)$$

Pointwise likelihood-ratio bounds in both directions imply (87) by summation. Every verifier's view contains its random tape, transcript, and output.

6.1 The randomized-witness protocol for SAT

Let φ be a Boolean formula whose encoding has length N , and let $m \leq N$ be its number of variables. On a satisfiable formula, let $w \in \{0, 1\}^m$ be the lexicographically first satisfying assignment, which the unbounded prover can find. Let $U_m \leftarrow \text{Unif}(\{0, 1\}^m)$ and set $\eta = 2^{-N}$. The prover sends the single message

$$W = \begin{cases} w, & \text{with probability } 1 - \eta, \\ U_m, & \text{with probability } \eta, \end{cases} \quad (88)$$

and the verifier accepts exactly when the received string satisfies φ . It rejects messages that are not m -bit assignments. On an unsatisfiable formula the prover sends a fixed m -bit string; malformed formula encodings are rejected without requesting a message. The prover sends at most one message and ignores unsolicited requests. The simulator samples U_m in place of W and runs the verifier to obtain its complete view.

Theorem 9 (Unrelativized conditional example). *The protocol in (88) has completeness error at most 2^{-N} , perfect soundness, and multiplicative simulation parameter $(\varepsilon, \delta) = (N \ln 2, 0)$ against every PPT cheating verifier. Consequently,*

$$\text{SAT} \in \text{MZK}[N \ln 2]. \quad (89)$$

If PH does not collapse to its second level, then

$$\text{SAT} \in \text{MZK}[N \ln 2] \setminus \text{SZK}. \quad (90)$$

Proof. Apply Lemma 3 to the deterministic witness prover. It sends $m \leq N$ bits with perfect completeness and soundness. Taking $h(N) = N$ gives (88), with leakage $\ln 2 \max\{m, N\} = N \ln 2$, zero additive simulation error, and completeness error at most 2^{-N} . For each cheating verifier and auxiliary input, the simulator samples U_m , runs the verifier on that message, and outputs its complete view. The lemma gives both multiplicative inequalities uniformly over auxiliary inputs; the explicit witness and nonwitness likelihood ratios appear in Section 2.2.

If $\text{SAT} \in \text{SZK}$, its NP-completeness and closure of SZK under polynomial-time many-one reductions give $\text{NP} \subseteq \text{SZK}$. The containment $\text{SZK} \subseteq \text{AM} \cap \text{coAM}$ [34] then gives $\text{NP} \subseteq \text{coAM}$, equivalently $\text{coNP} \subseteq \text{AM}$ [1, 37]. Boppana, Håstad, and Zachos show that this collapses PH to its second level [10]. The contrapositive proves (90). $\square$

An unconditional proof that $\text{SAT} \notin \text{SZK}$ remains a major open problem [1, 2, 10, 37].

6.2 Padding to obtain every fixed polynomial exponent

For an integer $k \geq 2$, define

$$\text{PadSAT}_k = \left\{ \varphi \# 1^{|\varphi|^k} : \varphi \in \text{SAT} \right\}. \quad (91)$$

Malformed encodings are rejected. If $s = |\varphi|$, the padded length is $N = s + 1 + s^k = \Theta(s^k)$, while an assignment has at most s bits.

Corollary 4 (Sublinear polynomial parameters). *For every fixed $\alpha > 0$, there is an NP-complete language L_α such that*

$$L_\alpha \in \text{MZK}[O(N^\alpha)]. \quad (92)$$

If PH does not collapse, then $L_\alpha \notin \text{SZK}$.

Proof. On a valid encoding (91), run the protocol with

$$\eta = 2^{-s}, \quad (93)$$

giving likelihood-ratio parameter $s \ln 2 = O(N^{1/k})$. The completeness error is at most $2^{-s} = 2^{-\Theta(N^{1/k})}$, which is negligible as a function of N , and soundness remains perfect. The polynomial-time padding map $\varphi \mapsto \varphi \# 1^{|\varphi|^k}$ preserves satisfiability, so PadSAT_k is NP-complete. Choose an integer $k \geq 2$ with $k > 1/\alpha$ and let $L_\alpha = \text{PadSAT}_k$. Then $N^{1/k} = O(N^\alpha)$, and the conditional exclusion from SZK follows as in Theorem 9. $\square$

Remark 3 (Why (93) matters). *Using $\eta = 2^{-N}$ on the padded input would give parameter $N \ln 2$. Tying the uniform-branch probability to the unpadded length s gives $O(N^{1/k})$ leakage.*

7 Polylogarithmic Leakage and Limited Nondeterminism

For each fixed integer $j \geq 1$, define

$$\mathcal{M}_j := \text{MZK}[O((\log_2(n+2))^j)],$$

where $O(\cdot)$ permits any fixed multiplicative constant and the unspecified additive simulation error is negligible. Monotonicity and Theorem 8 give

$$\text{SZK} = \mathcal{M}_1 \subseteq \mathcal{M}_2 \subseteq \cdots \subseteq \text{MZK}[\text{poly}].$$

Fixed polynomial changes of input length preserve these levels' leakage bounds: a reduction with output length $N = \Theta(n^d)$ pulls back leakage $O(N^\alpha)$ to $O(n^{d\alpha})$, and $O((\log N)^j)$ to $O((\log n)^j)$. Proof and simulation errors require separate tracking.

Write $\text{GC}(b, \mathbf{P})$ for languages verified in polynomial time using at most $b(n)$ guessed bits, and put

$$\beta_j \mathbf{P} := \text{GC}(O((\log(n+2))^j), \mathbf{P}).$$

Enumeration gives $\beta_1 \mathbf{P} = \mathbf{P}$. For $j \geq 2$, Corollary 3 gives the unconditional inclusion

$$\beta_j \mathbf{P} \subseteq \text{MZK}[O((\log n)^j), 0]. \quad (94)$$

We conditionally separate each level $j \geq 2$ from SZK, without separating consecutive levels.

Definition 10 (A subexponential nonuniform nondeterministic hypothesis). For some fixed integer $k \geq 1$, the tautology problem for k -DNF formulas on m declared variables has no nondeterministic circuit family of size $2^{o(m)}$:

$$k\text{-TAUT} \notin \text{NSIZE}(2^{o(m)}). \quad (95)$$

Circuit size is measured in the number of formula variables m ; nondeterministic input bits are existentially quantified. Repeated terms are removed, so a formula has an encoding of length $m^{O(k)}$. The circuit family must correctly recognize tautologies for every m , with no uniformity requirement on the circuits.

This follows from the stronger *Non-uniform Nondeterministic Strong Exponential Time Hypothesis* (NUNSETH) of Carmosino et al. [11, Definition 4]: for every $\xi > 0$, some k has no nondeterministic circuit family of size $O(2^{(1-\xi)m})$ recognizing k -DNF tautologies. Taking $\xi = 1/2$ implies (95), which is the only hypothesis used below.

Theorem 10 (Conditional separation at polylogarithmic leakage). *Assume Definition 10. For every fixed integer $j \geq 2$, there is a language*

$$L_{k,j} \in \beta_j \mathbf{P} \setminus \mathbf{SZK}.$$

It has a one-message general-verifier MZK proof with leakage $O((\log N)^j)$, zero additive simulation error, perfect soundness, and negligible completeness error, where N is its input length. Consequently $\mathbf{SZK} \subsetneq \mathcal{M}_j$ for every $j \geq 2$.

Proof. Fix the k in (95) and an integer $j \geq 2$. Use a canonical fixed-length encoding $\text{enc}_m(F)$ of k -DNFs on declared variables $x_1, \dots, x_m$, with length $\ell_k(m) = m^{O(k)}$. For example, use a self-delimiting header for m and one presence bit for each possible term of at most k literals. Polynomial-time conversion to and from the usual duplicate-free encoding preserves the variable bound and hence (95). Set

$$N_j(m) = 2^{\lceil m^{1/j} \rceil}.$$

Choose a fixed m_0 such that $N_j(m) \geq \ell_k(m) + 1$ for all $m \geq m_0$, and define

$$\begin{aligned} \text{pad}_j(F) &:= \text{enc}_m(F) 0 1^{N_j(m) - \ell_k(m) - 1}, \\ L_{k,j} &:= \{\text{pad}_j(F) : m \geq m_0, F \text{ is not a tautology}\}. \end{aligned}$$

All other binary strings are no-inputs. To recognize the format in time polynomial in the total length N , parse the self-delimiting header and reject unless N is a power of two and $m_0 \leq m \leq (\log_2 N)^j$. Checking this before allocating or enumerating the formula core handles oversized binary headers in polynomial time. Put $t = \log_2 N$ and check $(t-1)^j < m$ to verify $t = \lceil m^{1/j} \rceil$. Check $\ell_k(m) + 1 \leq N$, parse the fixed-length core, and verify the delimiter and padding. For fixed k, j , all these operations take polynomial time. A well-formed input is in $L_{k,j}$ exactly when some m -bit assignment falsifies F . Since $m \leq (\log_2 N)^j$, this proves $L_{k,j} \in \beta_j \mathbf{P}$. To apply Corollary 3, use the length-dependent budget

$$b_j(N) = \lceil \log_2(N+2) \rceil^j,$$

padding the assignment with ignored bits. This budget is polynomial-time computable and super-logarithmic in N , so the corollary supplies the stated proof system.

If $L_{k,j} \in \mathbf{SZK}$, the containment $\mathbf{SZK} \subseteq \mathbf{coAM}$ [1, 37] gives $\overline{L_{k,j}} \in \mathbf{AM} \subseteq \mathbf{NP/poly}$. The last inclusion follows by amplifying the AM error below 2^{-N-1} at polynomial cost and fixing Arthur's coins so the protocol is correct on all 2^N inputs of length N simultaneously. Existentially guessing Merlin's response gives a nondeterministic circuit of size $N^{O(1)}$.

At length $N_j(m)$, compose this circuit with $F \mapsto \text{pad}_j(F)$, hardwiring the header and padding. It accepts precisely the m -variable k -DNF tautologies and has size

$$N_j(m)^{O(1)} = 2^{O(m^{1/j})} = 2^{o(m)}.$$

The finitely many $m < m_0$ can be hardwired separately. This contradicts (95). The separator belongs to $\mathcal{M}_j$ by (94), and $\mathbf{SZK} \subseteq \mathcal{M}_j$ by monotonicity, proving strict containment. $\square$

Coin-fixing is nonuniform, so the argument uses (95) or NUNSETH; it does not establish the same separation from ordinary uniform NSETH.

8 Fully Relativized Separation

The logarithmic leakage threshold is tight relative to oracles. Fix a polynomially bounded integer profile $\ell(n) \geq 1$, computable in polynomial time from 1^n , with $\ell(n) = \omega(\log n)$. We construct a singleton-per-length oracle A and a unary language $L_\ell(A)$ such that

$$L_\ell(A) \in \text{UP}^A \cap \text{coUP}^A \quad \text{but} \quad L_\ell(A) \notin \text{QSZK}^A. \quad (96)$$

An explicit classical protocol then gives

$$L_\ell(A) \in \text{MZK}^A[(\ell(n) - 1) \ln 2, 0] \setminus \text{QSZK}^A, \quad (97)$$

and hence separates it from SZK^A . The oracle may depend on ℓ ; no monotonicity of ℓ is required.

Menda and Watrous [30, Theorem 4] obtain (96) for address length $\ell(n) = n$. We retain their output-state bounds and diagonalize directly at address length $\ell(n)$. The quantum background below fixes the conventions needed for this construction.

8.1 Unambiguous computation and the oracle language

8.1.1 The classes UP and coUP

For a nondeterministic Turing machine M , let $\text{acc}_M(x)$ denote the number of accepting computation paths of M on input x . A language L belongs to *unambiguous polynomial time*, denoted UP , if there is a polynomial-time nondeterministic Turing machine M such that

$$\text{acc}_M(x) = \begin{cases} 1, & x \in L, \\ 0, & x \notin L. \end{cases} \quad (98)$$

This class was introduced by Valiant [43] and satisfies $\text{UP} \subseteq \text{NP}$. Define

$$\text{coUP} := \{L : \overline{L} \in \text{UP}\}. \quad (99)$$

Membership in $\text{UP} \cap \text{coUP}$ requires an unambiguous machine for each of L and $\overline{L}$.

For an oracle $A \subseteq \{0, 1\}^*$, allow M unit-cost membership queries to A with the same running-time and accepting-path requirements to define UP^A ; set $\text{coUP}^A := \{L : \overline{L} \in \text{UP}^A\}$. Oracle answers are fixed by A and introduce no nondeterministic paths.

8.1.2 The classical oracle model

The class MZK^A gives every participant access to the same fixed oracle A . The prescribed prover P^A is computationally unbounded; the honest verifier V^A and each cheating verifier V^{*A} are polynomial-time oracle machines. For each V^{*A} , a strict-PPT oracle simulator $S_{V^*}^A$ receives the same auxiliary input and must satisfy both simulation inequalities on every yes-input without interacting with the prover. Writing queries counts toward running time. A complete view records the verifier's coins and messages; oracle queries and answers may also be included, since they are reconstructible relative to A . Soundness quantifies over all computationally unbounded oracle provers.

8.1.3 The unary witness language

Let $\Sigma = \{0, 1\}$ and

$$\mathcal{A} = \{A \subseteq \Sigma^* : |A \cap \Sigma^r| = 1 \text{ for every } r \geq 1\}. \quad (100)$$

For $A \in \mathcal{A}$, write a_r for the unique string in $A \cap \Sigma^r$. The input length is n ; the relevant oracle address length is $r = \ell(n)$. Define

$$L_\ell(A) := \{0^n : n \geq 1 \text{ and the first bit of } a_{\ell(n)} \text{ is } 1\}. \quad (101)$$

Nonunary strings and the empty string are no-inputs.

Lemma 4. *For every polynomially bounded, polynomial-time computable integer profile $\ell(n) \geq 1$ and every $A \in \mathcal{A}$, $L_\ell(A) \in \text{UP}^A \cap \text{coUP}^A$.*

Proof. On 0^n with $n \geq 1$, the UP^A machine computes $r = \ell(n)$, guesses $y \in \Sigma^{r-1}$, and accepts iff $1y \in A$. The complement machine instead accepts iff $0y \in A$. The singleton property gives exactly one accepting path for the appropriate machine and none for the other. On nonunary or empty inputs, the first machine rejects and the complement machine accepts deterministically. Both machines run in polynomial time. This is the witness argument of Menda and Watrous [30, Section 4.1], with address length r in place of input length n . $\square$

8.2 Relativized quantum state distinguishability

8.2.1 States, trace distance, and fidelity

We use the standard conventions [29, 32, 44]. Let $\mathcal{H}$ be a finite-dimensional complex Hilbert space. A quantum state on $\mathcal{H}$ is a density operator $\rho \in \text{L}(\mathcal{H})$ satisfying $\rho \succeq 0$ and $\text{Tr}(\rho) = 1$. Pure states have the form $|\psi\rangle\langle\psi|$; discarding work qubits by partial trace can produce a mixed state. For an operator X , its trace norm is

$$\|X\|_1 := \text{Tr} \sqrt{X^\dagger X}. \quad (102)$$

The trace distance between states ρ and σ is

$$\Delta(\rho, \sigma) := \frac{1}{2} \|\rho - \sigma\|_1. \quad (103)$$

Our fidelity convention is

$$F(\rho, \sigma) := \|\sqrt{\rho}\sqrt{\sigma}\|_1. \quad (104)$$

The Fuchs–van de Graaf inequalities [16] give

$$1 - F(\rho, \sigma) \leq \Delta(\rho, \sigma) \leq \sqrt{1 - F(\rho, \sigma)^2}. \quad (105)$$

These unsquared-fidelity conventions agree with Menda and Watrous [30, Section 2].

8.2.2 Oracle query gates and uniform circuit families

The restriction $B = A \cap \Sigma^m$ of an oracle A is an m -bit black box. The corresponding membership-query gate is the unitary

$$K_B|z, b, u\rangle = |z, b \oplus \mathbf{1}[z \in B], u\rangle, \quad (106)$$

where $z \in \Sigma^m$, b is the answer qubit, and u denotes all untouched workspace qubits. A relativized quantum circuit may query several address lengths, each using the corresponding restriction of A . A family $\{Q^x : x \in \Gamma^*\}$ is polynomial-time uniform if a deterministic classical polynomial-time algorithm maps x to an explicit description of Q^x , using a fixed finite universal set for its nonoracle gates. The input x is hardwired: Q^x may have no input qubits and may output a mixed state by tracing out designated work qubits.

8.2.3 The relativized QSZK characterization

Quantum statistical zero-knowledge (QSZK) is defined by quantum interactive proofs whose verifier views on yes-inputs are reproducible by efficient quantum circuits up to negligible trace distance. We use its complete problem QSD, *Quantum State Distinguishability*, due to Watrous [45]. The completeness proof relativizes by treating oracle-query gates as gates available to both verifier and simulator. Theorem 1 of Menda and Watrous [30] states that, for every language $L \subseteq \Gamma^*$ and oracle A , $L \in \text{QSZK}^A$ if and only if there is a polynomial-time uniform family of pairs of relativized circuits (Q_0^x, Q_1^x) , with equally sized output registers and output states $\rho_0^x(A), \rho_1^x(A)$, satisfying

$$\begin{aligned} x \in L &\implies \Delta(\rho_0^x(A), \rho_1^x(A)) \geq \frac{2}{3}, \\ x \notin L &\implies \Delta(\rho_0^x(A), \rho_1^x(A)) \leq \frac{1}{3}. \end{aligned} \tag{107}$$

Thus, to prove $L \notin \text{QSZK}^A$, it suffices to defeat every such uniform circuit family on at least one input.

8.2.4 The imported black-box bounds

We use two quantitative bounds of Menda and Watrous. Here r is the black-box address length, independent of the public input length. The bounds allow arbitrary operations independent of the black box between queries; they impose no uniformity or circuit-size requirement.

Lemma 5 (Average fidelity). *Let Q be a quantum circuit with no input qubits, making at most T queries to an r -bit black box. Let $\rho(B)$ be its output state when the black box is $B \subseteq \Sigma^r$. Then*

$$2^{-r} \sum_{x \in \Sigma^r} F(\rho(\{x\}), \rho(\emptyset)) \geq 1 - \frac{2T}{2^{r/2}}. \tag{108}$$

Lemma 5 is Lemma 2 of Menda and Watrous [30].

Lemma 6 (Stability of output-state discrimination). *Let Q_0, Q_1 have no input qubits, each make at most T queries to an r -bit black box, and output states $\rho_0(B), \rho_1(B)$ on equally sized registers. Put*

$$d_B := \Delta(\rho_0(B), \rho_1(B)). \tag{109}$$

If

$$T \leq \frac{2^{r/2}}{20736}, \tag{110}$$

then at least $(2/3)2^r$ strings $x \in \Sigma^r$ satisfy

$$|d_{\{x\}} - d_{\emptyset}| < \frac{1}{6}. \tag{111}$$

This is Lemma 3 of Menda and Watrous [30], derived from Lemma 5, the Fuchs–van de Graaf inequalities, Jensen’s inequality, Markov’s inequality, and the triangle inequality.

Lemma 7 (A singleton that defeats the circuit pair). *Fix a public input 0^n whose membership is determined by the first bit of the singleton in an r -bit black box, where $r \geq 1$. Under the hypotheses of Lemma 6, at least $2^r/6$ choices $x \in \Sigma^r$ cause (Q_0, Q_1) to violate (107) on 0^n when the black box is $\{x\}$.*

Proof. Let S be the set of at least $(2/3)2^r$ strings supplied by Lemma 6. If $d_\emptyset \leq 1/2$, then $|S \cap 1\Sigma^{r-1}| \geq (2/3 - 1/2)2^r = 2^r/6$, and every x in this set has $d_{\{x\}} < 2/3$, violating the yes-condition for 0^n . If $d_\emptyset \geq 1/2$, then the same count gives at least $2^r/6$ strings in $S \cap 0\Sigma^{r-1}$, each satisfying $d_{\{x\}} > 1/3$, violating the no-condition. These cases cover all $d_\emptyset$. This is the black-box argument of Menda and Watrous [30, Section 4.1]. $\square$

8.3 Finite-extension oracle diagonalization

Theorem 11 (Oracle separation from quantum SZK). *For every polynomially bounded, polynomial-time computable integer profile $\ell(n) \geq 1$ with $\ell(n) = \omega(\log n)$, there is an oracle $A \in \mathcal{A}$ such that*

$$L_\ell(A) \in \text{UP}^A \cap \text{coUP}^A \quad \text{and} \quad L_\ell(A) \notin \text{QSZK}^A. \quad (112)$$

Proof. Enumerate all clocked polynomial-time generators of pairs of relativized quantum circuits with equally sized output registers. The generators are classical and oracle-independent, as in (107); the generated circuits query A . Replace malformed outputs by a fixed valid pair. For family i , let an integer polynomial $p_i(n) \geq 1$ bound each circuit's query count and largest queried address length on inputs of length n . Every family in (107) occurs in this enumeration with some polynomial clock.

Fix the empty string's oracle membership arbitrarily and set $B_0 = n_0 = 0$. After stage $i - 1$, all slices of positive length at most B_{i-1} contain one fixed string. Choose $n_i > n_{i-1}$ so that

$$r_i := \ell(n_i) > B_{i-1}, \quad p_i(n_i) \leq \frac{2^{r_i/2}}{20736}, \quad (113)$$

and put

$$B_i := \max\{B_{i-1}, r_i, p_i(n_i)\}. \quad (114)$$

Such n_i exist because $\ell(n) \rightarrow \infty$ and $2^{\ell(n)/2}/p_i(n) \rightarrow \infty$. Neither condition requires ℓ to be monotone.

At every still-unassigned length at most B_i other than r_i , choose an arbitrary singleton and fix that slice. In the i th circuit pair on 0^{n_i} , hardwire every query gate at lengths other than r_i using these fixed answers. The resulting pair has the same output states for every choice of the r_i -bit black box and makes at most $p_i(n_i)$ queries to it. The hardwired operations are independent of that black box, so Lemma 7 supplies a string $a_{r_i} \in \Sigma^{r_i}$ whose singleton makes this pair fail (107) on 0^{n_i} . Fix the slice at r_i to $\{a_{r_i}\}$.

All slices through B_i are now fixed. They determine both the i th pair's output states on 0^{n_i} and that input's membership in $L_\ell(A)$. Later stages preserve these slices, so the failure persists even if ℓ repeats or decreases at other input lengths. Since $B_i > B_{i-1}$, the union of the finite stages defines an oracle in $\mathcal{A}$. Every uniform circuit family fails on its selected input; (107) gives $L_\ell(A) \notin \text{QSZK}^A$. Lemma 4 supplies the two unambiguous upper bounds. $\square$

The stage lengths depend on ℓ , so this argument constructs the oracle for the requested profile. It does not pad a language for a previously fixed oracle.

8.4 The relativized multiplicative zero-knowledge protocol

Fix ℓ and A from Theorem 11. The verifier rejects nonunary and empty inputs without requesting a prover message. On 0^n with $n \geq 1$, set $q = \ell(n) - 1$; the verifier receives one q -bit string Y and accepts iff $1Y \in A$, rejecting malformed messages. There are no verifier challenges. The prover sends at most one message and ignores unsolicited messages.

On a yes-input, let $y^* \in \Sigma^q$ be the unique string satisfying $1y^* \in A$. The unbounded oracle prover can find y^* by exhaustive search. Set $\eta = 2^{-q}$ and prescribe

$$Y = \begin{cases} y^*, & \text{with probability } 1 - \eta, \\ U_q, & \text{with probability } \eta, \end{cases} \quad U_q \leftarrow \text{Unif}(\Sigma^q). \quad (115)$$

This mixture uses at most $2q$ independent fair random bits. On a unary no-input the prover may send any fixed q -bit string. The honest-verifier simulator samples U_q and runs the verifier relative to A , outputting its complete view.

Theorem 12 (Sharp oracle threshold for multiplicative leakage). *For every polynomially bounded, polynomial-time computable integer profile $\ell(n) \geq 1$ with $\ell(n) = \omega(\log n)$, there is an oracle $A \in \mathcal{A}$ for which*

$$L_\ell(A) \in \text{MZK}^A[(\ell(n) - 1) \ln 2, 0] \setminus \text{SZK}^A. \quad (116)$$

The language also belongs to $\text{UP}^A \cap \text{coUP}^A$ and lies outside SZK^A . The prover sends $q = \ell(n) - 1$ bits once. The protocol has perfect soundness, completeness error at most $2^{1-\ell(n)}$, and zero additive simulation error. At length zero the verifier rejects immediately and the leakage is zero.

Proof. Let R_{y^*} and Q be the real and simulated message distributions. For every $y \in \Sigma^q$,

$$R_{y^*}(y) = (1 - \eta)\mathbf{1}[y = y^*] + \eta 2^{-q}, \quad Q(y) = 2^{-q}. \quad (117)$$

At y^* ,

$$\frac{R_{y^*}(y^*)}{Q(y^*)} = (1 - \eta)2^q + \eta \leq 2^q, \quad (118)$$

and away from y^* ,

$$\frac{Q(y)}{R_{y^*}(y)} = \eta^{-1} = 2^q. \quad (119)$$

The opposite ratio is at most 1 in each case. Hence $R_{y^*} \approx_{q \ln 2, 0} Q$.

For $q \geq 1$, the honest verifier's rejection probability is $\eta(1 - 2^{-q}) \leq 2^{-q}$; for $q = 0$, both message distributions are the point mass on the empty string and the verifier accepts with probability one. Since $q = \ell(n) - 1 = \omega(\log n)$, the completeness error is negligible. On a nonempty unary no-input there is no y satisfying $1y \in A$; nonunary and empty inputs are rejected immediately. Thus soundness is perfect. As in (7), any cheating oracle verifier's complete view is a randomized postprocessing, relative to A , of the sole prover message. Its strict-PPT oracle simulator samples one U_q and runs the verifier with the same oracle and auxiliary input, supplying the sample in the sole message slot. Randomized postprocessing makes the full views, including oracle activity, symmetrically $(q \ln 2, 0)$ -close. The prover supplies no further samples even if requested. Computing q , sampling U_q , and running the verifier take polynomial time because ℓ is polynomial-time computable and polynomially bounded. Hence $L_\ell(A) \in \text{MZK}^A[(\ell(n) - 1) \ln 2, 0]$.

Finally, the Statistical Difference characterization of SZK [37] relativizes. Implementing its two classical oracle samplers as quantum circuits produces diagonal states whose trace distance equals the original statistical distance. The characterization in (107) therefore gives $\text{SZK}^A \subseteq \text{QSZK}^A$. Theorem 11 gives $L_\ell(A) \notin \text{QSZK}^A$, hence $L_\ell(A) \notin \text{SZK}^A$, as well as membership in $\text{UP}^A \cap \text{coUP}^A$. $\square$

Sharpness at logarithmic leakage. The collapse in Theorem 8 relativizes to every fixed oracle A :

$$\text{HVMZK}^A[O(\log n)] = \text{MZK}^A[O(\log n)] = \text{SZK}^A.$$

Its view replay, bounded sampling, prefix-entropy construction, and reductions to Statistical Difference and Entropy Difference all use the same oracle and preserve their error bounds. Conversely, Theorem 12 separates the classes for every profile in its hypotheses, with an oracle that may depend on the profile. This includes integer profiles of orders $\log n \log \log n$ and $(\log n)^2$: with $t = \lceil \log_2(n+2) \rceil$, take

$$\ell_1(n) = t \lceil \log_2(t+1) \rceil, \quad \ell_2(n) = t^2.$$

More generally, let $\varepsilon(n) \geq 0$ be a rational-valued, polynomially bounded profile, computable in polynomial time from 1^n , with $\varepsilon(n) = \omega(\log n)$. Taking $\ell(n) = 1 + \lfloor \varepsilon(n) \rfloor$ gives leakage $\lfloor \varepsilon(n) \rfloor \ln 2 \leq \varepsilon(n)$. Thus there is an oracle A and a language in $\text{MZK}^A[\varepsilon, 0] \setminus \text{QSZK}^A$ for every such budget. For the singleton language itself, $\ell(n) = O(\log n)$ allows a deterministic oracle machine to enumerate all $2^{\ell(n)}$ addresses, so $L_\ell(A) \in \text{P}^A$ for every $A \in \mathcal{A}$.

9 Complete Problems and a Canonical Protocol Characterization

The union of polynomial-leakage classes equals PSPACE and has TQBF as a complete problem. Suitable padding gives completeness at every fixed polynomial budget. At intermediate budgets, such as $\varepsilon(n) = (\ln n)^2$, we give a canonical reduction for protocols with specified leakage, simulation-error, and proof-error profiles. The reduction records the source input length; it does not establish ordinary fixed-budget completeness measured at the target encoding length.

Views contain the verifier's complete random tape and full transcript; conversations have polynomially bounded length. Completeness, soundness, and unspecified additive simulation errors are negligible, as in our earlier definitions and the standard interactive-proof and zero-knowledge model [27].

9.1 Polynomial leakage: PSPACE and TQBF

Define

$$\begin{aligned} \text{MZK}[\text{poly}] &:= \bigcup_{c>0} \text{MZK}[n^c], \\ \text{HVMZK}[\text{poly}] &:= \bigcup_{c>0} \text{HVMZK}[n^c]. \end{aligned} \tag{120}$$

The second union imposes the honest-verifier restriction. Increasing the exponent absorbs polynomial constants and degrees; finitely many short inputs are handled separately.

Theorem 13 (Polynomial-leakage collapse). *Even with pure multiplicative simulation,*

$$\text{MZK}[\text{poly}] = \text{HVMZK}[\text{poly}] = \text{PSPACE}. \tag{121}$$

Consequently, TQBF is complete for both union classes under polynomial-time many-one reductions.

Proof. By Shamir's theorem, $\text{IP} = \text{PSPACE}$ [39], so every language in PSPACE has an interactive proof with polynomially many rounds and polynomial communication. Sequential repetition makes completeness and soundness errors negligible while preserving polynomial communication. Normalize

the message slots and public abort rules as required by Lemma 3, and let $m(n)$ count all prover bits in this interface. Applying the lemma with $h(n) = n$ gives pure symmetric multiplicative leakage

$$\varepsilon(n) = \ln 2 \cdot \max\{m(n), n\} = \text{poly}(n) \quad (122)$$

against every PPT verifier. Hence $\text{PSPACE} \subseteq \text{MZK}[\text{poly}]$.

Conversely,

$$\text{MZK}[\text{poly}] \subseteq \text{HVMZK}[\text{poly}] \subseteq \text{IP} = \text{PSPACE}, \quad (123)$$

because removing the simulation requirement leaves an interactive proof. The class equality and PSPACE -completeness of true quantified Boolean formulas (TQBF) [40] give the completeness claim. $\square$

9.1.1 Fixed polynomial budgets and padding

Fix an amplified, normalized TQBF interactive proof whose prover communication, including abort encodings, on a formula of length ℓ is at most ℓ^D , where $D \geq 1$ is a constant. Theorem 13 gives

$$\text{TQBF} \in \text{MZK}[O(n^D)]. \quad (124)$$

Thus TQBF is complete for $\text{MZK}[O(n^D)]$, since that class is contained in PSPACE .

For smaller exponents, fix $c > 0$, choose a rational b with $0 < b < c$, and choose an integer $k > D/b$. Define

$$\text{PadTQBF}_k := \{\varphi \# 1^{|\varphi|^k} : \varphi \in \text{TQBF}\}. \quad (125)$$

Reject inputs outside this padding format. On a padded input of length $N = \ell + 1 + \ell^k = \Theta(\ell^k)$, run the original protocol on φ with the uniform branch of Lemma 3 chosen with probability

$$\theta(N) = 2^{-\lceil N^b \rceil}. \quad (126)$$

For fixed rational $b = u/v > 0$, the integer $\lceil N^b \rceil$ is computed exactly by finding the least integer a with $a^v \geq N^u$. Thus the branch probability is sampled exactly with polynomially many fair bits. This probability is negligible in the padded length N , as are the original protocol's errors because $N = \Theta(\ell^k)$ for fixed k . Lemma 3 gives

$$\begin{aligned} \varepsilon(N) &\leq \ln 2 \cdot \max\{\ell^D, \lceil N^b \rceil\} \\ &= O\left(N^{\max\{D/k, b\}}\right) = O(N^c). \end{aligned} \quad (127)$$

The padding map is polynomial time for fixed k , so PadTQBF_k remains PSPACE -complete.

Corollary 5. *For every fixed $c > 0$, some polynomially padded version of TQBF is complete for $\text{MZK}[O(N^c)]$ under ordinary polynomial-time many-one reductions.*

When $c < 1$, choosing 2^{-N} instead of (126) would cost $\Theta(N)$ leakage and would not prove membership in $\text{MZK}[O(N^c)]$.

9.2 A canonical reduction for fixed protocol profiles

At intermediate leakage budgets, such as $\varepsilon(n) = (\ln n)^c$ for fixed $c > 1$, TQBF is not known to satisfy the same budget. Our canonical promise problem therefore retains the verifier and its interaction with the prover. An online prover chooses each message using only the common input, observed interaction history, and its own randomness. It has no access to unrevealed verifier coins.

9.2.1 Views generated by a prover

Let $\mathcal{V}$ be a finite verifier circuit with q uniform random bits and a fixed alternating message interface. Let $\mathcal{C}(\mathcal{V})$ be the set of full-view distributions generated by allowed, possibly unbounded, online prover strategies in the underlying interactive-proof model. If a_i is the i th prover message and h_{i-1} is the public history visible immediately before it, every $R \in \mathcal{C}(\mathcal{V})$ has the form

$$R(r, a_1, \dots, a_t) = 2^{-q} \prod_{i=1}^t K_i(a_i \mid h_{i-1}) \quad (128)$$

on transcripts consistent with $\mathcal{V}$, and is zero otherwise, for conditional message distributions K_i chosen by one prover strategy. Verifier messages are determined by r and the preceding transcript and are implicit in (128). Each K_i depends on r only through messages already in h_{i-1} .

9.2.2 Hockey-stick divergence

For $\alpha \geq 1$, the *hockey-stick divergence* is the f -divergence [13, 31, 35, 38]

$$\text{HS}_\alpha(P\|Q) := \sup_E (P(E) - \alpha Q(E)) = \sum_z (P(z) - \alpha Q(z))_+. \quad (129)$$

The equality follows by choosing $E = \{z : P(z) > \alpha Q(z)\}$. Consequently,

$$P \preceq_{\varepsilon, \delta} Q \iff \text{HS}_{e^\varepsilon}(P\|Q) \leq \delta. \quad (130)$$

Symmetric (ε, δ) -closeness is equivalent to this bound in both directions.

9.2.3 Definition of the promise problem

Fix profiles $\varepsilon(n) \geq 0$, $0 \leq \delta(n) < 1$, and proof-error bounds $c_{\text{ip}}(n), s_{\text{ip}}(n) \in [0, 1]$ satisfying

$$c_{\text{ip}}(n) + s_{\text{ip}}(n) < 1 \quad \text{for every } n \geq 1. \quad (131)$$

Fix a polynomial B as part of the problem specification, independently of the instance. An instance of

$$\text{IMDA}_{\varepsilon, \delta}^{c_{\text{ip}}, s_{\text{ip}}}[B] \quad (\text{Interactive Max-Divergence Acceptance}) \quad (132)$$

is an explicitly encoded triple $(1^n, \mathcal{V}, \mathcal{S})$ with $n \geq 1$ and total bit length $N \leq B(n)$. We use Boolean circuits over a fixed finite basis with independent uniform random-input bits. The encoding explicitly lists the verifier's round circuits, its random-tape length, and all message widths; it fixes the alternating message interface and the handling of malformed messages and early aborts. All these lengths are bounded by the explicit encoding, excluding succinct exponential interfaces. The sampler circuit outputs a string in the same fixed-length encoding of full views, including any designated malformed-output symbol. Invalid encodings are outside the promise.

Let Q be the output distribution of $\mathcal{S}$, and let $A_{\mathcal{V}}$ be the set of valid accepting full views. The set $\mathcal{C}(\mathcal{V})$ admits only factorizations induced by prover strategies allowed in the underlying interactive-proof model.

The promise consists of the following two cases:

$$\text{YES} : \exists R \in \mathcal{C}(\mathcal{V}) \text{ such that } R(A_{\mathcal{V}}) \geq 1 - c_{\text{ip}}(n), \quad (133)$$

$$\text{HS}_{e^{\varepsilon(n)}}(R\|Q) \leq \delta(n), \quad \text{HS}_{e^{\varepsilon(n)}}(Q\|R) \leq \delta(n); \quad (134)$$

$$\text{No} : \sup_{R \in \mathcal{C}(\mathcal{V})} R(A_{\mathcal{V}}) \leq s_{\text{ip}}(n). \quad (135)$$

By (131), a yes-witness has acceptance strictly above the no threshold, so the two cases are disjoint. Inputs satisfying neither are outside the promise. All four profiles are evaluated at the tag n , which need not equal the complete encoding length N .

Theorem 14 (Fixed-profile canonical reduction). *Let a promise problem Π have an honest-verifier protocol (P, V) and a strict-PPT simulator S with symmetric simulation parameters $(\varepsilon(n), \delta(n))$, completeness error at most $c_{\text{ip}}(n)$, and soundness error at most $s_{\text{ip}}(n)$. On nonempty inputs, there is a fixed polynomial B and a polynomial-time map*

$$f(x) = (1^{|x|}, \mathcal{V}_x, \mathcal{S}_x) \quad (136)$$

that sends Π_Y to the yes set and Π_N to the no set of $\text{IMDA}_{\varepsilon, \delta}^{c_{\text{ip}}, s_{\text{ip}}}[B]$. Every profile is preserved exactly as a function of the source length $n = |x|$. This includes $\varepsilon(n) = (\ln n)^2$.

Proof. Compile the fixed polynomial-time machines $V(x, \cdot)$ and $S(x, \cdot)$ into Boolean circuits $\mathcal{V}_x$ and $\mathcal{S}_x$, including their complete message interface and view encoding. Strict polynomial running time bounds the random bits and total circuit size. Choose one polynomial B that bounds the encoding lengths for all inputs; it may depend on the fixed machines, but not on x .

If $x \in \Pi_Y$, the real view $R_x = \text{View}_{P,V}(x)$ belongs to $\mathcal{C}(\mathcal{V}_x)$. Completeness gives (133), while symmetric simulation and (130) give both inequalities in (134). If $x \in \Pi_N$, every allowable prover strategy against $\mathcal{V}_x$ is a strategy against V on x . Soundness therefore gives (135). The tag is $|x|$, so every bound retains its source-length argument. Inputs with $n = 0$ can be handled separately. $\square$

The promise exactly records the acceptance and simulation constraints at each encoded instance: a yes-witness induces acceptance at least $1 - c_{\text{ip}}(n)$ and a view symmetrically $(\varepsilon(n), \delta(n))$ -close to the output of $\mathcal{S}$; on a no-instance every strategy accepts with probability at most $s_{\text{ip}}(n)$. This does not establish membership in the ordinary input-length class $\text{HVMZK}[\varepsilon, \delta]$, whose profiles are evaluated at N . A unary tag does not identify $\varepsilon(n)$ with $\varepsilon(N)$ or $\delta(n)$ with $\delta(N)$, and the per-instance existential witness does not supply a uniform prescribed prover for the canonical problem.

For $\delta = 0$, the promise (134) has the equivalent pointwise form

$$e^{-\varepsilon(n)}Q(z) \leq R(z) \leq e^{\varepsilon(n)}Q(z) \quad \text{for every } z. \quad (137)$$

Remark 4 (Fixed errors and encoding length). *The theorem requires the specified proof-error bounds: some negligible errors need not meet a previously fixed negligible pair $c_{\text{ip}}, s_{\text{ip}}$. For constant proof errors, one may fix $c_{\text{ip}} = s_{\text{ip}} = 1/3$. Repetition requires separate parameter analysis because product simulation generally increases leakage and accumulates additive error. The bound $N \leq B(n)$ controls representation size without identifying fixed profiles at lengths n and N .*

9.3 The scope of the canonical reduction

Theorem 13 gives general-verifier completeness of TQBF for $\text{MZK}[\text{poly}]$. Theorem 14 compiles an honest-verifier protocol with specified profiles into a promise evaluated at the source length. It also applies to general-verifier protocols by restricting simulation to the honest verifier. This does not establish ordinary completeness for either fixed-profile class.

The encoded sampler $\mathcal{S}$ simulates only the encoded honest verifier. General-verifier membership would additionally require simulating cheating verifiers in the canonical interaction, or an honest-verifier-to-general-verifier transformation with established parameter costs. This requirement is separate from the length and proof-error conventions.

9.4 Why an ordinary two-sampler problem does not yet follow

Statistical Difference is an interaction-free complete problem for SZK on two efficient samplers [37]. An analogous reduction for superlogarithmic leakage faces both quantitative and structural obstacles. For real view R , simulator distribution Q , and accepting event A ,

$$R(A) \leq e^\varepsilon Q(A) + \delta \implies Q(A) \geq e^{-\varepsilon}(R(A) - \delta). \quad (138)$$

For $R(A) \geq 1 - c_{\text{ip}}$ and $\varepsilon(n) = (\ln n)^2$, the guaranteed simulator mass is only

$$e^{-(\ln n)^2}(1 - c_{\text{ip}} - \delta) = n^{-\ln n}(1 - c_{\text{ip}} - \delta), \quad (139)$$

which is superpolynomially small. This bound does not support the polynomial-time rejection sampling used in the logarithmic-leakage collapse.

The real view may require an unbounded prover and need not have a PPT sampler. A sampler can output valid transcripts yet correlate an early prover message with an unrevealed verifier coin, violating the factorization (128). The set $\mathcal{C}(\mathcal{V})$ includes only views generated by a prover.

The canonical reduction therefore retains the interaction with an online prover. Reducing IMDA to a promise problem on a constant number of ordinary PPT samplers would require additional arguments, such as an efficient conditioning theorem.

10 Implications for Fraction Knowledge Complexity

The real-to-simulator inequality $R \preceq_{\varepsilon, \delta} Q$ yields fraction simulation while preserving the prover, verifier, and simulator output distribution. We select good simulator tapes, appending unused random bits when necessary.

Theorem 15. *Let (P, V) be an interactive proof and S one strict finite-coin PPT simulator. For every yes-input x of length n , write $R = \text{View}_{P,V}(x)$ and $Q = S(x)$. Suppose a polynomially bounded profile $\varepsilon(n) \geq 0$ satisfies*

$$R \preceq_{\varepsilon, \delta} Q, \quad 0 \leq \delta < 1. \quad (140)$$

Ideal real-valued selection gives a good-tape event of density exactly

$$\rho = e^{-\varepsilon} \quad (141)$$

whose conditional output has statistical distance at most δ from R . For finite tapes, fix $\xi(n) = 2^{-p(n)}$, where p is a positive integer-valued polynomial of degree at least one. Appending polynomially many unused random bits gives, for each x , an event $G = G_x$ whose conditional output Q_G satisfies

$$\Pr[G] \geq e^{-\varepsilon}, \quad (142)$$

$$\text{SD}(Q_G, R) \leq \delta + 2\xi. \quad (143)$$

The appended tape length depends only on n ; the good set may depend on x . The simulator's output law remains Q . If $\delta(n)$ is negligible, the protocol has fraction statistical knowledge complexity

$$k \leq \frac{\varepsilon}{\ln 2}. \quad (144)$$

Proof. Fix a yes-input x and suppress the input and length parameters. For finite distributions, (140) is equivalent to

$$\sum_z (R(z) - e^\varepsilon Q(z))_+ \leq \delta. \quad (145)$$

The left-hand side is the hockey-stick divergence $\sup_E (R(E) - e^\varepsilon Q(E))$, attained by $E = \{z : R(z) > e^\varepsilon Q(z)\}$.

Define

$$\mu(z) = \min\{R(z), e^\varepsilon Q(z)\}, \quad m = \sum_z \mu(z). \quad (146)$$

Equation (145) gives

$$m = 1 - \sum_z (R(z) - e^\varepsilon Q(z))_+ \geq 1 - \delta. \quad (147)$$

Set

$$\rho = e^{-\varepsilon}, \quad \tau(z) = \rho \mu(z). \quad (148)$$

Then $\tau \leq Q$ has mass ρm . Complete it within Q to mass ρ : when $\rho m < 1$, set

$$\sigma = \tau + \frac{\rho(1-m)}{1-\rho m}(Q - \tau). \quad (149)$$

Since $\rho \leq 1$, the coefficient lies in $[0, 1]$, so $\tau \leq \sigma \leq Q$ and $\sum_z \sigma(z) = \rho$. If $\rho m = 1$, take $\sigma = Q = R$. Both $R' = \sigma/\rho$ and R dominate μ , so the overlap formula gives

$$\text{SD}(R, R') \leq 1 - m \leq \delta. \quad (150)$$

Ideally, retaining a $\sigma(z)/Q(z)$ fraction of the tapes producing each $z \in \text{supp}(Q)$ gives selected subdistribution σ , density ρ , and conditional output R' .

Finite-tape realization. Pad the simulator tape to its polynomial bound $r = r(n)$ and append $t = t(n)$ unused bits. Each original tape now has 2^t copies, with output law still Q . Put $N = 2^{r+t}$ and select

$$N_z = \lfloor N\sigma(z) \rfloor \quad (151)$$

extended preimages of each z , possible since $\sigma \leq Q$. The selected subdistribution $\tilde{\sigma}$ has mass $\rho - \lambda$, where

$$0 \leq \lambda < \frac{|\text{supp}(Q)|}{N} \leq 2^{-t}, \quad (152)$$

because the original sampler has at most 2^r possible outputs. Add unused tapes until exactly $M = \lceil N\rho \rceil$ are selected; the initial count is at most $N\rho$, and $M \leq N$. Let G be this event, γ its output subdistribution, and $g = \Pr[G] = M/N$. Then $\gamma \geq \tilde{\sigma}$ and

$$g = \rho + \omega \geq \rho, \quad 0 \leq \omega < 1/N. \quad (153)$$

Choose an integer-valued polynomial $h(n) \geq \varepsilon(n)/\ln 2$ and set $t(n) = p(n) + h(n)$. This choice is uniform over all yes-inputs of length n and satisfies

$$2^{-t} \leq \xi e^{-\varepsilon} = \xi \rho. \quad (154)$$

Then $\lambda, \omega \leq \xi \rho$. Since $g \geq \rho$, the distributions $R' = \sigma/\rho$ and $Q_G = \gamma/g$ share the subdistribution $\tilde{\sigma}/g$ of mass $(\rho - \lambda)/g$. Hence

$$\text{SD}(Q_G, R') \leq 1 - \frac{\rho - \lambda}{g} = \frac{\lambda + \omega}{g} \leq 2\xi. \quad (155)$$

Together with (150) and $g \geq \rho$, this proves both finite-tape bounds. For negligible δ , the error $\delta + 2\xi$ is negligible and $\log_2(1/\Pr[G]) \leq \varepsilon/\ln 2$, proving (144). Fraction knowledge requires only existence of G , not efficient recognition. $\square$

10.1 Implications for AM and coAM

The bound $\varepsilon/\ln 2$ is tight for exact simulation. When $e^{-\varepsilon}$ is dyadic, take $R = \delta_a$ and $Q(a) = e^{-\varepsilon}$. Then $R \preceq_{\varepsilon,0} Q$, but any simulator-tape event whose conditional output equals R contains only tapes outputting a , so its density is at most $e^{-\varepsilon}$. For other ε , dyadic choices of $Q(a)$ approach this bound from above.

For negligible δ , Theorem 15 gives $k \leq \varepsilon/\ln 2$. The logarithmic fraction-knowledge bound of Petrank and Tardos [34] then gives

$$\text{MZK}[O(\log n)] \subseteq \text{AM} \cap \text{coAM}.$$

Section 4 gives the stronger SZK upper bound under symmetric closeness.

10.2 One-Sided Characterization and the Symmetric Converse

Recall that $\text{HVoMZK}^{\rightarrow}[\kappa, \nu]$ (Definition 4) requires only

$$R \preceq_{\kappa, \nu} Q, \tag{156}$$

with uniform strict-PPT simulators and no simulator-to-real requirement. This honest-verifier notion matches our fraction-knowledge convention. We omit the second argument when it is negligible.

Proposition 3 (One-sided correspondence). *Both implications below preserve the protocol.*

1. *If a simulator has a good-tape event of density at least 2^{-k} and its conditional output is within statistical distance ν of the real view, then*

$$R \preceq_{k \ln 2, \nu} Q. \tag{157}$$

2. *If $R \preceq_{\kappa, \delta} Q$ with polynomially bounded $\kappa \geq 0$, then for every inverse-exponential negligible $\xi = 2^{-p(n)}$, appending polynomially many unused random bits to the simulator gives a good-tape event of density at least*

$$e^{-\kappa} \tag{158}$$

whose conditional output is within $\delta + 2\xi$ of R .

Consequently, for nonnegative polynomially bounded k , with negligible statistical errors and uniform strict-PPT simulators on both sides,

$$\text{SKC}_{\text{frac}}[k] = \text{HVoMZK}^{\rightarrow}[k \ln 2]. \tag{159}$$

The equality is exact in k : rounding changes only the negligible error, not the simulator's output law.

Proof. For the good event G , write $\rho = \Pr[G] \geq 2^{-k}$ and let Q_G be the conditional simulator output. Partitioning simulator tapes gives

$$Q = \rho Q_G + (1 - \rho) Q_{\overline{G}}, \tag{160}$$

for some distribution $Q_{\overline{G}}$. Thus $Q(E) \geq \rho Q_G(E)$ for every event E , and statistical distance at most ν gives

$$\begin{aligned} R(E) &\leq Q_G(E) + \nu \\ &\leq \rho^{-1} Q(E) + \nu \leq 2^k Q(E) + \nu. \end{aligned} \tag{161}$$

The reverse implication is Theorem 15 with $\varepsilon = \kappa$. $\square$

Fraction knowledge does not control excess simulator mass. Consider

$$R = \delta_a, \quad Q = \frac{1}{2}\delta_a + \frac{1}{2}\delta_b. \quad (162)$$

Selecting tapes outputting a gives density $1/2$ and conditional law R : the pair has exact fraction knowledge one and $R \preceq_{\ln 2, 0} Q$. But $E = \{b\}$ gives

$$Q(E) = \frac{1}{2}, \quad R(E) = 0. \quad (163)$$

Hence $Q \preceq_{\varepsilon, \delta} R$ requires $\delta \geq 1/2$ for every finite ε .

This separates the two inequalities for the *same simulator*, not protocols or language classes: R itself is efficiently samplable and provides a perfect symmetric simulator for this example.

10.3 The class-level converse

A class-level transformation may change both protocol and simulator. Theorem 8 gives

$$\text{MZK}[O(1)] = \text{MZK}[O(\log n)] = \text{SZK}. \quad (164)$$

At density one every simulator tape is good, so

$$\text{SKC}_{\text{frac}}[0] = \text{HVSZK} = \text{SZK}. \quad (165)$$

Consequently,

$$\begin{aligned} & \text{SKC}_{\text{frac}}[O(1)] \subseteq \text{MZK}[O(1)] \\ \iff & \text{SKC}_{\text{frac}}[O(1)] = \text{SKC}_{\text{frac}}[0], \\ & \text{SKC}_{\text{frac}}[O(\log n)] \subseteq \text{MZK}[O(\log n)] \\ \iff & \text{SKC}_{\text{frac}}[O(\log n)] = \text{SKC}_{\text{frac}}[0]. \end{aligned} \quad (166)$$

The forward implication uses (164); the reverse uses $\text{SKC}_{\text{frac}}[0] = \text{MZK}[0] \subseteq \text{MZK}[O(1)]$. The logarithmic case is identical.

Petrank and Tardos ask whether [34]

$$\begin{aligned} & \text{SKC}_{\text{frac}}[1] \stackrel{?}{=} \text{SKC}_{\text{frac}}[0], \\ & \text{SKC}_{\text{frac}}[O(\log n)] \stackrel{?}{=} \text{SKC}_{\text{frac}}[0]. \end{aligned} \quad (167)$$

Sahai and Vadhan also discuss this hierarchy [37]. A general symmetric converse for constant fraction knowledge would therefore resolve an open hierarchy problem. A converse preserving the protocol is stronger and does not follow from equality of language classes.

11 Towards a Symmetric Converse for Knowledge Complexity

Fraction knowledge controls only real-to-simulator mass, and constant forward KL divergence need not imply constant approximate max-divergence with negligible error.

11.1 What fraction knowledge does imply

For exact fraction simulation, $\rho \geq 2^{-k}$ and $Q_G = R$, so (160) gives

$$Q = \rho R + (1 - \rho)Q_{\overline{G}}. \quad (168)$$

Thus $Q(z) \geq \rho R(z)$ pointwise, and

$$D_\infty(R\|Q) := \ln \max_{z: R(z) > 0} \frac{R(z)}{Q(z)} \leq \ln \frac{1}{\rho} \leq k \ln 2. \quad (169)$$

In particular,

$$D_{\text{KL}}(R\|Q) \leq k \ln 2. \quad (170)$$

With statistical error $\text{SD}(R, Q_G) \leq \nu < 1$, up to ν mass of R may lie outside $\text{supp}(Q)$. Smooth R by defining

$$\begin{aligned} \mu(z) &= \min\{R(z), Q_G(z)\}, \\ m &= \sum_z \mu(z) \geq 1 - \nu, \quad R' = \mu/m. \end{aligned} \quad (171)$$

Then $\text{SD}(R, R') \leq \nu$, while

$$Q \geq \rho Q_G \geq \rho \mu = \rho m R'. \quad (172)$$

Consequently,

$$D_\infty(R'\|Q) \leq k \ln 2 + \ln \frac{1}{1 - \nu}, \quad (173)$$

with the same bound for $D_{\text{KL}}(R'\|Q)$. Both bounds are *real-to-simulator*, as in Proposition 3.

11.2 There is no reverse KL bound

The bad-tape distribution $Q_{\overline{G}}$ is unrestricted. For $0 < \rho < 1$, take

$$R = \delta_a, \quad Q = \rho \delta_a + (1 - \rho) \delta_b. \quad (174)$$

Then $D_{\text{KL}}(R\|Q) = \ln(1/\rho)$, but $D_{\text{KL}}(Q\|R) = \infty$. Even replacing R by $R_\gamma = (1 - \gamma)\delta_a + \gamma\delta_b$, where $0 < \gamma < 1$, gives

$$D_{\text{KL}}(Q\|R_\gamma) = \rho \ln \frac{\rho}{1 - \gamma} + (1 - \rho) \ln \frac{1 - \rho}{\gamma}, \quad (175)$$

which diverges as $\gamma \rightarrow 0$, while forward KL remains bounded for fixed ρ . Petrank and Tardos note this bad-mass obstruction in fraction knowledge [34].

11.3 Bounded forward KL does not yield approximate max-divergence

Constant forward KL need not give a constant approximate-max-divergence bound with negligible error. Fix $K > 0$, let $\tau = 1/n$ for $n \geq 2$, and define

$$\begin{aligned} X(a) &= \tau, & Y(a) &= \tau e^{-K/\tau}, \\ X(b) &= 1 - \tau, & Y(b) &= 1 - \tau e^{-K/\tau}. \end{aligned} \quad (176)$$

Their divergence is

$$\begin{aligned} D_{\text{KL}}(X\|Y) &= K + (1 - \tau) \ln \frac{1 - \tau}{1 - \tau e^{-K/\tau}} \\ &= K + O(\tau). \end{aligned} \quad (177)$$

But $X \preceq_{\varepsilon, \delta} Y$, applied to $\{a\}$, requires

$$\delta \geq \tau \left(1 - e^{\varepsilon - K/\tau}\right). \quad (178)$$

For constant ε , this is eventually at least $1/(2n)$. Conversely, negligible δ forces $\varepsilon \geq K/\tau - o(1) = \Omega(n)$.

Likewise, $\text{SD}(X, X') < \tau/2$ implies $X'(a) > \tau/2$, so

$$D_\infty(X' \| Y) \geq \ln \frac{\tau/2}{\tau e^{-K/\tau}} = \frac{K}{\tau} - \ln 2 = \Omega(n). \quad (179)$$

Every negligible smoothing function is eventually smaller than $1/(2n)$, so smoothing does not remove the obstruction.

For an exact uniform finite-coin example, set $K = \ln 2$ and $\tau = 2^{-\lceil \log_2 n \rceil}$ in (176). Then $Y(a) = \tau 2^{-1/\tau}$; strict polynomial-time samplers use $O(n + \log n)$ bits. Since $\tau = \Theta(1/n)$, the same calculations give bounded KL and an $\Omega(n)$ approximate-max-divergence lower bound.

11.4 Consequence for the proposed converse

Fraction knowledge supplies only (157). Equations (174) and (178) rule out a distribution-free converse through KL divergence. We next obtain a symmetric converse by assuming that one online prover generates the simulator's full view distribution.

12 When a Prover Can Generate the Simulator's Views

12.1 The prover and the product protocol

Let $\mathcal{Y}$ be the set of yes-instances and $Q_x = S(x)$ the unconditioned simulator output. Assume there is one uniform, possibly unbounded online prover $\hat{P}$ such that

$$Q_x = \text{View}_{\hat{P}, V}(x) \quad \text{for every } x \in \mathcal{Y}. \quad (180)$$

The quantifier order is $\exists \hat{P} \forall x \in \mathcal{Y}$. Each message uses only the common input, the observed interaction prefix, and the prover's own private randomness. Valid transcripts alone do not imply this condition. For example, if the prover sends a bit before the verifier tosses a fresh fair coin, an offline sampler can always make the two bits equal. Those transcripts are valid, but an online prover can match the later coin only with probability $1/2$.

For a polynomial-time computable integer $t = t(n)$, define $\Pi^{(t)} = (P^{(t)}, V^{(t)})$: the verifier completes all t sequential copies of V with independent tapes and then accepts their OR. The prescribed prover chooses a private $I \leftarrow [t]$, uses P in copy I and $\hat{P}$ elsewhere, with independent randomness and no dependence on other copies' transcripts. The simulator runs the original strict-PPT S independently t times, yielding $Q_x^{\otimes t}$. We identify the repeated verifier's full view with the tuple of copy views; the final OR is determined by that tuple.

Proposition 4 (Symmetrization from an online prover). *Let $\Pi = (P, V)$ have completeness error $c(n)$ and soundness error $s(n)$. Suppose a strict-PPT simulator S gives fraction knowledge $k(n)$ with statistical error $\nu(n)$, and one uniform, possibly unbounded online prover $\hat{P}$ generates its full unconditioned output distribution on every yes-input, as in (180). If $k(n) = O(\log n)$ and c, s, ν are*

negligible, then a polynomial t makes $\Pi^{(t)}$ honest-verifier symmetric multiplicative zero knowledge with exponent $O(k+1)$ and negligible error. Specifically, for all sufficiently large n , put

$$\bar{\kappa}(n) = k(n) \ln 2 + \ln 2, \quad \eta(n) = \bar{\kappa}(n) + 1, \quad (181)$$

and choose any polynomial-time computable, polynomially bounded integer-valued $t(n)$ satisfying

$$t(n) \geq n^2 e^{2\bar{\kappa}(n)}. \quad (182)$$

Then the repeated real view $T_{t,x}$ and simulator $Q_x^{\otimes t}$ obey

$$T_{t,x} \approx_{\eta, \Delta} Q_x^{\otimes t}, \quad (183)$$

where

$$\Delta(n) = e^{\eta(n)} \nu(n) + \exp\left(-\frac{2t(n)(1 - e^{-\eta(n)})^2}{e^{2\bar{\kappa}(n)}}\right) \quad (184)$$

is negligible. The repeated protocol has completeness error at most $c(n)$ and soundness error at most $t(n)s(n)$.

12.2 Smoothing the good conditional distribution

Proof. Fix a yes-input x of sufficiently large length that $\nu(n) \leq 1/2$, and abbreviate its real view by $R = \text{View}_{P,V}(x)$. Let G be the simulator's good-tape event, so

$$\rho = \Pr[G] \geq 2^{-k}, \quad \text{SD}(R, Q_G) \leq \nu. \quad (185)$$

Partitioning simulator tapes gives

$$Q = \rho Q_G + (1 - \rho) Q_{\bar{G}}, \quad Q \geq \rho Q_G \quad (186)$$

pointwise. Define

$$\begin{aligned} \mu(z) &= \min\{R(z), Q_G(z)\}, \\ m &= \sum_z \mu(z), \\ R'(z) &= \mu(z)/m. \end{aligned} \quad (187)$$

The overlap formula for statistical distance gives $m \geq 1 - \nu$ and

$$\text{SD}(R, R') \leq 1 - m \leq \nu. \quad (188)$$

Moreover,

$$Q \geq \rho Q_G \geq \rho \mu = \rho m R'. \quad (189)$$

Thus R' is supported on $\text{supp}(Q)$ and

$$R' \leq e^\kappa Q, \quad \kappa := \ln \frac{1}{\rho m} \leq k \ln 2 + \ln \frac{1}{1 - \nu}. \quad (190)$$

Eventually $\nu \leq 1/2$, so $\kappa \leq \bar{\kappa}$ from (181).

12.3 Random coordinate embedding

Conditioning on the selected coordinate gives the real view

$$T_t = \frac{1}{t} \sum_{i=1}^t Q^{\otimes(i-1)} \otimes R \otimes Q^{\otimes(t-i)}. \quad (191)$$

Replace R by R' to obtain

$$T'_t = \frac{1}{t} \sum_{i=1}^t Q^{\otimes(i-1)} \otimes R' \otimes Q^{\otimes(t-i)}. \quad (192)$$

By convexity and invariance under tensoring with a common distribution,

$$\text{SD}(T_t, T'_t) \leq \text{SD}(R, R') \leq \nu. \quad (193)$$

Each summand replaces only one coordinate, so averaging keeps the smoothing error at ν .

For $Q(z) > 0$, set $L(z) = R'(z)/Q(z)$, and set $L(z) = 0$ otherwise. Equation (190) implies

$$0 \leq L \leq e^\kappa, \quad \mathbb{E}_{z \leftarrow Q}[L(z)] = 1. \quad (194)$$

For every point in the support of $Q^{\otimes t}$,

$$\frac{T'_t(z_1, \dots, z_t)}{Q^{\otimes t}(z_1, \dots, z_t)} = \frac{1}{t} \sum_{i=1}^t L(z_i). \quad (195)$$

Since this average is at most e^κ , every event E satisfies

$$T'_t(E) \leq e^\kappa Q^{\otimes t}(E). \quad (196)$$

For the reverse direction, let

$$B_\eta = \left\{ (z_1, \dots, z_t) : \frac{1}{t} \sum_{i=1}^t L(z_i) < e^{-\eta} \right\}. \quad (197)$$

Under $Q^{\otimes t}$, the $L(z_i)$ are independent, bounded by e^κ , and have mean one. Hoeffding's inequality gives

$$\Pr_{Q^{\otimes t}}[B_\eta] \leq \beta_t(\eta, \kappa) := \exp\left(-\frac{2t(1 - e^{-\eta})^2}{e^{2\kappa}}\right). \quad (198)$$

On B_η^c , the likelihood ratio is at least $e^{-\eta}$, so

$$\begin{aligned} Q^{\otimes t}(E) &\leq Q^{\otimes t}(E \cap B_\eta^c) + \beta_t(\eta, \kappa) \\ &\leq e^\eta T'_t(E) + \beta_t(\eta, \kappa). \end{aligned} \quad (199)$$

Transferring both inequalities across (193) gives, for every event E ,

$$T_t(E) \leq e^\kappa Q^{\otimes t}(E) + \nu, \quad (200)$$

$$Q^{\otimes t}(E) \leq e^\eta T_t(E) + e^\eta \nu + \beta_t(\eta, \kappa). \quad (201)$$

Since $\eta \geq \bar{\kappa} \geq \kappa$ and $e^\eta \nu + \beta_t \geq \nu$, these prove (183) with uniform error bound (184).

12.4 Negligibility, completeness, and soundness

Since $\eta = \bar{\kappa} + 1$, we have $1 - e^{-\eta} \geq 1 - e^{-1}$. By (182), the second error term is at most

$$\exp(-2(1 - e^{-1})^2 n^2), \quad (202)$$

which is negligible. Because $k = O(\log n)$, $e^\eta = e^{\bar{\kappa}+1} = 2e 2^k$ and the required t are polynomial, so $e^\eta \nu$ is negligible as well.

On yes-instances, rejection requires copy I , which uses P , to reject; completeness error is at most c . On no-instances, a single-copy cheating prover can internally simulate all preceding copies and then continue with the resulting state. This gives the same marginal strategy in that copy even when the repeated prover correlates its actions across copies. The simulated history is independent of the current verifier's fresh tape, so soundness bounds each copy's acceptance probability by s . A union bound gives negligible soundness error ts . $\square$

12.5 What the argument assumes

Proposition 4 requires one online prover to generate the simulator's entire unconditioned distribution on every yes-input. Fraction simulation alone does not supply such a prover, so the general symmetric converse remains open. Replacing Q by a prover view within distance σ can cost $t\sigma$ under products; this additional error must remain negligible.

13 Conclusion

With negligible proof and additive simulation errors, logarithmic multiplicative leakage preserves the language class: $\text{HVMZK}[O(\log n)] = \text{MZK}[O(\log n)] = \text{SZK}$. This threshold is tight relative to oracles: every polynomial-time computable, polynomially bounded superlogarithmic profile admits a separation even from quantum SZK . Unrelativized, the stated subexponential nonuniform hypothesis separates every fixed polylogarithmic level of degree at least two from SZK , and the union over polynomial leakage bounds equals PSPACE . Unconditional unrelativized separations at these polylogarithmic levels, and strictness between successive levels, remain open.

Fraction statistical knowledge complexity exactly characterizes one-sided multiplicative simulation. Its symmetric converse remains tied to the knowledge-complexity collapse question. Our canonical reduction preserves fixed protocol profiles at a unary source-length tag; an ordinary fixed-profile complete problem at total encoding length remains a separate question. Finally, Appendix A shows that constant multiplicative leakage can reduce communication for the same distributed task even though the corresponding language-class regime collapses.

14 Acknowledgements

We thank Noam Mazon, Salil Vadhan, and Yupeng Zhang for helpful discussions. This project began while Daniel Alabi was supported by the Simons Foundation (965342, D.A.) as part of the Junior Fellowship from the Simons Society of Fellows. Marshall Ball was supported in part by the U.S. National Science Foundation through CAREER Award No. 2443735 and Award No. 2552209, the latter under the joint NSF–United States–Israel Binational Science Foundation (BSF) program. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- [1] William Aiello and Johan Håstad. Statistical zero-knowledge languages can be recognized in two rounds. *Journal of Computer and System Sciences*, 42(3):327–345, June 1991.
- [2] S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2006.
- [3] L Babai. Trading group theory for randomness. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, STOC '85, page 421–429, New York, NY, USA, 1985. Association for Computing Machinery.
- [4] László Babai and Shlomo Moran. Arthur-merlin games: a randomized proof system, and a hierarchy of complexity class. *J. Comput. Syst. Sci.*, 36(2):254–276, April 1988.
- [5] Michael Backes, Aniket Kate, Sebastian Meiser, and Tim Ruffing. Secrecy without perfect randomness: Cryptography with (bounded) weak sources. In *Applied Cryptography and Network Security (ACNS 2015)*, volume 9092 of *Lecture Notes in Computer Science*, pages 675–695. Springer, 2015.
- [6] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *Journal of Computer and System Sciences*, 68(4):702–732, 2004.
- [7] Amos Beimel, Kobbi Nissim, and Eran Omri. Distributed private data analysis: On simultaneously solving how and what, 2011. Full version of the CRYPTO 2008 paper, pages 451–468.
- [8] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, pages 1–10. ACM, 1988.
- [9] Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Statistical difference beyond the polarizing regime. In *Theory of Cryptography – TCC 2019, Part II*, volume 11892 of *Lecture Notes in Computer Science*, pages 311–332. Springer, 2019.
- [10] Ravi B. Boppana, Johan Håstad, and Stathis Zachos. Does co-NP have short interactive proofs? *Information Processing Letters*, 25(2):127–132, May 1987.
- [11] Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, ITCS '16, pages 261–270. ACM, 2016.
- [12] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory 2nd Edition (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, July 2006.
- [13] Imre Csiszár. Eine informationstheoretische ungleichung und ihre anwendung auf den beweis der ergodizität von markoffschen ketten. *A Magyar Tudományos Akadémia Matematikai Kutató Intézetének Közleményei*, 8(1–2):85–108, 1963.
- [14] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *EUROCRYPT*, pages 486–503, 2006.

- [15] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, pages 265–284, 2006.
- [16] Christopher A. Fuchs and Jeroen van de Graaf. Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Transactions on Information Theory*, 45(4):1216–1227, May 1999.
- [17] Oded Goldreich and Johan Håstad. On the complexity of interactive proofs with bounded communication. *Information Processing Letters*, 67(4):205–214, 1998.
- [18] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to prove all np statements in zero-knowledge and a methodology of cryptographic protocol design (extended abstract). In Andrew M. Odlyzko, editor, *Advances in Cryptology — CRYPTO’ 86*, pages 171–185, Berlin, Heidelberg, 1987. Springer Berlin Heidelberg.
- [19] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in np have zero-knowledge proof systems. *J. ACM*, 38(3):690–728, July 1991.
- [20] Oded Goldreich, Rafail Ostrovsky, and Erez Petrank. Computational complexity and knowledge complexity. *SIAM Journal on Computing*, 27(4):1116–1141, August 1998.
- [21] Oded Goldreich and Erez Petrank. Quantifying knowledge complexity. *Computational Complexity*, 8(1):50–98, 1999.
- [22] Oded Goldreich, Amit Sahai, and Salil Vadhan. Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing, STOC ’98*, pages 399–408, New York, NY, USA, 1998. ACM.
- [23] Oded Goldreich, Salil Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *Computational Complexity*, 11(1–2):1–53, 2002.
- [24] Oded Goldreich and Salil P. Vadhan. Comparing entropies in statistical zero knowledge with applications to the structure of SZK. In *Proceedings of the 14th Annual IEEE Conference on Computational Complexity*, pages 54–73. IEEE Computer Society, 1999.
- [25] S Goldwasser, S Micali, and C Rackoff. The knowledge complexity of interactive proof-systems. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing, STOC ’85*, page 291–304, New York, NY, USA, 1985. Association for Computing Machinery.
- [26] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- [27] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, February 1989.
- [28] Thomas Holenstein and Renato Renner. One-way secret-key agreement and applications to circuit polarization and immunization of public-key encryption. In *Advances in Cryptology – CRYPTO 2005*, Lecture Notes in Computer Science, pages 478–493. Springer, 2005.
- [29] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalii. *Classical and Quantum Computation*. American Mathematical Society, USA, 2002.
- [30] Sanketh Menda and John Watrous. Oracle separations for quantum statistical zero-knowledge. arXiv:1801.08967, 2018. Manuscript.

- [31] Tetsuzo Morimoto. Markov processes and the h-theorem. *Journal of the Physical Society of Japan*, 18(3):328–331, 1963.
- [32] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [33] Tatsuaki Okamoto. On relationships between statistical zero-knowledge proofs. *Journal of Computer and System Sciences*, 60(1):47–108, 2000.
- [34] Erez Petrank and Gábor Tardos. On the knowledge complexity of NP. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pages 494–503. IEEE Computer Society, 1996.
- [35] Alfréd Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, pages 547–561, Berkeley, CA, 1961. University of California Press.
- [36] Guy N. Rothblum, Eran Omri, Junye Chen, and Kunal Talwar. PINE: Efficient verification of a euclidean norm bound of a Secret-Shared vector. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 6975–6992, Philadelphia, PA, USA, August 2024. USENIX Association.
- [37] Amit Sahai and Salil P. Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM*, 50(2):196–249, March 2003.
- [38] Igal Sason and Sergio Verdú. f -divergence inequalities. *IEEE Transactions on Information Theory*, 62(11):5973–6006, 2016.
- [39] Adi Shamir. $IP = PSPACE$. *Journal of the ACM*, 39(4):869–877, October 1992.
- [40] Larry J. Stockmeyer and Albert R. Meyer. Word problems requiring exponential time: Preliminary report. In *Proceedings of the Fifth Annual ACM Symposium on Theory of Computing*, pages 1–9, New York, NY, USA, 1973. Association for Computing Machinery.
- [41] Kunal Talwar. Differential secrecy for distributed data and applications to robust differentially secure vector summation. In *3rd Symposium on Foundations of Responsible Computing (FORC 2022)*, volume 218 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.
- [42] Salil P. Vadhan. *A Study of Statistical Zero-Knowledge Proofs*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, 1999.
- [43] Leslie G. Valiant. Relative complexity of checking and evaluating. *Information Processing Letters*, 5(1):20–23, May 1976.
- [44] Thomas Vidick and Stephanie Wehner. *Introduction to Quantum Cryptography*. Cambridge University Press, 2023.
- [45] John Watrous. Limits on the power of quantum statistical zero-knowledge. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 459–468. IEEE Computer Society, 2002.

A A Communication Separation for Distributed Simulation

Multiplicative simulation can unconditionally reduce communication for a fixed randomized distributed functionality. Parties here compute a common output from private inputs; this result concerns a different model from the public-input MZK and SZK proof classes. The proof combines the network method of Beimel, Nissim, and Omri [7] with the Hellinger cut-and-paste identity [6, Lemma 6.3].

Model and simulation requirement. There are $m \geq 4$ parties with inputs $x_i \in \{0, 1\}$, private point-to-point channels, and no correlated setup. Their private random tapes are independent of the inputs and of one another. A static passive coalition T follows the protocol and pools its views. A fixed finite public schedule specifies each message's order, sender, recipient, and positive bit length, independently of inputs and randomness. Communication counts all these bits; timing, silence, lengths, and termination convey no additional information. The lower bound permits arbitrary rounds, local computation, and finite private randomness.

Set $t = \lfloor (m-1)/3 \rfloor$. All parties agree almost surely on a common output O ; correctness error below concerns its distribution. The intended functionality is

$$F_m(x) = \text{Bern}\left(\frac{1}{4} + \frac{\sum_{i=1}^m x_i}{2m}\right). \quad (203)$$

Write $V_T(x)$ for the coalition's inputs, complete private random tapes, incident messages, and local outputs. Multiplicative simulation means that one simulator, given only $(1^m, T, x_T, Y)$, satisfies

$$(V_T(x), O) \approx_{\varepsilon, \delta} (S(1^m, T, x_T, Y), Y), \quad Y \leftarrow F_m(x), \quad (204)$$

for all x and $|T| \leq t$. The same simulator must handle all honest inputs given only x_T and the supplied ideal output. Statistical simulation replaces $\approx_{\varepsilon, \delta}$ by statistical distance at most η . Exact correctness means $O \sim F_m(x)$; statistical simulation permits at most η error in this output law.

Theorem 16 (Distributed communication separation). *For every power of two $m \geq 4$:*

1. *There is an exactly correct protocol for F_m using*

$$C_{\text{mult}} = (m-1)(\log_2 m + 2)$$

bits, with uniform strict-PPT $(\ln 3, 2^{-m})$ simulators against all passive coalitions of size at most t .

2. *Every protocol satisfying statistical simulation with $\eta \leq 1/(100m^2)$ communicates at least $m(t+1)/2 = \Omega(m^2)$ bits, even with computationally unbounded simulators.*

The communication advantage is $\Omega(m/\log m)$ for the same functionality, even allowing η correctness error in the statistical protocol's common output law.

Proof. The star protocol. The center, party 1, samples $J \leftarrow [m]$ and sends its $\log_2 m$ -bit encoding to every leaf. Each party forms a report B_i : party J reports x_J with probability $3/4$ and $1 - x_J$ otherwise; other reports are fair dummy bits. Each leaf sends its report to the center, which forms its own report locally and sends $O = B_J$ to every leaf. All parties output O . Two independent private bits per party suffice for the reports, with one unused for dummy reports. The stated cost follows, and

$$\Pr[O = 1] = \frac{1}{m} \sum_{j=1}^m \left(\frac{1}{4} + \frac{x_j}{2} \right).$$

Fix T and x . A neutral experiment D samples J and the corrupted tapes jointly as in the protocol: a corrupted center's index coins determine J . Corrupted reports use x_T and these tapes as prescribed; honest reports are independent fair bits. The output is B_J , so D is efficiently samplable from (m, T, x_T) .

Compare the real and neutral laws P_x, D of $W = (J, B_1, \dots, B_m, r_T)$, retaining all corrupted tapes r_T and marginalizing honest tapes. The coalition view and output are determined by W and x_T . If $J \in T$, the laws agree; otherwise only the selected honest report changes probability from $1/2$ to $1/4$ or $3/4$. Thus, pointwise,

$$\frac{1}{2}D(w) \leq P_x(w) \leq \frac{3}{2}D(w). \quad (205)$$

Summing over either output $B_J = y$ and normalizing gives

$$\frac{1}{3}D(w \mid B_J = y) \leq P_x(w \mid B_J = y) \leq 3D(w \mid B_J = y). \quad (206)$$

An ideal sampler S_∞ returns the coalition view under $D \mid (B_J = y)$. Averaging (206) with the real output probabilities $F_m(x)(y)$ proves pure $\ln 3$ simulation in (204).

After averaging the sampled tapes, the neutral output conditional on J is fair if $J \notin T$ and is a $3/4$ -randomized response to x_J if $J \in T$. Each neutral output therefore has probability at least $1/4$. Cap rejection sampling at $8m$ independent trials and return a fixed placeholder on failure. For each supplied y , statistical distance from S_∞ is at most $q_m = (3/4)^{8m}$. The two event inequalities lose at most $3q_m$ and q_m , respectively. Since $3q_m \leq 2^{-m}$, this gives the uniform strict-PPT simulator.

The statistical lower bound. Let G be the protocol's undirected active-channel graph. If $1 \leq \deg_G(i) \leq t$, corrupt $T = N_G(i)$ and choose $j \notin T \cup \{i\}$, possible since $t \leq m - 2$. Fix other inputs to zero and let P_{ab} be the law of V_T on $(x_i, x_j) = (a, b)$. Party i has no channel to the remaining honest parties $H = [m] \setminus (T \cup \{i\})$.

Fix a coalition view v , including tapes and outputs. Consistency with this boundary transcript imposes separate constraints on the tapes of i and H , the latter including internal communication within H . The corrupted parties' messages and outputs are determined by their fixed inputs, tapes, and received messages in v . Independent initial tapes therefore give nonnegative factors

$$P_{ab}(v) = \alpha_a(v)\beta_b(v)\gamma(v), \quad P_{00}(v)P_{11}(v) = P_{01}(v)P_{10}(v). \quad (207)$$

Here α_a and β_b are the respective probabilities of the honest tape constraints; γ includes the corrupted-tape probability and consistency of the corrupted parties' actions. Including T 's local output in v does not condition on an external ideal output.

Define $h^2(P, Q) = 1 - \sum_v \sqrt{P(v)Q(v)}$. Equation (207) gives the exact cut-and-paste identity $h^2(P_{00}, P_{11}) = h^2(P_{01}, P_{10})$. Inputs 01 and 10 have identical ideal output laws and x_T , hence identical simulated views. Thus $\text{SD}(P_{01}, P_{10}) \leq 2\eta$. Using $h^2 \leq \text{SD} \leq \sqrt{2h^2}$ yields

$$\text{SD}(P_{00}, P_{11}) \leq 2\sqrt{\eta}. \quad (208)$$

Because T is nonempty, its view determines the common output. The ideal success probabilities under 00 and 11 differ by $1/m$; η correctness and data processing give

$$\frac{1}{m} \leq 2\eta + \text{SD}(P_{00}, P_{11}) \leq 2\eta + 2\sqrt{\eta} < \frac{1}{m}, \quad (209)$$

where the last inequality uses $\eta \leq 1/(100m^2)$.

An isolated party's output law cannot depend on other inputs, yet changing one other bit changes the ideal probability by $1/(2m) > 2\eta$. Thus every vertex has degree at least $t + 1$, yielding $m(t + 1)/2$ active edges, each carrying at least one bit. $\square$

Statistical implementation is feasible by unconditional secure computation with an honest majority [8]. Evaluate a fixed Boolean sampler for F_m with jointly generated secret random bits and a fixed communication pattern. XORing all parties' private uniform strings inside the secure computation supplies a uniform sampler seed even conditioned on the corrupted parties' strings.

Sublogarithmic leakage and aggregation. Take the common-output vector functionality $F_m(x)^{\otimes k}$ and repeat the star protocol independently $k = \lfloor \sqrt{\log_2 m} \rfloor$ times. The product of uncapped simulators gives pure leakage $k \ln 3$; capping changes it by at most kq_m in statistical distance. Hence

$$\varepsilon_k = k \ln 3 = o(\log m), \quad \delta_k \leq 3^k kq_m = 2^{-\Omega(m)}, \quad C_k = O(m(\log m)^{3/2}). \quad (210)$$

For statistical simulation with the same bound $\eta \leq 1/(100m^2)$, the lower bound remains $\Omega(m^2)$: the 01 and 10 ideal output-vector laws still agree, and the first output coordinate retains the $1/m$ gap. For $s = \sum_i x_i$, the following estimator is unbiased and, by independence, has variance at most m^2/k :

$$\hat{s} = 2m \left(\frac{1}{k} \sum_{a=1}^k Y_a - \frac{1}{4} \right) \quad \text{satisfies} \quad \mathbb{E}|\hat{s} - s| \leq \sqrt{\text{Var}(\hat{s})} \leq \frac{m}{\sqrt{k}} = o(m).$$

The lower bound permits arbitrary redesign of the statistical protocol and simulator within this distributed, passive-security, fixed-pattern model.

Statement on AI Use

The authors used ChatGPT, including GPT-5.6 Sol, to generate initial proof drafts for some results and to assist with exposition and technical review. The authors revised the proof drafts and checked the mathematical statements. They take responsibility for the results, proofs, citations, and final presentation. AI assistance does not constitute independent verification of the claims.