

A Uniform Min-Max Theorem with Applications in Cryptography*

Salil Vadhan Colin Jia Zheng
School of Engineering and Applied Sciences
Harvard University
Cambridge, Massachusetts
{salil,colinz}@seas.harvard.edu

Abstract

We present a new, more constructive proof of von Neumann’s Min-Max Theorem for two-player zero-sum game — specifically, an algorithm that builds a near-optimal mixed strategy for the second player from several best-responses of the second player to mixed strategies of the first player. The algorithm extends previous work of Freund and Schapire (Games and Economic Behavior ’99) with the advantage that the algorithm runs in $\text{poly}(n)$ time even when a pure strategy for the first player is a distribution chosen from a set of distributions over $\{0,1\}^n$. This extension enables a number of additional applications in cryptography and complexity theory, often yielding uniform security versions of results that were previously only proved for nonuniform security (due to use of the non-constructive Min-Max Theorem).

We describe several applications, including: a more modular and improved uniform version of Impagliazzo’s Hardcore Theorem (FOCS ’95); regularity theorems that provide efficient simulation of distributions within any sufficiently nice convex set (extending a result of Trevisan, Tulsiani and Vadhan (CCC ’09)); an improved version of the Weak Regularity Lemma of Frieze and Kannan; a Dense Model Theorem for uniform algorithms; and showing impossibility of constructing Succinct Non-Interactive Arguments (SNARGs) via black-box reductions under uniform hardness assumptions (using techniques from Gentry and Wichs (STOC ’11) for the nonuniform setting).

Keywords: Min-Max Theorem, multiplicative weights, KL projection, indistinguishability, Hardcore Theorem, Succinct Non-Interactive Argument, efficient simulation

*Supported by NSF grant CCF-1116616 and US-Israel BSF grant 2010196.

1 Introduction

Von Neumann’s Min-Max Theorem (or Linear Programming Duality, finite-dimensional Hahn-Banach Theorem) has proved to be an extremely useful tool in theoretical computer science. Consider a zero-sum game between two players where for every mixed strategy V for Player 1 (as a distribution over his strategy space \mathcal{V}), Player 2 has a response $W \in \mathcal{W}$ that guarantees $\mathbb{E}[F(V, W)] \geq 0$, where F (payoff) can be an arbitrary function. The Min-Max Theorem says that there must exist a Player 2’s mixed strategy W^* (as a distribution over his strategy space \mathcal{W}) that guarantees $\mathbb{E}[F(V, W^*)] \geq 0$ for *all* strategies $V \in \mathcal{V}$ of Player 1.

The Min-Max Theorem gives rise to a number of results in cryptography and complexity theory such as Impagliazzo’s Hardcore Theorem [Imp], equivalence of different notions of computational entropy [BSW], the Dense Model Theorem [RTTV], leakage-resilient cryptography [DP2, FR], efficient simulation of high entropy distributions [TTV], impossibility of constructing succinct non-interactive arguments (SNARGs) via black-box reductions [GW], and simple construction of pseudorandom generators from one-way functions [VZ]. In a typical application like these, Player 1 chooses V from a convex set \mathcal{V} of distributions over $\{0, 1\}^n$, and Player 2 chooses W from a set \mathcal{W} of (possibly randomized) boolean functions $\{0, 1\}^n \rightarrow \{0, 1\}$ and receives payoff $F(V, W) = \mathbb{E}[W(V)]$ i.e. function W ’s expected output when input is drawn from the distribution V . For example, \mathcal{V} contains all high entropy distributions over $\{0, 1\}^n$ and \mathcal{W} contains all boolean functions of small circuit size.

A limitation of the Min-Max Theorem is that it is highly non-constructive; it only asserts the existence of the optimal strategy W^* but does not say how it can be found (algorithmically). Consequently, applications of the Min-Max Theorem only give rise to results about nonuniform boolean circuits, rather than uniform algorithms (e.g. we set cryptographic protocols based on nonuniform hardness rather than uniform hardness assumptions).

To overcome this, we consider the natural algorithmic task of constructing such an optimal strategy W^* for Player 2, given an efficient algorithm for F . When the sizes of strategy spaces \mathcal{V} and \mathcal{W} are small (e.g. polynomial) this can be done by linear programming, for which efficient algorithms are well-known. However, applications in cryptography and complexity theory such as ones just mentioned involve exponentially large strategy spaces, and an optimal strategy W^* cannot be found in polynomial time in general. Thus we also require that, given any mixed strategy V for Player 1, not only does there exist a strategy $W \in \mathcal{W}$ for Player 2 with $\mathbb{E}[F(V, W)] \geq 0$, but such response W can be obtained efficiently by an oracle (or an efficient uniform algorithm).

Assuming such an oracle, Freund and Schapire [FS] show how to find an approximately optimal W^* for Player 2 in polynomial time and by making $O((\log |\mathcal{V}|)/\epsilon^2)$ adaptive oracle queries, using the idea of multiplicative weight updates. However, their algorithm still falls short in some of aforementioned applications where \mathcal{V} is a set of distributions over $\{0, 1\}^n$, and thus \mathcal{V} can have doubly-exponentially many vertices. For example, consider the set of distributions on $\{0, 1\}^n$ of min-entropy at least k ; the vertices of \mathcal{V} are uniform distributions on a subset of size 2^k , and there are $\binom{2^n}{2^k}$ such subsets.

We present a Uniform Min-Max Theorem that efficiently finds an approximately optimal strategy W^* for Player 2, given an oracle that for any of Player 1’s mixed strategy $V \in \mathcal{V}$ returns some Player 2’s strategy that guarantees reasonable payoff, even when \mathcal{V} is a (sufficiently nice) set of distributions over $\{0, 1\}^n$. Our theorem is inspired by the proof of Uniform Hardcore Theorem of Barak, Hardt, and Kale [BHK]. Like [BHK], the underlying algorithm uses “relative entropy (KL) projections” together with multiplicative weight updates (a technique originally due to Herbster

and Warmuth [HW]). Our contribution is providing the right abstraction: formulating this algorithm as providing a Uniform Min-Max Theorem. An advantage of this formulation is that it is more modular, and not specific to the Hardcore Theorem. Consequently, it immediately enables a number of applications, including (but not limited to) deriving uniform versions of many of the aforementioned results, where we now deal with algorithms rather than nonuniform boolean circuits. Even for the Hardcore Theorem, where the uniform version was already known [Hol1, BHK], there are several advantages to deducing it using the Uniform Min-Max Theorem. Furthermore, even in nonuniform settings, replacing the use of standard Min-Max Theorem with the Uniform Min-Max Theorem can often lead to improved, optimal parameters.

Uniform Hardcore Theorem. Impagliazzo’s Hardcore Theorem ([Imp] and later strengthened in [KS, Hol1, BHK]) is a fundamental result in complexity theory that says if a boolean function f is somewhat hard on average, then there must be a subset of inputs (the hardcore) on which f is extremely hard, and outside of which f is easy. There are two approaches to proving the theorem. One is constructive [Imp, KS, Hol1, BHK] and leads to a *Uniform Hardcore Theorem* where hardness of f is measured against uniform algorithms, rather than nonuniform boolean circuits, and has found several applications in cryptography [KS, Hol1, Hol2, HHR, HRV]. However, the existing proofs turn out to be adhoc and do not achieve all of the optimal parameters simultaneously for a Uniform Hardcore Theorem. Another approach due to Nisan [Imp] (and strengthened in [Hol1]) uses the (non-constructive) Min-Max Theorem and has the advantage of simplicity, but is restricted to the nonuniform measure of hardness.

In Section 4, we show that by replacing the use of Min-Max Theorem in the proof of Nisan [Imp] or Holenstein [Hol1] with our Uniform Min-Max Theorem, we obtain a new proof of the Uniform Hardcore Theorem with the advantages of (i) optimal hardcore density; (ii) optimal complexity blow-up; and (iii) modularity and simplicity.

Construction of Pseudorandom Generators from One-Way Functions. Recently, we [VZ] obtained a simplified and more efficient construction of pseudorandom generators from arbitrary one-way functions, building on the work of Haitner, Reingold, and Vadhan [HRV]. Key to the simplification is a new characterization of a computational analogue of Shannon entropy, whose proof in the nonuniform setting involves the Min-Max Theorem. Using the Uniform Min-Max Theorem instead, we proved our characterization of pseudoentropy in the uniform setting, and hence obtain (simpler) pseudorandom generator from arbitrary one-way functions that are secure against efficient algorithms. See Section 5 for a more detailed discussion.

Regularity Theorems for Distributions Restricted to a Convex Set We apply the Uniform Min-Max Theorem to show a generalization and quantitative improvement to the “regularity theorem” of Trevisan, Tulsiani, and Vadhan [TTV] which (informally) says that any high min-entropy distribution X is indistinguishable from some high min-entropy, *low complexity* distribution Y . The result of [TTV] is itself a quantitative improvement of regularity and “decomposition” theorems in additive combinatorics [GT, TZ]. It is shown in [TTV] that such results can be used to deduce the Dense Model Theorem [TZ, RTTV, Gow], Impagliazzo’s Hardcore Theorem [Imp], and other results, by replacing any unknown distribution X with an “equivalent” distribution Y that can be efficiently analyzed and manipulated.

Our result is more general than [TTV] in the sense that we are no longer restricted to distributions of high min-entropy. We show that for any sufficiently nice convex set of distributions \mathcal{V} , every distribution $X \in \mathcal{V}$ is indistinguishable from some distribution $Y \in \mathcal{V}$ where Y has “low complexity”, for various notions of complexity and indistinguishability. In the case of min-entropy distributions, we obtain a high min-entropy Y with lower complexity than [TTV]. This also yields an improved and optimal Weak Regularity Lemma for graphs of density $o(1)$ (Section 6.2).

Average-case versions of our regularity theorems can be used to deduce “low complexity” versions of a technical lemma of [GW]. We note that our average-case regularity theorem for circuit complexity is a strengthening of a recent result of Pietrzak and Jethchev [PJ], with a simpler proof. The low circuit complexity version of the [GW] lemma (with slightly weaker parameters) was initially proved by Pietrzak and Jethchev [PJ], and an interactive extension was proved by Chung, Lui, and Pass [CLP] for applications in the context of distributional zero-knowledge.

Uniform Dense Model Theorem. A celebrated result of Green and Tao [GT] shows that there exist arbitrarily long arithmetic progressions of prime numbers. A key new component of their proof is the Dense Model Theorem which, in the generalized form of Tao and Ziegler [TZ], says if X is a pseudorandom distribution and D is a distribution dense in X , then D is indistinguishable to a distribution M that is dense in the uniform distribution. Using the Min-Max Theorem, Reingold et al. [RTTV] provided another proof of Dense Model Theorem where the indistinguishability and complexity blow-ups are polynomial (rather than exponential); a similar proof was given by Gowers [Gow]. The polynomial blow-ups are crucial for applications in leakage-resilient cryptography [DP2, DP1, FOR], and for connections to computational differential privacy [MPRV]. Using the Uniform Min-Max Theorem, we show how to obtain a Dense Model Theorem where the distinguishers are efficient (uniform) algorithms, with polynomial blow-ups in running time and indistinguishability.

Impossibility of Black-Box Construction of Succinct Non-interactive Argument. A result of Gentry and Wichs [GW] shows that there is no black-box construction of succinct non-interactive arguments (SNARGs) from any natural cryptographic assumption. Their result relies on the (mild) assumption that there exist *hard subset membership problems*, which is equivalent to the existence of subexponentially hard one-way functions. One limitation is that they need to assume nonuniformly secure one-way functions, in part due to their use of the non-constructive Min-Max theorem (in [GW] Lemma 3.1).

In Section 8, we show how to obtain the analogous result in the *uniform setting* by using the Uniform Min-Max Theorem. More specifically, assuming that there exist subexponentially hard one-way functions that are secure against uniform algorithms, we show that there is no construction of SNARGs whose security can be reduced in a black-box way to a cryptographic assumption against uniform algorithms (unless the assumption is already false).

1.1 Paper Organization

Basic notions from information theory including KL projection are defined in Section 2. In Section 3 we state and prove the Uniform Min-Max Theorem, and show that it also implies the standard Min-Max Theorem. In Section 4, 5, 6, 7, 8, we describe a number of applications of the Uniform Min-Max Theorem.

2 Preliminaries

Notations. For a natural number n , $[n]$ denotes the set $\{1, \dots, n\}$, U_n denotes the uniform distribution on binary strings of length n . For a finite set Σ , U_Σ denotes the uniform distribution on Σ . For a distribution X , $\text{supp}(X)$ denotes the support of X , and $x \leftarrow X$ means x is a random sample drawn from distribution X . We write $\text{Avg}_{a \leq i \leq b}$ as a shorthand for the average over all $i \in \{a, \dots, b\}$. $\text{Conv}(\cdot)$ denotes the convex hull.

For more background on entropy and proofs of the lemmas stated below, see [CT].

Definition 2.1 (Entropy). For a distribution X , the (*Shannon*) *entropy* of X is defined to be

$$H(X) = \mathbb{E}_{x \leftarrow X} \left[\log \frac{1}{\Pr[X = x]} \right].$$

The *min-entropy* of X is defined to be

$$H_\infty(X) = \min_{x \in \text{supp}(X)} \left(\log \frac{1}{\Pr[X = x]} \right).$$

The notion of *KL divergence* from distribution A to distribution B is closely related to Shannon entropy; intuitively it measures how dense A is within B , on average (with 0 divergence representing maximum density, i.e. $A = B$, and large divergence meaning that A is concentrated in a small portion of B).

Definition 2.2 (KL divergence). For distributions A and B , the *KL divergence from A to B* is defined to be

$$\text{KL}(A \parallel B) = \mathbb{E}_{a \leftarrow A} \left[\log \frac{\Pr[A = a]}{\Pr[B = a]} \right],$$

or conventionally $+\infty$ if $\text{supp}(A) \not\subseteq \text{supp}(B)$.

While the KL divergence is *not* a metric (it is not symmetric and does not satisfy the triangle inequality), it does satisfy nonnegativity, and equals zero only if the distributions are identical:

Proposition 2.3 (Nonnegativity of KL divergence). *For all distributions A and B , $\text{KL}(A \parallel B) \geq 0$. In particular, $\text{KL}(A \parallel B) = 0$ if and only if $A = B$.*

Definition 2.4 (Conditional KL divergence). For joint distributions (X, A) and (Y, B) , the *conditional KL divergence from $A|X$ to $B|Y$* is defined to be

$$\text{KL}((A|X) \parallel (B|Y)) = \mathbb{E}_{(x,a) \leftarrow (X,A)} \left[\log \frac{\Pr[A = a|X = x]}{\Pr[B = a|Y = x]} \right].$$

Thus, conditional KL divergence captures the expected KL divergence from $A|_{X=x}$ to $B|_{Y=x}$, over $x \leftarrow X$. Like Shannon entropy, KL divergence has a chain rule:

Proposition 2.5 (Chain rule for KL divergence). $\text{KL}(X, A \parallel Y, B) = \text{KL}(X \parallel Y) + \text{KL}((A|X) \parallel (B|Y))$.

Definition 2.6 (KL projection). Let X be a distribution on Σ , and \mathcal{V} be a non-empty closed convex set of distributions on Σ . $Y^* \in \mathcal{V}$ is called a *KL projection of X on \mathcal{V}* if

$$Y^* = \arg \min_{Y \in \mathcal{V}} \text{KL}(Y \parallel X).$$

A nice property of KL projection is the following geometric structure (see [CT], Chap 11, Section 6):

Theorem 2.7 (Pythagorean Theorem). *Let \mathcal{V} be a non-empty closed convex set of distributions on Σ . Let Y^* be a KL projection of X on \mathcal{V} . Then for all $Y \in \mathcal{V}$,*

$$\text{KL}(Y \parallel Y^*) + \text{KL}(Y^* \parallel X) \leq \text{KL}(Y \parallel X).$$

In particular,

$$\text{KL}(Y \parallel Y^*) \leq \text{KL}(Y \parallel X).$$

Assuming $\text{KL}(Y^* \parallel X)$ is finite, then the Pythagorean Theorem implies that the KL projection is unique:

Lemma 2.8. *The KL projection is unique.*

Proof. Suppose Y^* and Y are both KL projections of X on \mathcal{V} . Then by the Pythagorean Theorem (Theorem 2.7) $\text{KL}(Y \parallel Y^*) = 0$, which implies $Y = Y^*$ by Proposition 2.3. \square

Finding the exact KL projection is often computationally infeasible, so we consider *approximate KL projection*:

Definition 2.9 (Approximate KL projection). We say Y^* is a σ -approximate KL projection of X on \mathcal{V} , if $Y^* \in \mathcal{V}$ and for all $Y \in \mathcal{V}$,

$$\text{KL}(Y \parallel Y^*) \leq \text{KL}(Y \parallel X) + \sigma.$$

3 A Uniform Min-Max Theorem

Consider a zero-sum game between two players, where the space of pure strategies for Player 1 is \mathcal{V} , the space of pure strategies for Player 2 is \mathcal{W} , and \mathcal{V} is an arbitrary subset of distributions over $[N]$. In this section we present a Uniform Min-Max Theorem that efficiently finds an approximately optimal strategy $W^* \in \text{Conv}(\mathcal{W})$ for Player 2, given an oracle which, when fed any of Player 1's mixed strategies $V \in \text{Conv}(\mathcal{V})$, returns a strategy for Player 2 that guarantees good payoff. Our algorithm is inspired by the proof of Uniform Hardcore Theorem of Barak, Hardt, and Kale [BHK]. Like [BHK], our algorithm uses “relative entropy (KL) projections” together with multiplicative weight updates (a technique originally due to Herbster and Warmuth [HW]).

We first state the theorem and mention how it implies standard Min-Max Theorem.

Theorem 3.1 (A Uniform Min-Max Theorem). *Consider a two-player zero-sum game where the sets of pure strategies for Player 1 and Player 2 are $\mathcal{V} \subseteq \{\text{distributions over } [N]\}$ and \mathcal{W} , and the payoff to Player 2 is defined to be $F(V, W) = \mathbb{E}_V [f(V, W)]$ for some function $f : [N] \times \mathcal{W} \rightarrow [-k, k]$. Then for every $0 < \epsilon \leq 1$ and S , Algorithm 3.1 (Finding Universal Strategy) always outputs a mixed strategy W^* for Player 2 such that*

$$F(V, W^*) \geq \text{Avg}_{1 \leq i \leq S} F(V^{(i)}, W^{(i)}) - O(k\epsilon)$$

for all Player 1 strategies $V \in \mathcal{V}$ where $\text{KL}(V \parallel V_1) \leq S \cdot \epsilon^2$. (This holds regardless of the arbitrary choice of $W^{(i)}$ and $V^{(i+1)}$ in the algorithm.)

In particular, taking $S \geq (\log N - \min_{V \in \mathcal{V}} H(V)) / \epsilon^2$ where we set $V^{(1)} = U_{[N]} \in \text{Conv}(\mathcal{V})$ yields that for all $V \in \mathcal{V}$,

$$F(V, W^*) \geq \text{Avg}_{1 \leq i \leq S} F(V^{(i)}, W^{(i)}) - O(k\epsilon).$$

$S \leftarrow (n - \min_{V \in \mathcal{V}} H(V)) / \epsilon^2$
 Choose an initial strategy $V^{(1)} \in \text{Conv}(\mathcal{V})$ for Player 1
for $i \leftarrow 1$ **to** S **do**
 $W^{(i)} \leftarrow \widehat{W}(V^{(i)})$
 Weight Update:
 Let $V^{(i)'}$ be such that $\Pr[V^{(i)'} = x] \propto e^{-\epsilon \cdot f(x, W^{(i)}) / 2k} \cdot \Pr[V^{(i)} = x]$
 Projection:
 $V^{(i+1)} \leftarrow$ an arbitrary ϵ^2 -approximate KL projection of $V^{(i)'}$ on $\text{Conv}(\mathcal{V})$
end
 Let W^* be the mixed strategy for Player 2 uniform over $W^{(1)}, \dots, W^{(S)}$
return W^*

Algorithm 3.1: Finding Universal Strategy

We now describe how Theorem 3.1 implies the original Min-Max Theorem, which says

$$\max_{W \in \text{Conv}(\mathcal{W})} \min_{V \in \mathcal{V}} F(V, W) = \min_{V \in \text{Conv}(\mathcal{V})} \max_{W \in \mathcal{W}} F(V, W).$$

For each i , take $W^{(i)}$ to be Player 2's best response to Player 1's mixed strategy $V^{(i)}$, i.e. $F(V^{(i)}, W^{(i)}) = \max_{W \in \mathcal{W}} F(V^{(i)}, W)$. Theorem 3.1 says for every $\lambda = O(k\epsilon) > 0$, by setting an appropriate $V^{(1)}$ and sufficiently large S , there exists $W^* \in \text{Conv}(\mathcal{W})$ with

$$\begin{aligned} \min_{V \in \mathcal{V}} F(V, W^*) &\geq \text{Avg}_{1 \leq i \leq S} F(V^{(i)}, W^{(i)}) - \lambda \\ &= \text{Avg}_{1 \leq i \leq S} \max_{W \in \mathcal{W}} F(V^{(i)}, W) - \lambda \\ &\geq \min_{V \in \text{Conv}(\mathcal{V})} \max_{W \in \mathcal{W}} F(V, W) - \lambda, \end{aligned}$$

where the last inequality holds because $\max_{W \in \mathcal{W}} F(V^{(i)}, W) \geq \min_{V \in \text{Conv}(\mathcal{V})} \max_{W \in \mathcal{W}} F(V, W)$ for every i . Thus, for every $\lambda > 0$,

$$\max_{W \in \text{Conv}(\mathcal{W})} \min_{V \in \mathcal{V}} F(V, W) \geq \min_{V \in \text{Conv}(\mathcal{V})} \max_{W \in \mathcal{W}} F(V, W) - \lambda$$

Taking $\lambda \rightarrow 0$ gives the Min-Max Theorem.

Proof of Theorem 3.1. Consider any $V \in \mathcal{V}$ such that $\text{KL}(V \parallel V_1) \leq S \cdot \epsilon^2$. We show in Lemma A.1 that

$$\text{KL}(V \parallel V^{(i)}) - \text{KL}(V \parallel V^{(i)'}) \geq (\log e) \epsilon \left(\frac{F(V^{(i)}, W^{(i)}) - F(V, W^{(i)})}{2k} - \epsilon \right).$$

Since $V^{(i+1)}$ is an ϵ^2 -approximate KL projection of $V^{(i)'}$ on $\text{Conv}(\mathcal{V})$, by definition we have $\text{KL}(V \parallel V^{(i+1)}) \leq \text{KL}(V \parallel V^{(i)'}) + \epsilon^2$. Therefore

$$\text{KL}(V \parallel V^{(i)}) - \text{KL}(V \parallel V^{(i+1)}) \geq (\log e)\epsilon \left(\frac{F(V^{(i)}, W^{(i)}) - F(V, W^{(i)})}{2k} - \epsilon \right) - \epsilon^2.$$

Summing over $i = 1, \dots, S$ and telescoping, we obtain

$$\begin{aligned} \text{KL}(V \parallel V^{(1)}) - \text{KL}(V \parallel V^{(S+1)}) &\geq (\log e)\epsilon \sum_{i=1}^S \left(\frac{F(V^{(i)}, W^{(i)}) - F(V, W^{(i)})}{2k} - \epsilon \right) - S\epsilon^2 \\ &= (\log e)S\epsilon \left(\frac{\text{Avg}_{1 \leq i \leq S} F(V^{(i)}, W^{(i)}) - F(V, W^*)}{2k} - \epsilon \right) - S\epsilon^2. \end{aligned}$$

Since $\text{KL}(V \parallel V^{(S+1)}) \geq 0$ and $\text{KL}(V \parallel V_1) \leq S \cdot \epsilon^2$, rearranging yields

$$\frac{\text{Avg}_{1 \leq i \leq S} F(V^{(i)}, W^{(i)}) - F(V, W^*)}{2k} \leq \frac{\text{KL}(V \parallel V^{(1)}) + S\epsilon^2}{(\log e)S\epsilon} + \epsilon = O(\epsilon).$$

□

Next we describe an average case variant where the set \mathcal{V} of strategies for Player 1 is a set of distributions of the form (X, C) where C may vary, but the marginal distribution of X is fixed. This is convenient for a number of applications (e.g. Section 5 and 8) that involve distinguishers on such joint distributions (X, C) .

Theorem 3.2 (Uniform Min-Max Theorem – Average Case). *Let \mathcal{V} be a subset of distributions over $[N] \times [q]$ of the form (X, C) where C may vary, but the marginal distribution of X is fixed. That is, for every $(X, C), (X', C') \in \mathcal{V}$ and every $x \in [N]$ we have $\sum_c \Pr[(X, C) = (x, c)] = \sum_c \Pr[(X', C') = (x, c)]$.*

Consider a two-player zero-sum game where the sets of pure strategies for Player 1 and Player 2 are \mathcal{V} and \mathcal{W} , and the payoff to Player 2 is defined to be $F((X, C), W) = \mathbb{E}_{X, C} [f((X, C), W)]$ for some function $f : [N] \times [q] \times \mathcal{W} \rightarrow [-k, k]$. Then for every $0 < \epsilon \leq 1$ and S , Algorithm 3.2 (Finding Universal Strategy – Average Case) always outputs a mixed strategy W^ for Player 2 such that*

$$F((X, C), W^*) \geq \text{Avg}_{1 \leq i \leq S} F((X, C^{(i)}), W^{(i)}) - O(k\epsilon)$$

for all Player 1 strategies $(X, C) \in \mathcal{V}$ where $\text{KL}(X, C \parallel X, C^{(1)}) \leq S \cdot \epsilon^2$. (This holds regardless of the arbitrary choice of $W^{(i)}$ and $C^{(i+1)}$ in the algorithm.)

In particular, taking $S \geq (\log q - \min_{(X, C) \in \mathcal{V}} H(C|X)) / \epsilon^2$ where we set $(X, C^{(1)}) = (X, U_{[q]}) \in \text{Conv}(\mathcal{V})$ ($U_{[q]}$ being independent of X) yields that for all $(X, C) \in \mathcal{V}$,

$$F((X, C), W^*) \geq \text{Avg}_{1 \leq i \leq S} F((X, C^{(i)}), W^{(i)}) - O(k\epsilon).$$


```

Choose an initial strategy  $(X, C^{(1)}) \in \text{Conv}(\mathcal{V})$  for Player 1
for  $i \leftarrow 1$  to  $S$  do
    Obtain an arbitrary strategy  $W^{(i)} \in \mathcal{W}$  for Player 2, in response to  $(X, C^{(i)})$ 
    Weight Update:
    Let  $C^{(i)'}$  be such that  $\forall x, a, \Pr[C^{(i)'} = a | X = x] \propto e^{-\epsilon \cdot f(x, a, W^{(i)})/2k} \cdot \Pr[C^{(i)} = a | X = x]$ 
    Projection:
     $(X, C^{(i+1)}) \leftarrow$  an arbitrary  $\epsilon^2$ -approximate KL projection of  $(X, C^{(i)'})$  on  $\text{Conv}(\mathcal{V})$ 
end
Let  $W^*$  be the mixed strategy for Player 2 uniform over  $W^{(1)}, \dots, W^{(S)}$ 
return  $W^*$ 

```

Algorithm 3.2: Finding Universal Strategy – Average Case

Proof. Note that Algorithm 3.2 is the same as Algorithm 3.1, except for the difference that here we update $C^{(i)}$ instead of $V^{(i)}$. We show that the combined effect of the update and KL projection steps is identical in the two algorithms. Note that we can write $V^{(i)'}$ as $(X^{(i)'}, g_i(X^{(i)'}))$ for the randomized function g_i where $\Pr[g_i(x) = a] \propto e^{\epsilon \cdot f(x, a, W^{(i)})/2k} \cdot \Pr[C^{(i)} = a | X = x]$ for every x and a . For the same function g_i , we have $(X, g_i(X)) = (X, C^{(i)'})$. Thus, we can apply the following lemma. \square

Lemma 3.3. *Let X' be a distribution on $[N]$ with $\text{supp}(X') \supseteq \text{supp}(X)$, and let $g : [N] \rightarrow [q]$ be a randomized function. Then the KL projection of $(X', g(X'))$ on $\text{Conv}(\mathcal{V})$ equals the KL projection of $(X, g(X))$ on $\text{Conv}(\mathcal{V})$.*

Proof. Consider any $(X, C) \in \text{Conv}(\mathcal{V})$. We have

$$\begin{aligned}
& \text{KL}(X, C \parallel X', g(X')) \\
&= \text{KL}(X \parallel X') + \text{KL}((C|X) \parallel (g(X')|X')) && \text{(by the chain rule for KL divergence)} \\
&= \text{KL}(X \parallel X') + \text{KL}((C|X) \parallel (g(X)|X)) && \text{(by definition of conditional KL divergence)} \\
&= \text{KL}(X \parallel X') + \text{KL}(X, C \parallel X, g(X)). && \text{(by the chain rule for KL divergence)}
\end{aligned}$$

Thus the KL projections are the same. \square

4 Application: Uniform Hardcore Theorem

A fundamental result in complexity theory is Impagliazzo's Hardcore Theorem [Imp], which, in the strengthened version due to Klivans and Servedio [KS] and Holenstein [Hol1], says that every function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ that is δ -hard for poly-sized boolean circuits (that is, every poly-sized circuit fails to compute f on at least δ fraction of inputs) must be *extremely* hard on a subset of inputs of density at least 2δ (the *hardcore set*) (and may be easy elsewhere). In this section, we provide a simplified proof of a hardcore theorem with optimal parameters, where hardness is defined with respect to *uniform* algorithms rather than boolean circuits. Following [Imp], we will deal with hardcore distributions instead of hardcore sets, which are equivalent up to a negligible additive difference in density, where density of a distribution is defined as follows:

Definition 4.1 (Density of distribution). Let X and Y be distributions over some finite set Σ . We say X is δ -dense in Y if $\Pr[Y = x] \geq \delta \cdot \Pr[X = x]$ for all $x \in \Sigma$. We say X is δ -dense if it is δ -dense in U_Σ (equivalently, having min-entropy at least $\log|\Sigma| - \log(1/\delta)$). We denote by $\mathcal{C}_{m,\delta}$ the set of all δ -dense distributions on $\{0, 1\}^m$.

The asymptotically optimal nonuniform Hardcore Theorem is due to [KS], using techniques from boosting and an idea of iteratively increasing hardcore size due to Wigderson, and can be stated as follows:

Theorem 4.2 (Hardcore Theorem [KS]). *Let $(X, B)^1$ be a joint distribution over $\{0, 1\}^n \times \{0, 1\}$ and $\epsilon > 0$. Let B be (t, δ) -hard given X , i.e. for every size t circuit P it holds that $\Pr[P(X) = B] \leq 1 - \delta$. Then there is a joint distribution (\hat{X}, \hat{B}) that is 2δ -dense in (X, B) , such that for every size $t' = t/O(\log(1/\delta)/\epsilon^2)$ circuit A it holds that $\Pr[A(\hat{X}) = \hat{B}] \leq (1 + \epsilon)/2$.*

Theorem 4.2 is asymptotically optimal as it achieves optimal hardcore density 2δ , as well as optimal complexity blow-up $O(\log(1/\delta)/\epsilon^2)$, where the lower bound of $\Omega(\log(1/\delta)/\epsilon^2)$ is due to Lu, Tsai, and Wu [LTW]².

The original paper of Impagliazzo [Imp] contains both a non-trivial constructive proof, as well as a much simpler, yet non-constructive proof due to Nisan that uses the Min-Max Theorem. Nisan's proof has an appealing simplicity: Assume for contradiction that there is no hardcore distribution of high density. Then, by the Min-Max Theorem there is a *universal* predictor A^* such that for every (\hat{X}, \hat{B}) that is dense in (X, B) it holds that $\Pr[A^*(\hat{X}) = \hat{B}] > (1 + \epsilon)/2$. A^* is a distribution over circuits of size t , and its prediction probability is taken over this distribution as well as (\hat{X}, \hat{B}) . By subsampling we can assume that A^* is uniform over a multiset of $S = O((1/\epsilon^2) \log(1/\epsilon\delta))$ circuits of size t , while changing the advantage ϵ by at most a constant fraction. Given the universal predictor A^* , one can build a good predictor for B , contradicting the hardness of B given X , as formalized in Lemma 4.3:

Lemma 4.3 (From universal circuit to predictor [Imp]). *Let (X, B) be a joint distribution on $\{0, 1\}^n \times \{0, 1\}$. Let A^* be the uniform distribution over a multiset of S circuits of size t . Suppose for every joint distribution (\hat{X}, \hat{B}) that is δ -dense in (X, B) it holds that $\Pr[A^*(\hat{X}) = \hat{B}] > (1 + \epsilon)/2$. Then there is a circuit P of size $O(S \cdot t)$ such that $\Pr[P(X) = B] > 1 - \delta$.*

Specifically, we can let $P(x) = \text{majority}\{A(x) : A \in A^\}$. Equivalently, $P(x)$ outputs 1 with probability*

$$\frac{1}{2} \left(1 + \text{sign} \left(\Pr[A^*(x) = 1] - \frac{1}{2} \right) \right).$$

Unfortunately, both proofs in [Imp] yield a suboptimal hardcore density of δ . Following Nisan's proof using Min-Max Theorem, Holenstein [Hol1] proves the hardcore theorem with optimal hardcore density of 2δ (Theorem 4.2), by strengthening the above lemma to Lemma 4.4 below (using a trick from Levin's proof of the XOR Lemma).

Lemma 4.4 (From universal circuit to *optimal* predictor [Hol1]). *Let (X, B) be a joint distribution on $\{0, 1\}^n \times \{0, 1\}$. Let A^* be the uniform distribution over a multiset of S circuits of size t . Suppose*

¹The version we state is a slight generalization of the version in [KS], which only allows B to be a deterministic boolean function of X . However, the more general version follows readily from almost the same proof.

²[LTW] showed a black-box lower bound on the number of t' -sized circuits that a black-box reduction needs to obtain to construct some P with $\Pr[P(X) = B] > 1 - \delta$.

for every joint distribution (\hat{X}, \hat{B}) that is 2δ -dense in (X, B) it holds that $\Pr[A^*(\hat{X}) = \hat{B}] > (1 + \epsilon)/2$. Then there is a circuit P of size $O(S \cdot t)$ such that $\Pr[P(X) = B] > 1 - (1 - \epsilon)\delta$.

Specifically, we can let $P(x)$ output 1 with probability $p(x)$ truncated at 0 and 1 (i.e. $\min\{\max\{p(x), 0\}, 1\}$), for

$$p(x) = \frac{1}{2} \left(1 + \frac{\Pr[A^*(x) = 1] - \frac{1}{2}}{\phi} \right)$$

where ϕ is the least number s.t. $\Pr_{X,B}[\Pr_{A^*}[A^*(X) = B] \leq 1/2 + \phi] \geq 2\delta$. (WLOG ϕ is a multiple of $1/S$.)

One drawback of proofs based on the standard Min-Max Theorem is the suboptimal complexity blow-up (due to suboptimal settings of S from the probabilistic construction of the multiset defining A^*). By replacing the use of Min-Max Theorem with the Uniform Min-Max Theorem, we immediately achieve optimal complexity blow-up (by replacing the probabilistic construction of the multiset with a smarter online learning/boosting algorithm).

Another drawback of proofs based on the standard Min-Max Theorem is that they are non-constructive. Indeed, a constructive proof such as the one by Impagliazzo [Imp] can be interpreted as a hardcore theorem for the *uniform* setting of hardness, where the hardness is with respect to efficient algorithms rather than small circuits. (See Theorem 4.5 below for the exact formulation). This *Uniform Hardcore Theorem* is needed for several important applications ([KS, Hol1, Hol2, HHR, HRV]). Building on the constructive proof in [Imp], Holenstein [Hol1] also shows a *uniform* hardcore theorem with optimal hardcore density, but is rather involved and fails to achieve the optimal complexity blow-up $O(\log(1/\delta)/\epsilon^2)$. Subsequently, Barak, Hardt, and Kale ([BHK]) gave an alternative proof of uniform hardcore theorem achieving optimal complexity blow-up of $O(\log(1/\delta)/\epsilon^2)$ (but without optimal hardcore density), based on ideas of multiplicative weights and Bregman projection.

As an application of the Uniform Min-Max Theorem (which itself is inspired by [BHK]), we offer a new proof of the Uniform Hardcore Theorem. Essentially, our proof simply replaces the use of Min-Max Theorem in Holenstein's proof (of the non-uniform hardcore theorem, Theorem 4.2) with the Uniform Min-Max Theorem. Consequently it has the advantages of (i) optimal hardcore density 2δ ; (ii) optimal complexity blow-up $O(\log(1/\delta)/\epsilon^2)$; (iii) being simpler (e.g. compared to Holenstein's uniform proof [Hol1]), and more modular (e.g. compared to [BHK], as it avoids adapting the analysis of [HW] to the specific setting of the hardcore theorem).

Notation. For a distribution Z , let O_Z denote the oracle that gives a random sample from Z when queried.

Theorem 4.5 (Uniform Hardcore Theorem). *Let n be a security parameter, $m = m(n) = \text{poly}(n)$, $\delta = \delta(n)$, $\epsilon' = \epsilon'(n)$, $q = q(n)$ all computable in $\text{poly}(n)$ time, and $(X, B) = g(U_m)$ be a joint distribution where $g : \{0, 1\}^m \rightarrow \{0, 1\}^n \times \{0, 1\}$ is computable in $\text{poly}(n)$ time. Suppose that (X, B) has no hardcore distribution of density at least 2δ , i.e. there is a t -time oracle algorithm A such that for infinitely many n and every $C \in \mathcal{C}_{m, 2\delta}$,*

$$\Pr_{(x,b) \leftarrow g(C)} [A^{O_C}(x) = b] > \frac{1}{2} + \epsilon'.$$

Then there is a $\text{poly}(t, n, 1/\delta, 1/\epsilon')$ -time randomized algorithm P such that for infinitely many n ,

$$\Pr[P(X) = B] > 1 - \delta.$$

Moreover, P is constructed by making $O(\log(1/\delta)/\epsilon'^2)$ calls to A .

For the proof of Uniform Hardcore Theorem (as well as several results in Section 6), we will need the notion of *measures*, which are simply $[0, 1]$ bounded, unnormalized mass functions:

Definition 4.6 (Measure). A function $M : \mathcal{X} \rightarrow [0, 1]$ is said to be a *measure* for X if $\Pr[X = x] = M(x)/\sum_{x \in \mathcal{X}} M(x)$, and we denote this X by Φ_M . We say M is δ -dense if its *density* $\mu(M) = \sum_{x \in \mathcal{X}} M(x)/|\mathcal{X}|$ is at least δ . We denote by $\mathcal{M}_{m,\delta}$ the set of all δ -dense measures defined on $\{0, 1\}^m$. One can verify that if $M \in \mathcal{M}_{m,\delta}$ then $\Phi_M \in \mathcal{C}_{m,\delta}$ (but not conversely).

Proof of Theorem 4.5. We will apply Theorem 3.1 (Uniform Min-Max Theorem), with

- $\mathcal{V} = \mathcal{C}_{m,2\delta}$;
- $\mathcal{W} = \{(\text{deterministic}) \text{ circuits of size } tm + \text{poly}(t)\}$;
- $f(z, W) = I(W(x) = b)$, where $(x, b) = g(z)$ and $I(\cdot)$ is the indicator function.

This corresponds to the two-player zero-sum game where Player 1 chooses some distribution $C \in \mathcal{C}_{m,2\delta}$, and Player 2 chooses a $tm + \text{poly}(t)$ sized circuit W , with expected payoff $F(C, W) = \mathbb{E}[f(C, W)] = \Pr_{(x,b) \leftarrow g(C)} [W(x) = b]$ for Player 2. We will use A to show that Algorithm 3.1 (Finding Universal Strategy) with KL projection on the set $\mathcal{V} = \mathcal{C}_{m,2\delta}$ can be implemented efficiently, such that for infinitely many n , in each iteration we obtain some W with good prediction probability. This gives us an efficient universal predictor A^* of B given X , by the Uniform Min-Max Theorem. From the universal predictor, we then show how to obtain a $(1 - \delta)$ -predictor of B using Lemma 4.4.

In Algorithm 3.1, we start with an initial distribution $V^{(1)}$ that is uniform on $\{0, 1\}^m$. Let $\epsilon = \epsilon'/c$ for a sufficiently large constant c , and $\gamma = \epsilon/2S$. The number of iterations is

$$S = \left(m - \min_{C \in \mathcal{C}_{m,2\delta}} H(C) \right) / \epsilon^2 = (m - (m - \log(1/2\delta))) / \epsilon^2 = (\log(1/\delta) - 1) / \epsilon^2.$$

In each iteration we represent the distribution $V^{(i)}$ (the current C) by a circuit $M^{(i)}$ computing a measure for $V^{(i)}$. So we can take $M^{(1)}(x) = 1$ for all x . We will need the following claim to implement an iteration.

Claim 4.7. There is a randomized algorithm that, given oracle access to a measure $M \in \mathcal{M}_{m,2\delta}$, w.p. at least $1 - \gamma$ outputs a $tm + \text{poly}(t)$ sized (deterministic) boolean circuit W such that $\Pr_{(x,b) \leftarrow g(\Phi_M)} [W(x) = b] > 1/2 + \epsilon' - 4\epsilon$. The algorithm runs in $t + \text{poly}(n, s, t, 1/\delta, 1/\epsilon', \log(1/\gamma))$ time where s is a bound on the bit length of $M(x)$.

Proof of Claim 4.7. Given oracle access to M , we can generate t random samples of Φ_M in time $t' = t \cdot O((1/\delta) \log(t/\epsilon)) \cdot (s + m) + \text{poly}(n)$ and w.p. at least $1 - \epsilon$, using rejection sampling (see Lemma A.2). Thus we can eliminate all A 's oracle queries to O_{Φ_M} and obtain some t' time randomized algorithm A' such that $\Pr_{(x,b) \leftarrow g(\Phi_M)} [A'(x) = b] > 1/2 + \epsilon' - \epsilon$.

Write $A'(x) = A'(x; r)$ where r is the coin tosses of A' (which consists of coin tosses for A and at most t' random bits for the rejection sampling). For each r we compute an estimate $E(r)$ of $\Pr_{(x,b) \leftarrow g(\Phi_M)} [A'(x; r) = b]$ within $\pm\epsilon$ error with probability at least $\gamma/2q$, for $q = O((1/\epsilon) \log(1/\gamma))$. By a Chernoff bound, this can be done by testing $A'(\cdot; r)$ on $q' = O((1/\epsilon^2) \log(q/\gamma))$ random samples of $(x, b) \leftarrow g(\Phi_M)$ (which we generate with probability at least $1 - \Theta(\gamma/q)$, again using Lemma

A.2). We repeat this for q randomly chosen r , and if $E(r) > 1/2 + \epsilon' - 3\epsilon$ output a circuit W computing $A'(\cdot; r)$.

By union bound with probability at least $1 - \gamma/2$, all q estimates $E(r)$ are within ϵ error. By the Markov inequality, w.p. $1 - (1 - \Omega(\epsilon))^q \geq 1 - \gamma/2$ at least one of the r 's satisfies $\Pr_{(x,b) \leftarrow g(\Phi)}[A'(x; r) = b] > 1/2 + \epsilon' - 2\epsilon$ so $E(r) \geq 1/2 + \epsilon' - 3\epsilon$. Moreover we have $\Pr_{(x,b) \leftarrow g(C)}[A(x, r) = b] > 1/2 + \epsilon' - 4\epsilon$ whenever $E(r) > 1/2 + \epsilon' - 3\epsilon$. We conclude that w.p. at least $1 - \gamma$ we output the desired circuit, all in time $q'q \cdot (t' + O((1/\delta) \log(qq'/\gamma)) \cdot (s + m)) + \text{poly}(n) = \text{poly}(n, s, t, 1/\delta, 1/\epsilon', \log(1/\gamma))$. Finally, the circuit W is of size $tm + \text{poly}(t)$ as it simply runs A using the t fixed samples of Φ_M (which can be stored as tm nonuniform bits). \square

We now implement the i th iteration as follows. For technical convenience we assume that $e^{-\epsilon}$ has bit-length $\log(1/\epsilon)$ (if not, we replace ϵ by some $\tilde{\epsilon} = O(\epsilon)$ such that $e^{-\tilde{\epsilon}}$ has bit-length $\log(1/\epsilon)$).

1. **Obtaining Player 2's Response $W^{(i)}$:** Suppose that we have constructed a t_i sized circuit $M^{(i)}$ computing a measure for $V^{(i)}$, and outputs of $M^{(i)}$ have bit-length at most $O(i \cdot \log(1/\epsilon))$. Using Claim 4.7, we can obtain a (deterministic) circuit $W^{(i)}$ such that

$$\Pr_{(x,b) \leftarrow g(V^{(i)})} [W^{(i)}(x) = b] > \frac{1}{2} + \epsilon' - 4\epsilon,$$

in time $\text{poly}(t_i, n, t, 1/\delta, 1/\epsilon, \log(1/\gamma))$ and w.p. at least $1 - \gamma$. Note, however, that the circuit size of $W^{(i)}$ is $tm + \text{poly}(t)$, independent of t_i .

2. **Weight Update:** We represent the resulting distribution $V^{(i)'}$ by the circuit $M^{(i)'}(z) = \exp(-\epsilon \cdot I(W^{(i)}(x) = b)) \cdot M^{(i)}(z)$, where $(x, b) = g(z)$, which computes a measure for $V^{(i)'}$. Since $I(W^{(i)}(x) = b) \in \{0, 1\}$, $\exp(-\epsilon \cdot I(W^{(i)}(x) = b))$ has bit-length $\log(1/\epsilon)$. $M^{(i)}(z)$ has bit-length $O(i \cdot \log(1/\epsilon))$, thus multiplication takes time $\text{poly}(i \cdot \log(1/\epsilon))$. Thus $M^{(i)'}$ has circuit size $t'_i = t_i + tm + \text{poly}(t) + i \cdot \text{polylog}(1/\epsilon)$, bit-length at most $O(i \cdot \log(1/\epsilon) + \log(1/\epsilon))$, and can be constructed in similar time.
3. **KL Projection:** It is shown in Lemma A.3 (approximating KL projection on high min-entropy distributions, which is based on Lemma 2.3 of [BHK]) that given $M^{(i)'}$, w.p. $1 - \gamma$ one can generate a $t_{i+1} = t'_i + \text{polylog}(1/\epsilon)$ sized circuit $M^{(i+1)}$ computing a measure for a distribution $V^{(i+1)}$ that is an ϵ^2 -approximate KL projection of $V^{(i)'} = \Phi_{M^{(i)'}}$ on $\mathcal{C}_{m, 2\delta}$. Furthermore, outputs of $M^{(i+1)}$ have bit-length at most $O((i+1) \log(1/\epsilon))$. This can be done in time $\text{poly}(n, 1/\epsilon, \log(1/\delta), \log(1/\gamma)) \cdot t'_i$.

By union bound w.p. at least $1 - 2\gamma S = 1 - \epsilon$ all S iterations complete successfully. Since $t_1 = O(1)$ and $t_{i+1} = t_i + tm + \text{poly}(t) + i \cdot \text{polylog}(1/\epsilon)$, we have $t_i = \text{poly}(n, t, 1/\epsilon, \log(1/\delta))$ for all $i \in [S]$. Let A^* be the uniform distribution over $W^{(1)}, \dots, W^{(S)}$, thus A^* can be computed in total time $\text{poly}(n, t, 1/\delta, 1/\epsilon)$. By the Uniform Min-Max Theorem (Theorem 3.1), for all Player 1 strategies $C \in \mathcal{C}_{m, 2\delta}$,

$$\Pr_{(x,b) \leftarrow g(C)} [A^*(x) = b] > (1 - \epsilon) \left(\frac{1}{2} + \epsilon' - 4\epsilon \right) - O(\epsilon) \geq \frac{1 + \epsilon'}{2}.$$

Equivalently, for every joint distribution (\hat{X}, \hat{B}) that is 2δ -dense in $(X, B) = g(U_m)$ we have

$$\Pr[A^*(\hat{X}) = \hat{B}] > \frac{1 + \epsilon'}{2}$$

(since (\hat{X}, \hat{B}) equals $g(C)$ for some $C \in \mathcal{C}_{m, 2\delta}$).

From Universal Weak Predictor to $(1 - \delta)$ -Predictor. Now that we have a universal weak predictor A^* as the uniform distribution over $S = O(\log(1/\delta)/\epsilon'^2)$ circuits, applying Lemma 4.3 already proves a version of the Uniform Hardcore Theorem with suboptimal hardcore density.

To achieve optimal hardcore density, we apply Lemma 4.4 by guessing the value of $\phi \in [0, 1/2]$, which is a multiple of $1/S$. More concretely, for each $\lambda = 1/S, 2/S, \dots, 1/2$, we compute some estimate E_λ of $\Pr[P_\lambda(X) = B]$, where P_λ denotes the predictor in Lemma 4.4 with ϕ set to λ . Our final (uniform) predictor P will run P_λ for the λ where the estimate E_λ is the highest.

We compute E_λ by taking $O((1/\epsilon'^2\delta^2) \log(1/\epsilon'\delta))$ samples of (X, B) and coins of P_λ , so that by a Chernoff bound, for each λ w.p. at least $1 - \epsilon'\delta/4$ we have $|E_\lambda - \Pr[P_\lambda(X) = B]| \leq \epsilon'\delta/4$. The probability that either E_ϕ or the highest estimate is off by more than $\pm\epsilon'\delta/4$ is at most $\epsilon'\delta/2$. So it follows from Lemma 4.4 that

$$\Pr[P(X) = B] \geq \Pr[P_\phi(X) = B] - \epsilon'\delta/2 - \epsilon'\delta/2 > 1 - (1 - \epsilon')\delta - \epsilon'\delta = 1 - \delta$$

completing the proof. □

5 Application: Construction of Pseudorandom Generator Construction from One-Way Functions

Recently, we [VZ] obtained a simplified and more efficient construction of pseudorandom generator from arbitrary one-way functions, building on the work of Haitner, Reingold, and Vadhan [HRV]. Key to the simplification is a characterization of conditional pseudoentropy, defined as follows (for the nonuniform setting):

Definition 5.1 (Nonuniform (conditional) pseudoentropy, informal). Let (X, B) be jointly distributed random variables. We say that B has *nonuniform (conditional) pseudoentropy at least k given X* if there exists a random variable C , jointly distributed with X such that

1. (X, B) is indistinguishable from (X, C) by nonuniform polynomial-time algorithms (i.e. polynomial-sized boolean circuits).
2. $H(C|X) \geq k$, where $H(\cdot|\cdot)$ denotes conditional Shannon entropy.³

In the nonuniform setting, informally, we prove that B has nonuniform pseudoentropy at least $H(B|X) + \delta$ given X if and only if there is no polynomial-time nonuniform algorithm S such that $\text{KL}(X, B \parallel X, S(X)) \leq \delta$ (intuitively, given X it is hard to sample a distribution that is close to B). Our proof of the “if” direction (which is the one used in our PRG analysis) proceeds as follows. Assume for contradiction that B does not have nonuniform pseudoentropy given X . Then by the Min-Max Theorem there is a universal distinguisher D that is a convex combination of polynomial-time nonuniform algorithms and distinguishes (X, B) from all (X, C) where $H(C|X) \geq H(B|X) + \delta$. From such universal distinguisher D we show how to construct an efficient sampler S such that $\text{KL}(X, B \parallel X, S(X)) \leq \delta$, violating the assumption.

To obtain a pseudorandom generator from a one-way function that is secure against *uniform* algorithms, we prove an analogous result in the uniform setting, replacing the use of the Min-Max

³The *conditional (Shannon) entropy* of random variable Y given random variable Z is defined as $H(Y|Z) = \mathbb{E}_{z \sim Z}[H(Y|Z=z)]$.

Theorem by the Uniform Min-Max Theorem. To apply the Uniform Min-Max Theorem we show how to approximately compute the KL projection on the set of distributions of high conditional Shannon-entropy.

6 Application: Regularity Theorems for Distributions Restricted to a Convex Set

Another application of the Uniform Min-Max Theorem is to give a generalization and quantitative improvement to the “regularity theorem” of Trevisan, Tulsiani, and Vadhan [TTV] which (informally) says that any high min-entropy distribution X is indistinguishable from some high min-entropy, *low complexity* distribution Y . The result of [TTV] is itself a quantitative improvement of regularity and “decomposition” theorems in additive combinatorics [GT, TZ]. It is shown in [TTV] that such results can be used to deduce the Dense Model Theorem [TZ, RTTV, Gow], Impagliazzo’s Hardcore Theorem [Imp], and other results, by replacing any unknown distribution X with an “equivalent” distribution Y that can be efficiently analyzed and manipulated, thus translating the problem to a simpler one. It also implies the Weak Regularity Lemma in graph theory [FK], mostly by a translation of notation.

Our result is more general than [TTV] in the sense that we are no longer restricted to distributions of high min-entropy. We show that for any sufficiently nice convex set of distributions \mathcal{V} , every distribution $X \in \mathcal{V}$ is indistinguishable from some distribution $Y \in \mathcal{V}$ where Y has “low complexity”. In the case of min-entropy distributions, we obtain a high min-entropy Y with lower complexity than [TTV]. This also yields an improved and optimal Weak Regularity Lemma for graphs of density $o(1)$ (Section 6.2).

This section is divided into three parts, each proving results for a different notions of “complexity”: Section 6.1 for information-theoretic notion of complexity, Section 6.3 for circuit complexity, and Section 6.4 for time complexity of uniform algorithms.

In addition, using the Uniform Min-Max Theorem – Average Case (Theorem 3.2) we obtain average-case variants, which can be used to deduce “low complexity” versions of a technical lemma of [GW]. We note that the average-case variant for circuits is a strengthening of a recent result of Pietrzak and Jetchev [PJ], with a simpler proof. The low circuit complexity version of the [GW] lemma (with slightly weaker parameters) was initially proved by Pietrzak and Jetchev [PJ], and an interactive extension was proved by Chung, Lui, and Pass [CLP] for applications in the context of distributional zero-knowledge.

6.1 Regularity Theorems for Feature Complexity

Let \mathcal{W} be an arbitrary class of functions $W : \Sigma \rightarrow [0, 1]$ for some finite set Σ . Two distributions X and Y on Σ are ϵ -*indistinguishable by \mathcal{W}* if for every $W \in \mathcal{W}$, $|\mathbb{E}[W(X)] - \mathbb{E}[W(Y)]| < \epsilon$. For starters, we shall consider the setting where the complexity of a distribution Y is purely information-theoretic: We say Y has *feature complexity at most m w.r.t. \mathcal{W}* if its mass function $x \mapsto \Pr[Y = x]$ is a function of $W_1(x), \dots, W_m(x)$, for some $W_1, \dots, W_m \in \mathcal{W}$. Notice that we can assume \mathcal{W} to be closed under negation, i.e. if $W \in \mathcal{W}$ then we can add $1 - W$ to \mathcal{W} without affecting the meaning of complexity.

In order to obtain a low feature complexity approximation within a convex set \mathcal{V} of distributions on Σ , we require \mathcal{V} to be *permutation-invariant*. That is, for all permutations $\pi : \Sigma \rightarrow \Sigma$ we have

$X \in \mathcal{V} \iff \pi(X) \in \mathcal{V}$. Permutation invariance is a natural condition; for example, the set of high entropy distributions should be permutation-invariant for any reasonable notion of entropy. However, for a fixed distribution X_0 , $\{(X, C) : H(C|X) \geq k, X = X_0\}$ is not permutation-invariant in general. We will use the following properties of a permutation-invariant convex set:

Lemma 6.1. *Let \mathcal{V} be a permutation-invariant nonempty convex set of distributions on Σ . Then*

1. \mathcal{V} contains the uniform distribution on Σ .
2. Let X be a distribution on Σ having feature complexity at most m w.r.t. \mathcal{W} . Then the KL projection of X on \mathcal{V} also has feature complexity at most m w.r.t. \mathcal{W} .

Proof. 1. For any $Y \in \mathcal{V}$, the average of $\pi(Y)$ over all permutations π is still in \mathcal{V} (by convexity and permutation-invariance), and is clearly the uniform distribution.

2. Let Y^* denote the KL projection of X on \mathcal{V} . For all $x_1, x_2 \in \Sigma$ where $\Pr[X = x_1] = \Pr[X = x_2]$, we must also have $\Pr[Y^* = x_1] = \Pr[Y^* = x_2]$; otherwise, swapping $\Pr[Y^* = x_1]$ and $\Pr[Y^* = x_2]$ yields some $\widehat{Y}^* \in \mathcal{V}$ (by permutation-invariance) that is also a KL projection of X , violating the uniqueness of KL projection (Lemma 2.8). Therefore $\Pr[Y^* = x]$ is a function of $\Pr[X = x]$, and Y^* has feature complexity at most that of X . □

We show that every distribution $X \in \mathcal{V}$ is indistinguishable to some $Y \in \mathcal{V}$ of low feature complexity, as long as \mathcal{V} is permutation-invariant:

Theorem 6.2 (A regularity theorem for feature complexity). *Let Σ be a finite set, \mathcal{W} be an arbitrary class of functions $W : \Sigma \rightarrow [0, 1]$, \mathcal{V} be a permutation-invariant convex set of distributions on Σ , and $\epsilon > 0$. Then for every distribution $X \in \mathcal{V}$ there exists $Y \in \mathcal{V}$ such that*

1. X and Y are $O(\epsilon)$ -indistinguishable by \mathcal{W} ;
2. Y has feature complexity at most $S = (\log |\Sigma| - H(X))/\epsilon^2$ w.r.t. \mathcal{W} . That is, there exist $W_1, \dots, W_S \in \mathcal{W}$ and a function $\theta : [0, 1]^S \rightarrow [0, 1]$ such that $\forall x$,

$$\Pr[Y = x] = \theta(W_1(x), \dots, W_S(x)).$$

Remark. The main theorem of [TTV] (when considering feature complexity) is equivalent to Theorem 6.2 with \mathcal{V} being fixed to be the set of distributions of min-entropy at least $\log |\Sigma| - \log(1/\delta)$, and has a worse bound on the feature complexity of Y . For a distribution X with $H_\infty(X) = \log |\Sigma| - \log(1/\delta)$, [TTV] obtains a distribution Y with feature complexity at most $1/\epsilon^2 \delta^2$ such that Y is $O(\epsilon)$ -indistinguishable to X and $H_\infty(Y) \geq H_\infty(X)$, whereas Theorem 6.2 obtains such Y with feature complexity at most $\log(1/\delta)/\epsilon^2$.

Theorem 6.2 is interesting even if we do not require the low complexity Y to lie in \mathcal{V} . As mentioned in [TTV] (and pointed out by Elad Verbin), it easily follows from a Chernoff bound and a union bound that the uniform distribution over certain $O(\log |\mathcal{W}|/\epsilon^2)$ elements of Σ (which may not lie in \mathcal{V}) is ϵ -indistinguishable from X by \mathcal{W} . However, for large \mathcal{W} the feature complexity of $O(\log |\mathcal{W}|/\epsilon^2)$ is potentially much higher than $S = (\log |\Sigma| - H(X))/\epsilon^2$. Indeed, we do not use the fact that $Y \in \mathcal{V}$ when deducing the Weak Regularity Lemma of Frieze and Kannan [FK] from Theorem 6.2 (see Theorem 6.4 below); as shown in [TTV], the argument of Frieze and Kannan can be used to obtain a weaker variant of Theorem 6.2 where Y may not lie in \mathcal{V} , and the bound on S is worse.

Proof of Theorem 6.2. Suppose for contradiction that for every low feature complexity $Y \in \mathcal{V}$ there is some $W \in \mathcal{W}$ such that $\mathbb{E}[W(X)] - \mathbb{E}[W(Y)] \geq \epsilon'$ (recall that WLOG \mathcal{W} is closed under negation), where $\epsilon' = c \cdot \epsilon$ for a sufficiently large constant c . Consider the zero-sum game where Player 1 selects some distribution $Y \in \mathcal{V}$, Player 2 selects some $W \in \mathcal{W}$ and receives (expected) payoff $\mathbb{E}[W(X)] - \mathbb{E}[W(Y)]$. Consider Algorithm 3.1 (Finding Universal Strategy) where we set the initial strategy $V^{(1)}$ for Player 1 to be the uniform distribution on Σ (which lies in \mathcal{V} , by Lemma 6.1) and number of iterations to be S . Note that in each iteration the feature complexity of $V^{(i)}$ increases by at most one, due to the weight update using $W^{(i)}$, since KL projection on the permutation-invariant set \mathcal{V} does not increase feature complexity (Lemma 6.1). Hence by assumption, in each iteration there exists $W^{(i)} \in \mathcal{W}$ such that $\mathbb{E}[W^{(i)}(X)] - \mathbb{E}[W^{(i)}(V^{(i)})] \geq \epsilon'$. By the Uniform Min-Max Theorem (Theorem 3.1), W^* (the uniform distribution over $W^{(1)}, \dots, W^{(S)}$) satisfies

$$\mathbb{E}[W^*(X)] - \mathbb{E}[W^*(V)] \geq \epsilon' - O(\epsilon) > 0$$

for all Player 1 strategies $V \in \mathcal{V}$ such that $H(V) \geq H(X)$. Taking $V = X$ yields a contradiction. \square

6.2 Improved Weak Regularity Lemma for Graphs of Density $o(1)$

An information-theoretic application of [TTV] is deducing the Weak Regularity Lemma of Frieze and Kannan [FK]. Our Theorem 6.2, with the improved bound, implies a Weak Regularity Lemma with parameters stronger than [FK] for graphs that are $o(1)$ -dense. The Weak Regularity Lemma says that any graph $G = (V, E)$ is approximated within “cut-distance” σ by some edge-weighted graph G' on the vertices $\{1, \dots, t\}$, where t depends only on σ (i.e. independent of the size of G), and each vertex i corresponds to a block $V_i \subseteq V$ in a partition $\{V_1, \dots, V_t\}$ of V . The edge weight of (i, j) in the approximator G' is defined to be the *edge density* between V_i and V_j :

Definition 6.3 (Edge density). The *density* of a directed graph $G = (V, E)$ equals $|E| / |V|^2$. The *edge density between two sets of vertices* V_1, V_2 of G equals $d_G(V_1, V_2) = |(V_1 \times V_2) \cap E| / |V_1 \times V_2|$.

Theorem 6.4 (A Weak Regularity Lemma). *For every directed graph $G = (V, E)$ of density $\delta = |E| / |V|^2 > 0$ and $\sigma > 0$, there is a partition of V into $t = \exp(O(\delta/\sigma)^2 \log(1/\delta))$ disjoint sets V_1, \dots, V_t , such that for all $A, B \subseteq V$,*

$$\left| |(A \times B) \cap E| - \sum_{i,j} |A \cap V_i| |B \cap V_j| \cdot d_G(V_i, V_j) \right| < \sigma \cdot |V|^2.$$

Note that the only interesting setting of parameters is $\delta > \sigma$, $\delta > 1/|V|^{O(1)}$ (i.e. G has average degree greater than 1), because if $\delta \leq \sigma$ then the trivial partition $V_1 = V$ would work, and if $\sigma < \delta \leq 1/|V|^{O(1)}$ we could take $t = |V|$ and use the trivial partition into single vertices. As pointed out to us by Jacob Fox, the number of partitions $\exp(O(\delta/\sigma)^2 \log(1/\delta))$ in Theorem 6.4 (as a function of δ and σ) is optimal up to a constant factor, which can be shown by adapting a lower bound argument in [CF].

Theorem 6.4 is stronger than Frieze and Kannan [FK] when G has density $\delta = o(1)$. For example, when $|V| = N$ and $\delta = 2\sigma = 1/\text{poly}(\log N)$, Theorem 6.4 produces a partition of size $\text{poly}(\log N)$, whereas [FK] only yields a trivial partition into more than N sets.

Proof of Theorem 6.4. We apply Theorem 6.2 with $\Sigma = V \times V$, $\mathcal{W} = \{\chi_{S \times T}, 1 - \chi_{S \times T} : S, T \subseteq V\}$ (where $\chi_{S \times T}$ denotes the characteristic function of $S \times T$), \mathcal{V} being the set of all δ -dense distributions on Σ , $X = U_E \in \mathcal{V}$ (the uniform distribution on E), and $\epsilon = O(\sigma/\delta)$. By Theorem 6.2 there is some δ -dense distribution Y where:

1. Y has feature complexity at most $m = O((2 \log |V| - H(U_E))/\epsilon^2) = O((\delta/\sigma)^2 \log(1/\delta))$. That is, $\Pr[Y = e] = \phi(\chi_{S_1 \times T_1}(e), \dots, \chi_{S_m \times T_m}(e))$ for a function ϕ and sets $S_1, T_1, \dots, S_m, T_m \subseteq V$.
2. U_E and Y are ϵ -indistinguishable for \mathcal{W} . That is, for every $S, T \subseteq V$,

$$|\mathbb{E}[\chi_{S \times T}(U_E)] - \mathbb{E}[\chi_{S \times T}(Y)]| < \epsilon.$$

The fact that Y has feature complexity at most m yields a partition $\{V_1, \dots, V_t\}$, $t \leq 2^{2m}$, such that $\Pr[Y = e]$ has the same value for all $e \in V_i \times V_j$. (Specifically, the partition is the overlay of $S_1, T_1, \dots, S_m, T_m$, i.e. formed by taking the intersection of, for each i , either S_i or $V - S_i$, and either T_i or $V - T_i$.)

Consider any $A, B \subseteq V$. Taking $S = A$, $T = B$ in Item 2 yields

$$\left| \frac{1}{|E|} |(A \times B) \cap E| - \mathbb{E}[\chi_{A \times B}(Y)] \right| = |\mathbb{E}[\chi_{A \times B}(U_E)] - \mathbb{E}[\chi_{A \times B}(Y)]| < \epsilon.$$

Thus, by triangle inequality it suffices to show that

$$\left| \frac{1}{|E|} \sum_{e \in A \times B} \text{weight}(e) - \mathbb{E}[\chi_{A \times B}(Y)] \right| < \epsilon.$$

To do so, we randomly generate a set \tilde{A} as follows: For each i , w.p. $|V_i \cap A|/|V_i|$ include all elements of V_i in \tilde{A} , otherwise include none of the elements in \tilde{A} . Similarly generate a random \tilde{B} . Note that $\mathbb{E}_Y[\chi_{A \times B}(Y)] = \mathbb{E}_{\tilde{A}, \tilde{B}, Y}[\chi_{\tilde{A} \times \tilde{B}}(Y)]$ since within every $V_i \times V_j$, $\Pr[Y = e]$ is constant for all $e \in V_i \times V_j$, and

$$\frac{1}{|E|} \sum_{e \in A \times B} \text{weight}(e) = \frac{1}{|E|} \sum_{i,j} \frac{|V_i \cap A| \cdot |V_j \cap B| \cdot |(V_i \times V_j) \cap E|}{|V_i| \cdot |V_j|} = \mathbb{E}_{\tilde{A}, \tilde{B}, U_E}[\chi_{\tilde{A} \times \tilde{B}}(U_E)]$$

by linearity of expectation. Taking $S = \tilde{A}$, $T = \tilde{B}$ in Item 2 yields the required bound. □

6.3 Regularity Theorems for Circuit Complexity

In this section, we extend the notion of complexity to be computational and consider (boolean) circuit complexity. Let \mathcal{W} be the set of functions having low circuit complexity. Indeed, the highly constructive proof for Theorem 6.2 already provides a Y with low circuit complexity, as long as there exist approximate KL projections computed by small circuits. Thus we require \mathcal{V} to be *KL-projectable*:

Definition 6.5. Let \mathcal{V} be a convex set of distributions on $\{0, 1\}^n$. The ϵ -neighborhood of \mathcal{V} , denoted \mathcal{V}^ϵ , is the set of all distributions X on $\{0, 1\}^n$ such that for some $Y \in \mathcal{V}$ and for all $x \in \{0, 1\}^n$,

$$\Pr[X = x] \in [e^{-2\epsilon}, e^{2\epsilon}] \cdot \Pr[Y = x].$$

\mathcal{V} is said to be *KL-projectable* if for all $\epsilon > 0$, for every $X \in \mathcal{V}^\epsilon$ there exists some $Y \in \mathcal{V}$ such that

1. Y is an ϵ^2 -approximate KL projection of X on \mathcal{V} ;
2. If there is a size t circuit computing a measure M for X with outputs $M(x)$ of bit-length at most m , then there is a size $t + \text{poly}(m, \log(1/\epsilon))$ circuit M' computing a measure for Y with outputs $M'(x)$ of bit-length at most $m + \text{polylog}(1/\epsilon)$. (Recall that measures are $[0, 1]$ bounded, unnormalized mass functions; see Definition 4.6.)

Many natural convex sets of distributions are KL-projectable. Examples include the set of distributions with min-entropy at least k (Theorem A.3) and the set of distributions with Shannon entropy at least k (see [VZ]), for any $k > 0$.

We show that every distribution $X \in \mathcal{V}$ is indistinguishable, by all small circuits, to some $Y \in \mathcal{V}$ that has low circuit complexity, as long as \mathcal{V} is permutation-invariant and KL-projectable:

Theorem 6.6 (A regularity theorem for circuit complexity). *Let \mathcal{V} be a KL-projectable, permutation-invariant convex set of distributions on $\{0, 1\}^n$, $t > 0$, and $\epsilon > 0$. Then for every distribution $X \in \mathcal{V}$ there exists $Y \in \mathcal{V}$ such that*

1. X and Y are $O(\epsilon)$ -indistinguishable by size t circuits;
2. Y has low complexity: Y has a measure of circuit size $t' = S \cdot t + \text{poly}(S, \log(1/\epsilon))$, for $S = (n - H(X))/\epsilon^2$.

Proof. The proof is essentially the same as Theorem 6.2. Suppose for contradiction that for every low complexity $Y \in \mathcal{V}$ there is some size t circuit W such that $\mathbb{E}[W(X)] - \mathbb{E}[W(Y)] \geq \epsilon'$, where $\epsilon' = c \cdot \epsilon$ for a sufficiently large constant c . We will apply Theorem 3.1 (Uniform Min-Max Theorem), with

- $\mathcal{V} = \mathcal{V}$;
- $\mathcal{W} = \{(\text{deterministic}) \text{ circuits of size } t\}$;
- $f(z, W) = \mathbb{E}[W(X)] - W(z)$.

This corresponds to the two-player zero-sum game where Player 1 chooses some distribution $Y \in \mathcal{V}$, and Player 2 chooses a t sized circuit W , with expected payoff $F(Y, W) = \mathbb{E}[W(X)] - \mathbb{E}[W(Y)]$ for Player 2. We implement Algorithm 3.1 (Finding Universal Strategy) with KL projection on the set \mathcal{V} as follows. Start with an initial distribution $V^{(1)}$ that is uniform on $\{0, 1\}^n$ (which lies in \mathcal{V} , by Lemma 6.1). In each of the $S = (n - H(X))/\epsilon^2$ iterations we represent the distribution $V^{(i)}$ by a circuit $M^{(i)}$ computing a measure for $V^{(i)}$, where $M^{(i)}(x)$ has bit-length at most $i \cdot \text{polylog}(1/\epsilon)$. We implement the i th iteration as follows. For technical convenience we assume that $e^{-\epsilon}$ has bit-length at most $\log(1/\epsilon)$ (if not, we replace ϵ by some $\tilde{\epsilon} = O(\epsilon) > \epsilon$ such that $e^{-\tilde{\epsilon}}$ has bit-length at most $\log(1/\epsilon)$).

1. **Obtaining Player 2's Response $W^{(i)}$:** Suppose that we have constructed a $t_i \leq t'$ sized circuit $M^{(i)}$ computing a measure for $V^{(i)}$, and $M^{(i)}(x)$ has bit-length at most $i \cdot \text{polylog}(1/\epsilon)$. By assumption, there is a size t circuit $W^{(i)}$ such that

$$\mathbb{E}[W^{(i)}(X)] - \mathbb{E}[W^{(i)}(V^{(i)})] \geq \epsilon'.$$

2. **Weight Update:** We represent the resulting distribution $V^{(i)'}$ by the circuit $M^{(i)'}(x) = \exp(-\epsilon \cdot (1 - W^{(i)}(x))) \cdot M^{(i)}(x)$ that computes a measure for $V^{(i)'}$. Since $W^{(i)}(x) \in \{0, 1\}$, $\exp(-\epsilon \cdot (1 - W^{(i)}(x)))$ has bit-length at most $\log(1/\epsilon)$. $M^{(i)}(x)$ has bit-length at most $i \cdot \text{polylog}(1/\epsilon)$, thus multiplication takes time $i \cdot \text{polylog}(1/\epsilon)$. Thus $M^{(i)'}$ has circuit size $t'_i = t_i + t + i \cdot \text{polylog}(1/\epsilon)$, and bit-length at most $i \cdot \text{polylog}(1/\epsilon) + \log(1/\epsilon)$.
3. **KL Projection:** By KL-projectability of \mathcal{V} and the fact that $V^{(i)'} \in \mathcal{V}^\epsilon$, we have a circuit $M^{(i+1)}$ computing a measure for $V^{(i+1)}$ of size $t_{i+1} = t'_i + \text{poly}(i \cdot \text{polylog}(1/\epsilon), \log(1/\epsilon))$, and $M^{(i+1)}(x)$ has bit-length at most $i \cdot \text{polylog}(1/\epsilon) + \log(1/\epsilon) + \text{polylog}(1/\epsilon) = (i+1) \cdot \text{polylog}(1/\epsilon)$.

Note that $t_1 = O(1)$ and $t_{i+1} = t_i + t + \text{poly}(i \cdot \text{polylog}(1/\epsilon), \log(1/\epsilon))$, thus $t_i \leq S \cdot t + \text{poly}(S, \log(1/\epsilon))$ and the assumption that $t_i \leq t'$ is satisfied for all $i \in [S]$. By Theorem 3.1, W^* (the uniform distribution over $W^{(1)}, \dots, W^{(S)}$) satisfies

$$\mathbb{E}[W^*(X)] - \mathbb{E}[W^*(V)] \geq \epsilon' - O(\epsilon) > 0.$$

for all Player 1 strategies $V \in \mathcal{V}$ such that $H(V) \geq H(X)$. Taking $V = X$ yields a contradiction. \square

Remark. Most of our results in this section (Theorem 6.6, 6.8) hold not just for small circuits, but for an arbitrary class of distinguishers \mathcal{W} (like our Theorem 6.2) with a suitable definition of “complexity w.r.t. \mathcal{W} ” (see [TTV] Theorem 1.1 for one such example). However, we avoid stating results in greater generality since the appropriate definition of “complexity” may vary depending on the choice of \mathcal{V} (to account for the complexity of KL projections) and the application.

The above theorem also has an average-case variant. Rather than stating it in full generality (which would involve new definitions with the proof being essentially the same), we only state a special case where the low-complexity approximation is not confined in a convex set \mathcal{V} (i.e. \mathcal{V} is the universe). To express nonuniform complexity in the average-case setting, we extend the definition of measures:

Definition 6.7 (Conditional measure). For a joint distribution (X, C) , a function M is a *conditional measure for $C|X$* if for all $x \in \text{supp}(X)$, the function $f(y) = M(x, y)$ is a measure for $C|_{X=x}$.

Theorem 6.8 (A regularity theorem for circuit complexity – average case). *For every $t > 0$, $\epsilon > 0$, and joint distribution (X, B) over $\{0, 1\}^n \times \{0, 1\}^\ell$, there is a joint distribution (X, C) over $\{0, 1\}^n \times \{0, 1\}^\ell$ such that*

1. (X, B) and (X, C) are $O(\epsilon)$ -indistinguishable by size t circuits;
2. C has low complexity given X : $C|X$ has a conditional measure of circuit size $S \cdot t + (S \cdot \log(1/\epsilon))^2$, where $S = \ell/\epsilon^2$.

Proof. The proof is identical to Theorem 6.6, except we use the Uniform Min-Max Theorem – Average Case (Theorem 3.2), and do not need KL projections (thus the complexity of $C|X$ is lower compared to Theorem 6.6). \square

Theorem 6.8 is a slight strengthening of a recent result of Pietrzak and Jetchev [PJ], with a simpler proof. In [PJ] they obtain an (X, C) where given x , $C|_{X=x}$ is samplable by a circuit of size $O\left(\left(2^\ell \cdot \ell \cdot (1/\epsilon)^2 \log(1/\epsilon)\right)^2 \cdot t\right)$. Theorem 6.8 provides an (X, C) where given x , $C|_{X=x}$ can

be sampled by a circuit of size $O\left(2^\ell \cdot \left(\ell \cdot (1/\epsilon)^2 \cdot t + (\ell \cdot (1/\epsilon)^2 \log(1/\epsilon))^2\right)\right)$ by computing $M(x, y)$ for all $y \in \{0, 1\}^\ell$ and sampling $C|_{X=x}$ using its mass function.

Finally, we show an application of Theorem 6.8: deducing a technical lemma of [GW], which says if X and U are indistinguishable then for any short B jointly distributed with X , there is some C (the “auxilliary information”) jointly distributed with U such that (X, B) and (U, C) are indistinguishable. Not only is our proof simpler, but also it guarantees that C has low circuit complexity given U . This “low complexity” version (with slightly weaker parameters) was initially proved by Pietrzak and Jetchev [PJ], and an interactive extension was proved by Chung, Lui, and Pass [CLP] for applications in the context of distributional zero-knowledge.

Lemma 6.9 (Low circuit complexity version of [GW] Lemma 3.1). *Let X and U be distributions over $\{0, 1\}^n$, and B be a distribution over $\{0, 1\}^\ell$ jointly distributed with X . Suppose X and U are ϵ -indistinguishable by circuits of size t . Then there exists some $C \in \{0, 1\}^\ell$ jointly distributed with U such that:*

1. (X, B) and (U, C) are 2ϵ -indistinguishable by circuits of size $s = t/(2^\ell \cdot \ell \cdot (1/\epsilon)^2) - \ell \cdot ((1/\epsilon) \log(1/\epsilon))^2$.
2. C has low complexity given U : $C|U$ has a conditional measure of circuit size $\ell \cdot (1/\epsilon)^2 \cdot s + (\ell \cdot (1/\epsilon)^2 \log(1/\epsilon))^2$.

Proof. We first apply Theorem 6.8 to obtain a distribution $(X, P(X))$ such that (X, B) and $(X, P(X))$ are ϵ -indistinguishable by size s circuits, where P is a randomized function, and there is a size $\ell \cdot (1/\epsilon)^2 \cdot s + (\ell \cdot (1/\epsilon)^2 \log(1/\epsilon))^2$ circuit M computing a conditional measure for $P(X)|X$. Thus, given x , $P(x)$ can be sampled in time $s' = O\left(2^\ell \cdot \left(\ell \cdot (1/\epsilon)^2 \cdot s + (\ell \cdot (1/\epsilon)^2 \log(1/\epsilon))^2\right)\right)$ by computing $M(x, y)$ for all $y \in \{0, 1\}^\ell$ and sampling $P(x)$ from its mass function.

Let $C = P(U)$. Since P is efficient, indistinguishability of X and U implies that $(X, P(X))$ and $(U, P(U))$ are ϵ -indistinguishable by circuits of size s . (Otherwise, given an s -sized ϵ -distinguisher D for $(X, P(X))$ and $(U, P(U))$ we get an ϵ -distinguisher $T(x) = D(x, P(x))$ for X and U , of circuit size $O(s + s') \leq t$.) By triangle inequality, (X, B) and $(U, P(U)) = (U, C)$ must be 2ϵ -indistinguishable by circuits of size s . \square

6.4 Regularity Theorems for Time Complexity

In this section, we use the full strength of the Uniform Min-Max Theorem to obtain a low complexity approximation where complexity is measured using *uniform* algorithms. For simplicity, we only prove an average-case variant (Theorem 6.10) where no KL-projection is needed, which is the uniform analogue of Theorem 6.8. As an immediate corollary, we provide a “sampling” version of it (Theorem 6.11), which is cleaner and convenient for several applications, but involves exponential dependence on ℓ .

Theorem 6.10 (A regularity theorem for time complexity – average case). *Let n be a security parameter, $\ell = \ell(n)$, $t = t(n) \geq n$, $\epsilon = \epsilon(n) > 0$ all computable in $\text{poly}(n)$ time. Let $(X, B) = (X, B)(n)$ be a joint distribution on $\{0, 1\}^n \times \{0, 1\}^\ell$. Let A be a t -time randomized oracle algorithm. Then there is a $t' = \text{poly}(t, 1/\epsilon)$ -time randomized algorithm R such that w.p. $\Omega(\epsilon^2/\ell)$ over $M \leftarrow R(1^n)$ and $W \leftarrow A^M(1^n)$, if we interpret M as a deterministic circuit computing a conditional measure for $C|X$ and W as a randomized circuit $W : \{0, 1\}^n \times \{0, 1\}^\ell \rightarrow [0, 1]$, we have:*

$$\mathbb{E}[W(X, B)] - \mathbb{E}[W(X, C)] < \epsilon.$$

Proof. We will let R be an implementation of Algorithm 3.2 (Finding Universal Strategy – Average Case), using A as a subroutine. We then show R satisfies the desired properties by applying Theorem 3.2 (Uniform Min-Max Theorem – Average Case), with

- \mathcal{V} being the set of all joint distributions (X, C) on $\{0, 1\}^n \times \{0, 1\}^\ell$ (where the marginal distribution of X is fixed, and C may vary);
- $\mathcal{W} = \{\text{randomized circuits of size } t\}$;
- $f((x, y), W) = \mathbb{E}[W(X, B)] - \mathbb{E}[W(x, y)]$.

This corresponds to the two-player zero-sum game where Player 1 selects a distribution $(X, C) \in \mathcal{V}$, Player 2 selects a size t circuit W and receives expected payoff $F((X, C), W) = \mathbb{E}[W(X, B)] - \mathbb{E}[W(X, C)]$.

Our implementation of Algorithm 3.2 using A is as follows. We set the ϵ in Algorithm 3.2 to be $\epsilon' = \epsilon/c$ for a sufficiently large constant c , and start with an initial distribution $(X, C^{(1)}) = (X, U_\ell)$ (where U_ℓ is independent of X). In each of the $S = O(\ell/\epsilon^2)$ iterations we represent $C^{(i)}$ by a circuit $M^{(i)}$ computing a conditional measure for $C^{(i)}|X$, i.e. $M^{(i)}(x, y) \propto \Pr[C^{(i)} = y|X = x]$. So we can take $M^{(1)}(x, y) = 1$ for all x, y . We implement the i th iteration as follows, with $\gamma = 1/3S$:

1. **Obtaining Player 2's Response $W^{(i)}$:** Suppose that we have constructed a t_i -size circuit $M^{(i)}$ where $M^{(i)}(x, y)$ has bit-length $i \cdot \text{polylog}(1/\epsilon)$. There are two steps.

- (a) We run $A^{M^{(i)}}(1^n)$ to obtain a t -size randomized circuit $\widehat{W}^{(i)}$, and convert it into a $O(tm)$ -size deterministic circuit $\widetilde{W}^{(i)}$ by hardwiring $m = O((1/\epsilon^2) \log(1/\gamma))$ samples of the coins of $\widehat{W}^{(i)}$, so that w.p. at least $1 - \gamma$,

$$\mathbb{E}[\widetilde{W}^{(i)}(X, B)] - \mathbb{E}[\widetilde{W}^{(i)}(X, C^{(i)})] \geq \mathbb{E}[\widehat{W}^{(i)}(X, B)] - \mathbb{E}[\widehat{W}^{(i)}(X, C^{(i)})] - \epsilon'.$$

- (b) Our choice of $W^{(i)}$ is the following approximation to $\widetilde{W}^{(i)}$, so that $\exp(-\epsilon' \cdot (1 - W^{(i)}(x, y)))$ can be computed precisely and efficiently. First, we use Newton's method to compute a $\text{polylog}(1/\epsilon)$ -bit approximation $E(x, y) \in (0, 1]$ of $\exp(-\epsilon' \cdot (1 - \widetilde{W}^{(i)}(x, y)))$ within $\pm\epsilon'^2$ error, in time $O(tm) + \text{polylog}(1/\epsilon)$. We define $W^{(i)}$ to be such that $\exp(-\epsilon' \cdot (1 - W^{(i)}(x, y))) = E(x, y)$. Thus $|W^{(i)}(x, y) - \widetilde{W}^{(i)}(x, y)| \leq \epsilon'$, and

$$\mathbb{E}[W^{(i)}(X, B)] - \mathbb{E}[W^{(i)}(X, C^{(i)})] \geq \mathbb{E}[\widetilde{W}^{(i)}(X, B)] - \mathbb{E}[\widetilde{W}^{(i)}(X, C^{(i)})] - 2\epsilon'.$$

2. **Weight Update:** We represent the resulting distribution $C^{(i+1)}$ by the circuit $M^{(i+1)}(x, y) = \exp(-\epsilon' \cdot (1 - W^{(i)}(x, y))) \cdot M^{(i)}(x, y)$ computing a conditional measure for $C^{(i+1)}|X$. Since $\exp(-\epsilon' \cdot (1 - W^{(i)}(x, y))) = E(x, y)$ has bit-length $\text{polylog}(1/\epsilon)$ and $M^{(i)}(x, y)$ has bit-length $i \cdot \text{polylog}(1/\epsilon)$, multiplication takes time $i \cdot \text{polylog}(1/\epsilon)$. Thus $M^{(i+1)}$ has circuit size $t_{i+1} = t_i + O(tm) + i \cdot \text{polylog}(1/\epsilon)$ and bit-length $(i + 1) \cdot \text{polylog}(1/\epsilon)$, and can be constructed in similar time.
3. **KL projection:** Do nothing as Player 1 strategies can be arbitrary conditional distributions $C^{(i)}|_{X=x}$.

Now let R be the algorithm that chooses a random $i \leftarrow [S]$, runs the above implementation of Algorithm 3.2 for $i - 1$ iterations to construct and output $M^{(i)}$. Since $t_1 = O(1)$, we have $t_i = O(1) + S \cdot (O(tm) + S \cdot \text{polylog}(1/\epsilon))$ for all $i \in [S]$. Thus R runs in total time $\text{poly}(t, S, m, \log(1/\epsilon)) \leq t'$.

Suppose for contradiction that w.p. at least $1 - \gamma$ over coins of R used to generate $M^{(i)}$ and $A, A^{M^{(i)}}(1^n)$ outputs a randomized circuit $\widehat{W}^{(i)}$ s.t. $\mathbb{E}[\widehat{W}^{(i)}(X, B)] - \mathbb{E}[\widehat{W}^{(i)}(X, C^{(i)})] \geq \epsilon$. By union bound w.p. at least $1 - 2\gamma \cdot S = 1/3$, in all iterations we have

$$\mathbb{E}[W^{(i)}(X, B)] - \mathbb{E}[W^{(i)}(X, C^{(i)})] \geq \mathbb{E}[\widehat{W}^{(i)}(X, B)] - \mathbb{E}[\widehat{W}^{(i)}(X, C^{(i)})] - 3\epsilon' \geq \epsilon - 3\epsilon'.$$

Let W^* be the uniform distribution over $W^{(1)}, \dots, W^{(S)}$. By the Uniform Min-Max Theorem – Average Case (Theorem 3.2), w.p. at least $1/3$, W^* satisfies

$$\mathbb{E}[W^*(X, B)] - \mathbb{E}[W^*(X, C)] \geq \epsilon - 3\epsilon' - O(\epsilon') > 0$$

for all Player 1 strategies $(X, C) \in \mathcal{V}$. Taking $(X, C) = (X, B)$ yields a contradiction. \square

As an immediate corollary, we obtain a “sampling” version, which is cleaner, and convenient for several applications. Recall that for a distribution Z , we denote by O_Z the sampling oracle of Z , i.e. on each query O_Z returns a random sample of Z .

Theorem 6.11 (A regularity theorem for time complexity – average case (sampling version)). *Let n be a security parameter, $\ell = \ell(n)$, $t = t(n) \geq n$, $\epsilon = \epsilon(n) > 0$ all computable in $\text{poly}(n)$ time. Let $(X, B) = (X, B)(n)$ be a joint distribution on $\{0, 1\}^n \times \{0, 1\}^\ell$, and $Q = Q(n)$ be any $\text{poly}(n)$ -time samplable distribution on $\{0, 1\}^n$. Let A be a t -time randomized oracle algorithm. Then there is a $t' = \text{poly}(2^\ell, t, 1/\epsilon)$ -time randomized algorithm R that w.p. at least $\Omega(\epsilon^2/\ell)$ outputs a randomized circuit P of size at most t' satisfying:*

$$\mathbb{E}[A^{O_Q, P(Q)}(X, B)] - \mathbb{E}[A^{O_Q, P(Q)}(X, P(X))] < \epsilon.$$

Proof. Given a t -time randomized oracle algorithm A , we define a $2^\ell \cdot \text{poly}(t, 1/\epsilon)$ -time randomized oracle algorithm A' to which we apply Theorem 6.10, as follows. First define the randomized function $\widehat{A}(x, y; a)$ to equal $A(x, y)$ where we fix the outputs of the sampling oracle to be $a \in (\{0, 1\}^n \times \{0, 1\}^\ell)^t$. For every $M : \{0, 1\}^n \times \{0, 1\}^\ell \rightarrow [0, 1]$, let $A'^M(1^n)$ generate $a^{(1)}, \dots, a^{(m)}$ as $m = O((1/\epsilon^2) \cdot \log(c\ell/\epsilon^2))$ random samples of $(Q, P_M(Q))^t$, where P_M is the randomized function such that M is a conditional measure for $P_M(Q)|Q$, and c is a constant to be determined later. Recall that Q is $\text{poly}(n)$ -time samplable by assumption, and we can construct from M a circuit that samples $(Q, P_M(Q))$ by computing $M(x, y)$ for all $y \in \{0, 1\}^\ell$. We then let $A'^M(1^n)$ output a randomized circuit $W(x, y)$ computing the average of $\widehat{A}(x, y; a^{(i)})$ over all i . By a Chernoff bound, w.p. at least $\epsilon^2/c\ell$ over $W \leftarrow A'^M(1^n)$ we have

$$\mathbb{E}[W(X, B)] - \mathbb{E}[W(X, P_M(X))] \geq \mathbb{E}\left[A^{O_Q, P_M(Q)}(X, B)\right] - \mathbb{E}\left[A^{O_Q, P_M(Q)}(X, P_M(X))\right] - \epsilon/2.$$

By applying Theorem 6.10 to A' , there is a $\text{poly}(2^\ell, t, 1/\epsilon)$ -algorithm R such that w.p. $\Omega(\epsilon^2/\ell)$ over $M \leftarrow R(1^n)$ and $W \leftarrow A'^M(1^n)$ we have

$$\mathbb{E}[W(X, B)] - \mathbb{E}[W(X, C)] < \epsilon/2.$$

Thus w.p. at least $\Omega(\epsilon^2/\ell) - \epsilon^2/c\ell = \Omega(\epsilon^2/\ell)$ (for a sufficiently large c) over $M \leftarrow R(1^n)$,

$$\mathbb{E}[A^{O_{Q,P_M(Q)}}(X, B)] - \mathbb{E}[A^{O_{Q,P_M(Q)}}(X, P_M(X))] < \epsilon.$$

□

We now apply Theorem 6.11 to show Theorem 6.12, the uniform analogue of Theorem 6.9 (which in turn is the low circuit complexity version of [GW] Lemma 3.1). We do so mainly because it is convenient for applications, including (i) deriving a uniform Dense Model Theorem (see Section 7, Theorem 7.3); (ii) showing impossibility of constructing succinct non-interactive arguments (SNARGs) via black-box reductions under uniform hardness assumptions (see Section 8, Theorem 8.7).

Theorem 6.12 (Low time complexity version of [GW] Lemma 3.1). *Let n be a security parameter, $\ell = \ell(n)$, $s = s(n) \geq n$, $\epsilon = \epsilon(n) > 0$ all computable in $\text{poly}(n)$ time. Let $X = X(n)$ and $U = U(n)$ be $\text{poly}(n)$ -time samplable distributions on $\{0, 1\}^n$ that are ϵ -indistinguishable for s -time randomized algorithms. Let $B = B(n)$ be a distribution on $\{0, 1\}^\ell$ jointly distributed with X , and let $Q = Q(n)$ be any $\text{poly}(n)$ -time samplable distribution on $\{0, 1\}^n$. Let A be a t -time randomized oracle algorithm, for $t = s^{\Omega(1)}/\text{poly}(2^\ell, 1/\epsilon)$. Then there is a $t' = \text{poly}(2^\ell, t, 1/\epsilon)$ -time randomized algorithm R such that w.p. at least $\Omega(\epsilon^2/\ell)$, R outputs a randomized circuit P satisfying*

$$\mathbb{E}[A^{O_{Q,P(Q)}}(X, B)] - \mathbb{E}[A^{O_{Q,P(Q)}}(U, P(U))] < 2\epsilon.$$

Proof. By Theorem 6.11, there is a t' -time algorithm R that w.p. at least $\gamma = \Omega(\epsilon^2/\ell)$ outputs a randomized circuit P satisfying

$$\mathbb{E}[A^{O_{Q,P(Q)}}(X, B)] - \mathbb{E}[A^{O_{Q,P(Q)}}(X, P(X))] < 0.9\epsilon.$$

Since P is efficient, ϵ -indistinguishability of X and U implies that with probability at least $1 - \gamma/2$ over P ,

$$\mathbb{E}[A^{O_{Q,P(Q)}}(X, P(X))] - \mathbb{E}[A^{O_{Q,P(Q)}}(U, P(U))] < 1.1\epsilon.$$

Indeed, suppose that $A^{O_{Q,P(Q)}}$ achieves distinguishing advantage at least 1.1ϵ w.p. at least $\gamma/2$ over P , then we could obtain an ϵ -distinguisher for X and U by running R for $O((1/\gamma) \log(1/\epsilon))$ times, each time testing the distinguisher $T(x) = A^{O_{Q,P'(Q)}}(x, P'(x))$ where P' is the randomized circuit output by R (by running on $O((1/\epsilon^2) \log(1/\epsilon))$ random samples of X, U and $(Q, P'(Q))$), and finally taking the best one. This yields an ϵ -distinguisher for X and U that runs in time $O((1/\epsilon^2) \log(1/\epsilon) \cdot (\text{poly}(n) + (1/\gamma) \log(1/\epsilon) \cdot (t + \text{poly}(t')))) \leq s$, violating their indistinguishability.

Combining the two inequalities, we get with probability at least $\gamma/2$ over $P \leftarrow R$,

$$\mathbb{E}[A^{O_{Q,P(Q)}}(X, B)] - \mathbb{E}[A^{O_{Q,P(Q)}}(U, P(U))] < 2\epsilon.$$

□

7 Application: Dense Model Theorem

A celebrated result of Green and Tao [GT] shows that there exist arbitrarily long arithmetic progressions of prime numbers. A key new component of their proof is the Dense Model Theorem

which, in the generalized form of Tao and Ziegler [TZ], says if X is a pseudorandom distribution and D is a distribution dense in X , then D is indistinguishable to a distribution M that is dense in the uniform distribution. Like our results in Section 6.1, notions of indistinguishability and pseudorandomness in the Dense Model Theorem can be defined with respect to an arbitrary class of distinguishers \mathcal{W} , and are not restricted to classes of circuit distinguishers.

In the original proof, the indistinguishability (i.e. the bound on distinguishing probability) between D and M is exponentially larger than the indistinguishability between X and the uniform distribution, making it inapplicable for the typical complexity-theoretic or cryptographic settings of parameters. Using the Min-Max Theorem, Reingold et al. [RTTV] provided another proof where the indistinguishability and complexity blow-ups are only polynomial; a similar proof was given by Gowers [Gow]. These requirements are crucial for applications in leakage-resilient cryptography [DP2, DP1, FOR], and for connections to computational differential privacy [MPRV].

We now state a Dense Model Theorem due to Zhang [Zha], where the complexity blow-up $O((\delta/\epsilon)^2 \log(1/\delta))$ is asymptotically optimal.⁴

Recall from Definition 4.1 that for distributions X and Y on Σ , we say X is δ -dense in Y if $\Pr[Y = x] \geq \delta \cdot \Pr[X = x]$ for all $x \in \Sigma$, and say X is δ -dense if it is δ -dense in U_Σ . It will be convenient to denote by $\mathbf{Th}_t(x)$ the boolean threshold function i.e. $\mathbf{Th}_t(x) = 1$ if $x \geq t$ and $\mathbf{Th}_t(x) = 0$ if $x < t$.

Theorem 7.1 (Dense Model Theorem [Zha]). *Let Σ be a finite set, \mathcal{W} be an arbitrary class of functions $W : \Sigma \rightarrow [0, 1]$, $\epsilon > 0$, $\delta > 0$. Then the following holds for some $S = O((\delta/\epsilon)^2 \log(1/\delta))$.*

Let \mathcal{W}' be the set of all functions $W' : \Sigma \rightarrow \{0, 1\}$ defined by $W'(x) = \mathbf{Th}_t\left(\sum_{i=1}^S W_i(x)/S\right)$, for some $W_1, \dots, W_S \in \mathcal{W}$ and $t \in [0, 1]$. Let X be a distribution on Σ that is ϵ -indistinguishable from U_Σ by \mathcal{W}' . Let D be a distribution δ -dense in X . Then there is a δ -dense distribution M such that D and M are $O(\epsilon/\delta)$ -indistinguishable by \mathcal{W} .

A Min-Max Theorem based proof with a suboptimal blow-up of $S = O((\delta/\epsilon)^2 \log(1/\epsilon))$ proceeds as follows. (Note that we may assume $\delta > \epsilon$, else the conclusion of $O(\epsilon/\delta)$ -indistinguishability is trivial.) Assume for contradiction that for every δ -dense M there is a distinguisher $W \in \mathcal{W}$. By the Min-Max Theorem there is a *universal* distinguisher W^* such that $\mathbb{E}[W^*(D)] - \mathbb{E}[W^*(M)] \geq O(\epsilon/\delta)$ for every δ -dense M . By subsampling we can assume that W^* is the average over a multiset of $O((\delta/\epsilon)^2 \log(1/\epsilon))$ elements of \mathcal{W} , while changing the distinguishing advantage by at most a constant fraction. Given such universal distinguisher W^* we can construct an ϵ -distinguisher in \mathcal{W}' between X and U_Σ , as formalized in Lemma 7.2:

Lemma 7.2 (Implicit in [RTTV]). *Let Σ be a finite set, $\epsilon > 0$, $\delta > 0$. Let X, D be distributions on Σ , and D is δ -dense in X . Let $W^* : \Sigma \rightarrow [0, 1]$ be a function such that for every δ -dense distribution M we have*

$$\mathbb{E}[W^*(D)] - \mathbb{E}[W^*(M)] \geq O(\epsilon/\delta).$$

Then for some t as a multiple of $O(\epsilon/\delta)$, we have

$$\mathbb{E}[\mathbf{Th}_t(W^*(X))] - \mathbb{E}[\mathbf{Th}_t(W^*(U_\Sigma))] \geq \epsilon.$$

⁴Zhang [Zha] shows optimality by proving a black-box lower bound on the number of elements of \mathcal{W} that a black-box reduction needs to obtain to construct a distinguisher between X and the uniform distribution.

This proves a Dense Model Theorem, but with a suboptimal complexity blow-up of $O((\delta/\epsilon)^2 \log(1/\epsilon))$ (due to the probabilistic construction of the multiset defining W^*). Zhang [Zha] achieved optimal blow-up in Theorem 7.1 by adapting the technique of multiplicative weights with KL projection from Barak, Hardt, and Kale [BHK].

Replacing the use of the Min-Max Theorem in the above argument by our Uniform Min-Max Theorem (Theorem 3.1), we immediately obtain a simple proof of Theorem 7.1, with an optimal complexity blow-up that comes from the setting of

$$S = \frac{(\log |\Sigma| - \min_{M \in \mathcal{V}} H(M))}{\Omega(\epsilon/\delta)^2} = O\left(\frac{\log(1/\delta)}{(\epsilon/\delta)^2}\right)$$

in Theorem 3.1, with \mathcal{V} being the set of δ -dense distribution on Σ . Compared to [Zha], the proof using the Uniform Min-Max Theorem is more modular, and avoids adapting the analysis of [HW] and [BHK] to the specific setting of the Dense Model Theorem.

In the rest of the section, we prove a Uniform Dense Model Theorem where the distinguishers are (uniform) algorithms rather than (nonuniform) $[0, 1]$ -valued functions. Rather than directly applying the Uniform Min-Max Theorem and using Lemma 7.2, we follow [TTV] and deduce the Dense Model Theorem from a regularity theorem. Specifically, [TTV] shows how to deduce the nonuniform Dense Model Theorem from a nonuniform regularity theorem like Theorem 6.6; we prove our uniform Dense Model Theorem using a uniform regularity theorem (Theorem 6.12).

We begin with an overview of the proof of the nonuniform Dense Model Theorem in [TTV]. The distribution D being δ -dense in X means that there is a (possibly inefficient) binary random variable B jointly distributed with X such that $D = X|_{B=1}$, and $\Pr[B = 1] \geq \delta$. By a regularity theorem, there is an efficient randomized function P such that (X, B) and $(X, P(X))$ are indistinguishable. Since P is efficient, indistinguishability of X and U_n implies that $(X, P(X))$ and $(U_n, P(U_n))$ are also indistinguishable. So we can take $M = U_n|_{P(U_n)=1}$. M is δ -dense because $\Pr[P(U_n) = 1] \approx \Pr[P(X) = 1]$, again by indistinguishability of X and U_n . (Note that we use indistinguishability of X and U_n twice. In the uniform setting, the uniform distinguisher will have to determine which case to use, by testing whether $\Pr[P(U_n) = 1] \approx \Pr[P(X) = 1]$ or not.)

Theorem 7.3 (Uniform Dense Model Theorem). *Let n be a security parameter, $\epsilon = \epsilon(n)$, $\delta = \delta(n)$, $s = s(n) \geq n$ all computable in $\text{poly}(n)$ time. Let $X = X(n)$ and $U = U(n)$ be poly-time samplable distributions on $\{0, 1\}^n$ such that X and U are ϵ -indistinguishable for s -time randomized algorithms. Let $D = D(n)$ be a distribution with density δ in X . Then for some $t = s^{\Omega(1)}/\text{poly}(1/\epsilon, 1/\delta)$ and all t -time randomized oracle algorithms A , there is a distribution $M = M(n)$ that is $(\delta - O(\epsilon))$ -dense in U such that for all n ,*

$$\mathbb{E}[A^{O_M}(D)] - \mathbb{E}[A^{O_M}(M)] \leq O(\epsilon/\delta).$$

Moreover, M is constructive: $M = U|_{P(U)=1}$ for some randomized circuit P such that some $\text{poly}(t, 1/\epsilon)$ -time randomized algorithm R outputs P w.p. at least $\Omega(1/\epsilon^2)$.

Proof. WLOG we assume $1 > \delta > c\epsilon$ for a sufficiently large constant c . D having density δ in X means that there is a (possibly inefficient) binary random variable B jointly distributed with X such that $D = X|_{B=1}$, and $\delta_B = \Pr[B = 1] = \delta$. Consider any t -time randomized oracle algorithm A . Let A' be the randomized oracle algorithm where for every joint distribution (U, C) over $\{0, 1\}^n \times \{0, 1\}$, $A'^{O_{U,C}}$ on input (x, y) does the following:

1. Compute an estimate $\hat{\delta}_C$ of $\delta_C = \Pr[C = 1]$ such that $|\hat{\delta}_C - \delta_C| \leq \epsilon$ w.p. at least $1 - \epsilon$. To do so we take $O((1/\epsilon^2) \log(1/\epsilon))$ random samples of (U, C) and let $\hat{\delta}_C$ be the fraction on which C equals 1.
2. If $\hat{\delta}_C < \delta_B - 5\epsilon$ then return y ; if $\hat{\delta}_C > \delta_B + 5\epsilon$ then return $1 - y$.
3. Otherwise, $|\hat{\delta}_C - \delta_B| \leq 5\epsilon$, and
 - (a) If $y = 0$ then return zero.
 - (b) If $y = 1$ then simulate $A^{O_N}(x)$ for the distribution $N = U|_{C=1}$, and return the output. To simulate $A^{O_N}(x)$, we obtain t random samples of N w.p. at least $1 - \epsilon$, where each sample is generated using rejection sampling from $O_{U,C}$ for $O((1/\delta_C) \log(t/\epsilon))$ times, where $\delta_C \geq \delta_B - 4\epsilon \geq \delta/2$.

A' runs in time $t' = t + O((1/\epsilon^2) \log(1/\epsilon)) \cdot \text{poly}(n) + O((1/\delta) \log(t/\epsilon)) \cdot t \cdot \text{poly}(n)$. By Theorem 6.12, there is a $\text{poly}(t', 1/\epsilon)$ -time randomized algorithm R that w.p. at least $\Omega(\epsilon^2)$ outputs a randomized circuit P satisfying

$$\begin{aligned}
2\epsilon &> \mathbb{E} [A'^{O_{U,P(U)}}(X, B)] - \mathbb{E} [A'^{O_{U,P(U)}}(U, P(U))] \\
&\geq \Pr_{\hat{\delta}_C} [\hat{\delta}_C - \delta_B > 5\epsilon] \cdot (\delta_C - \delta_B) + \Pr_{\hat{\delta}_C} [\hat{\delta}_C - \delta_B < -5\epsilon] \cdot (\delta_B - \delta_C) \\
&\quad + \Pr_{\hat{\delta}_C} [|\hat{\delta}_C - \delta_B| \leq 5\epsilon] \cdot (\delta_B \cdot \mathbb{E}[A^{O_M}(D)] - \delta_C \cdot \mathbb{E}[A^{O_M}(M)] - \epsilon). \tag{1}
\end{aligned}$$

Take $(U, C) = (U, P(U))$ and $M = U|_{C=1}$. We claim that $\delta_C \geq \delta_B - 6\epsilon$, i.e. M is $(\delta - O(\epsilon))$ -dense in U . Indeed, if $\delta_C < \delta_B - 6\epsilon$ then a Chernoff bound implies

$$\Pr_{\hat{\delta}_C} [\hat{\delta}_C - \delta_B > 5\epsilon] \cdot (\delta_C - \delta_B) + \Pr_{\hat{\delta}_C} [\hat{\delta}_C - \delta_B < -5\epsilon] \cdot (\delta_B - \delta_C) > 5\epsilon$$

violating Eq. 1. By symmetry, we must have $\delta_C \in [\delta_B - 6\epsilon, \delta_B + 6\epsilon]$.

We now show that D and M are indistinguishable by A^{O_M} . Suppose that $\delta_C \in [\delta_B, \delta_B + 6\epsilon]$ (the case $\delta_C \in [\delta_B - 6\epsilon, \delta_B]$ is similar). Then $\Pr_{\hat{\delta}_C} [\hat{\delta}_C - \delta_B < -5\epsilon] \leq \epsilon$ and Eq. 1 implies

$$\begin{aligned}
2\epsilon &\geq \left(1 - \epsilon - \Pr_{\hat{\delta}_C} [|\hat{\delta}_C - \delta_B| \leq 5\epsilon]\right) \cdot (\delta_C - \delta_B) - \epsilon \\
&\quad + \Pr_{\hat{\delta}_C} [|\hat{\delta}_C - \delta_B| \leq 5\epsilon] \cdot (\delta_B (\mathbb{E}[A^{O_M}(D)] - \mathbb{E}[A^{O_M}(M)]) - (\delta_C - \delta_B) - \epsilon)
\end{aligned}$$

which simplifies to

$$\delta_B \cdot (\mathbb{E}[A^{O_M}(D)] - \mathbb{E}[A^{O_M}(M)]) < \frac{3\epsilon - (1 - \epsilon)(\delta_C - \delta_B)}{\Pr [|\hat{\delta}_C - \delta_B| \leq 5\epsilon]} + 2(\delta_C - \delta_B) + \epsilon. \tag{2}$$

Consider these cases:

- If $0 \leq \delta_C - \delta_B < 4\epsilon$, then $\Pr \left[\left| \hat{\delta}_C - \delta_B \right| \leq 5\epsilon \right] \geq 1 - \epsilon$ hence RHS of Eq. 2 is at most $3\epsilon/(1 - \epsilon) + (2 - (1 - \epsilon)/(1 - \epsilon))(\delta_C - \delta_B) + \epsilon \leq O(\epsilon)$.
- If $4\epsilon \leq \delta_C - \delta_B \leq 6\epsilon$, then RHS of Eq. 2 is at most $2(\delta_C - \delta_B) + \epsilon \leq O(\epsilon)$.

Thus we conclude that $\mathbb{E}[A^{OM}(D)] - \mathbb{E}[A^{OM}(M)] \leq O(\epsilon)/\delta_B \leq O(\epsilon/\delta)$.

□

8 Application: Impossibility of Black-Box Construction of Succinct Non-Interactive Argument

A result of Gentry and Wichs [GW] shows that there is no black-box construction of succinct non-interactive arguments (SNARGs) from any natural cryptographic assumption (formally, they consider *falsifiable* cryptographic assumptions: ones that are defined by a polynomial-time security game). Their result relies on the (mild) assumption that there exist *hard subset membership problems*, which is equivalent to the existence of subexponentially hard one-way functions. One limitation is that they need to work in the non-uniform setting, in part due to their use of the Min-Max Theorem (in [GW] Lemma 3.1). In this section we show how to obtain the analogous result in the *uniform setting* by using the Uniform Min-Max Theorem. More specifically, assuming that there exist subexponentially hard one-way functions that are secure against uniform algorithms, we show that there is no black-box construction of SNARGs based on cryptographic assumptions where security is measured against uniform algorithms (unless the assumption is already false).

A succinct non-interactive argument (SNARG) is a non-interactive argument system where the proof size is bounded by a fixed polynomial, for all instances and witnesses whose size can be an arbitrarily large polynomial. Formally,

Definition 8.1 (SNARG). Let L be an **NP** language associated with relation R . We say that a tuple (G, P, V) of probabilistic polynomial-time (PPT) algorithms is a *succinct non-interactive argument for R* if the following properties hold:

- **Completeness:** For all $(x, w) \in R$, if we choose $(\text{CRS}, \text{PRIV}) \leftarrow G(1^n), \Pi \leftarrow P(\text{CRS}, x, w)$, then

$$\Pr [V(\text{PRIV}, x, \Pi) = 0] = \text{negl}(n).$$

- **Soundness:** For every PPT algorithm (efficient *adversary*) A , if we choose $(\text{CRS}, \text{PRIV}) \leftarrow G(1^n), (X, \Pi) \leftarrow A(1^n, \text{CRS})$, then

$$\Pr [V(\text{PRIV}, X, \Pi) = 1 \wedge X \notin L] = \text{negl}(n).$$

- **Succinctness:** For all $(x, w) \in \text{supp}(X, W)$ and $\text{crs} \in \text{supp}(\text{CRS})$, the length of the proof $\pi = P(\text{crs}, x, w)$ is $|\pi| = \text{poly}(n)(|x| + |w|)^{o(1)}$. We also consider a weaker variant called *slightly succinct*, where we require the length of a proof to be $|\pi| = \text{poly}(n)(|x| + |w|)^\alpha + o(|x| + |w|)$ for some constant $\alpha < 1$.⁵

⁵Earlier versions of [GW] contained a minor bug in the definition of slight succinctness. We use the corrected definition from the current version of their paper.

Our notion of a falsifiable cryptographic assumption is analogous to [GW], except that the adversary A is a *uniform* algorithm instead of circuit:

Definition 8.2 (Falsifiable assumption). Given an interactive PPT algorithm Chal (the challenger), the *uniform falsifiable (cryptographic) assumption (associated with)* Chal states that for all (uniform) PPT algorithms H , the probability that Chal(1^n) outputs a special symbol win after interacting with $H(1^n)$ is at most $\text{negl}(n)$ for all sufficiently large n .

For any randomized (possibly inefficient) function H , we let $\text{Break}_H(n)$ denote the above probability and say that H *breaks the assumption* if $\text{Break}_H(n) \geq 1/\text{poly}(n)$ for infinitely many n .

Remark. An alternative definition of falsifiable assumption allows specifying a constant β , and says that the probability Chal(1^n) outputs win is at most $\beta + \text{negl}(n)$. However, it turns out that setting $\beta = 0$, i.e. our definition above, is without loss of generality [HH]. We adopt the simpler definition because it is convenient for our proof.

Next we define black-box reductions:

Definition 8.3 (Adversary and reduction). For a randomized function A and a constant $c \in \mathbb{N}$, we say (A, c) is a (G, P, V) -*adversary* if $|A(1^n, \text{crs})| \leq n^c$ and A violates the soundness condition infinitely often, i.e. if we choose $(\text{CRS}, \text{PRIV}) \leftarrow G(1^n)$, $(X, \Pi) \leftarrow A(1^n, \text{CRS})$, then

$$\Pr[V(\text{PRIV}, X, \Pi) = 1 \wedge X \notin L] \geq n^{-c}$$

for infinitely many n . We say (A, c) is an *a.e. (G, P, V)-adversary* if A violates soundness for all sufficiently large n .

A *uniform black-box reduction showing the soundness of (G, P, V) based on a falsifiable assumption* Chal is a family of (uniform) probabilistic oracle algorithms $\{\text{Red}_c\}$ (one for each $c \in \mathbb{N}$) such that for every (G, P, V) -adversary (A, c) , $\text{Red}_c^A(1^n)$ breaks the assumption and runs in time $\text{poly}_c(n)$ (i.e. a polynomial that depends on c).

For a probabilistic oracle algorithm Red, we say a query $(1^m, \text{crs})$ of $\text{Red}(1^n)$ has *length* m . In general, $\text{Red}(1^n)$ may make queries of various lengths. We say Red is *length-mapping* if for all n , all queries of $\text{Red}(1^n)$ are of the same length $m = m(n)$ and m is computable in time $\text{poly}(n)$; denote this m by $\text{query}_{\text{Red}}(n)$. Most reductions in cryptography set $m = n$ i.e. preserve length; that is, the security parameter of (G, P, V) is equal to that of the assumption.

Following [GW], our results assume the existence of *hard subset membership problem*.

Definition 8.4 (Uniformly hard subset membership problem). Let n be a security parameter, L be an **NP** language associated with relation R . We say $((X, W), U)$ is a *subset membership problem for R* if $(X, W) = (X, W)(n)$ is a $\text{poly}(n)$ -time samplable joint distribution whose support lies in R , and $U = U(n)$ a $\text{poly}(n)$ -time samplable distribution with $\Pr[U \notin L] \geq n^{-O(1)}$.

A subset membership problem $((X, W), U)$ is a *subexponentially hard* if X and U are $(2^{\Omega(n^\delta)}, 2^{-\Omega(n^\delta)})$ -indistinguishable for a constant $\delta > 0$. We say it is *exponentially hard* if the above occurs and $|x| + |w| = O(n^\delta)$ for every $(x, w) \in \text{supp}(X, W)$.

This is a relatively mild assumption; the existence of subexponentially hard subset membership problems is equivalent to the existence of subexponentially hard one-way functions.

Remark. Our definition of a hard subset membership problem is a variant of [GW] that is needed in the uniform setting, but also can be used in the nonuniform setting of [GW]. In [GW], they require that X is indistinguishable from a (not necessarily samplable) distribution U whose support is disjoint from L , whereas we require that U is samplable and allow it to hit L with probability up to $1 - n^{O(1)}$.

We now state the uniform analogue of the main result of [GW]. Compared to [GW], our Theorem 8.5 makes the weaker assumption of subexponentially hard subset membership problem with respect to *uniform* algorithms, with the conclusion that a *uniform* falsifiable assumption cannot be broken also being weaker (unless the assumption is false).

Theorem 8.5 (Main theorem). *Let L be an NP language associated with relation R that has a subexponentially hard subset membership problem, and (G, P, V) be a non-interactive proof system for R that satisfies the completeness and succinctness properties. Then for every uniform falsifiable assumption Chal , one of the following must hold:*

- *The assumption Chal is false, or*
- *There is no uniform black-box reduction showing the soundness of (G, P, V) based on Chal .*

The same conclusion also holds if we assume an exponentially hard subset membership problem, and (G, P, V) is only slightly succinct.

The same conclusion also holds if we require the uniform black-box reduction to work only for all (G, P, V) -adversary (A, c) where c is sufficiently large.

To prove it in the nonuniform setting, the main idea of [GW] is showing that any SNARG (G, P, V) has an inefficient adversary A that can be (efficiently) “simulated” i.e. there exists an efficient algorithm Sim (the simulator) such that $\text{Red}^A(1^n) \approx \text{Red}^{\text{Sim}}(1^n)$ for all PPT oracle algorithms Red (cf. [GW] Lemma 4.1). Thus, if there were a black-box reduction Red showing the soundness of (G, P, V) based on a falsifiable assumption, then Red^A would break the falsifiable assumption (since A is an adversary) and so would Red^{Sim} (since $\text{Red}^A(1^n) \approx \text{Red}^{\text{Sim}}(1^n)$). In other words, the assumption would be false.

To prove it in the uniform setting, we use a similar approach with several necessary tweaks. We show that there is an *adversary simulator* Sim , which is a PPT algorithm that with noticeable probability outputs a randomized circuit B_n that simulates some A_n , where A_n is an (inefficient) adversary on security parameter n :

Lemma 8.6 (Existence of adversary simulator). *Let L be an NP language associated with relation R that has a subexponentially hard subset membership problem $((X, W), U)$, and (G, P, V) be a non-interactive proof system for R that satisfies the completeness and succinctness properties. Let n be a security parameter, $((X, W), U) = ((X, W), U)(n)$, $(\text{PRIV}, \text{CRS}) = G(1^n)$, and $\Pi = P(\text{CRS}, X, W)$. Let $\ell = \ell(n) \geq n$ be a polynomial bound on the running time of $G(1^n)$ as well as the proof size $|\Pi|$, and c be a constant such that $|X| + |\Pi| \leq n^c$.*

Let Red be any length-mapping PPT oracle algorithm where $\text{query}_{\text{Red}}(k) = \omega(1)$. Then there is a PPT algorithm Sim such that for all polynomials $q(\cdot)$, for all sufficiently large k , and for $n = \text{query}_{\text{Red}}(k)$, w.p. at least $1/\text{poly}(k)$, $\text{Sim}(1^k)$ outputs a randomized circuit B_n such that there is a randomized function A_n satisfying:

- *(A_n, c) is a (G, P, V) -adversary on the security parameter n ;*

- $\text{Break}_{\text{Red}^{A_n}}(k) - \text{Break}_{\text{Red}^{B_n}}(k) < 1/q(k)$. (WLOG B_n only takes inputs $(1^n, \cdot)$.)

The same conclusion also holds if we assume an exponentially hard subset membership problem, and that (G, P, V) is only slightly succinct.

Note that Lemma 8.6 is only stated for length-mapping reductions (unlike [GW]). We remove this restriction by a general technique when we prove the main theorem in Section 8.2.

8.1 Proof of Existence of Adversary Simulator (Lemma 8.6)

The proof is set up as follows. Given a subexponentially hard subset membership problem $((X, W), U)$, we can WLOG assume that X and U are $(2^{d\ell}, 2^{-d\ell})$ -indistinguishable for a sufficiently large constant d , where $\ell = \ell(n)$ is a bound on the length of the proof output by $P(\text{crs}, x, w)$ for $(x, w) \in \text{supp}(X, W)$ and $\text{crs} \in \text{supp}(\text{CRS})$. (If X and U are only $(2^{n^\delta}, 2^{-n^\delta})$ -indistinguishable for some $\delta > 0$, we simply re-index, replacing $X(n)$ with $X((d\ell)^{1/\delta})$.) If $((X, W), U)$ is exponentially hard, we can also ensure that X and U are $(2^{d\ell}, 2^{-d\ell})$ -indistinguishable by re-indexing so that $\ell \leq \text{poly}(n) \cdot (|x| + |w|)^\alpha + o(|x| + |w|) = O(|x| + |w|)/d$ for all $(x, w) \in \text{supp}(X, W)$ and $\text{crs} \in \text{supp}(\text{CRS})$.

Overview of the Proof. Consider the joint distribution (CRS, X, Π) where $\text{CRS} = \text{CRS}(n)$ is the distribution of the common reference string, and $\Pi = \Pi(n)$ is the ℓ -bit proof produced by P for the instance/witness pair (X, W) . Using the fact that Π is short (by succinctness), and X and U are ϵ -indistinguishable for $\epsilon = 2^{-O(\ell)}$, we can apply Theorem 6.12 to conclude that, for every $2^{O(\ell)}$ -time oracle algorithm D , there is a $\text{poly}(2^\ell, 1/\epsilon)$ -time randomized algorithm R that outputs a randomized circuit F_n such that with probability at least $\Omega(\epsilon^2/\ell)$ over F_n ,

$$\mathbb{E}[D^{O_{Q, F_n(Q)}}(\text{CRS}, X, \Pi)] - \mathbb{E}[D^{O_{Q, F_n(Q)}}(\text{CRS}, U, F_n(\text{CRS}, U))] < 2\epsilon \quad (\star)$$

where Q can be any poly-time samplable distribution.

An adversary A_n can be defined to be $A_n(1^n, \text{crs}) = (U, F_n(\text{crs}, U))$ for any F_n where (\star) holds, for an appropriate choice of D . (Note that F_n depends on our choice of D .) If we take D to be the verifier V , then we can show that such A_n breaks soundness on security parameter n . Indeed, V accepts (X, Π) with high probability, so by (\star) it must also accept $(U, F_n(\text{CRS}, U)) = A_n(1^n, \text{CRS})$ with high probability. (Some extra work is needed to deal with the fact that V can access its private coins PRIV in addition to CRS .)

Thus we only need to argue that, for an appropriate choice of D , such A_n is simulated by some randomized circuit B_n generated by a PPT algorithm Sim ; then combining the two choices of D will yield the desired adversary A_n . Our choice of B_n is the randomized circuit such that $B_n(1^n, \text{CRS}) = (X, \Pi)$. If we appropriately construct D from the reduction Red and challenger Chal , then using (\star) we can show that

$$\text{Break}_{\text{Red}^{A_n}}(k) - \text{Break}_{\text{Red}^{B_n}}(k) \leq \text{poly}(k) \cdot 2^{-O(\ell)},$$

where $\ell = \ell(n)$ for $n = \text{query}_{\text{Red}}(k)$. (If $\text{Break}_{\text{Red}^{A_n}}(k) - \text{Break}_{\text{Red}^{B_n}}(k) > \text{poly}(k) \cdot 2^{-O(\ell)}$, then we could use Red and Chal to construct a $2^{-O(\ell)}$ -distinguisher between $(\text{CRS}, B_n(1^n, \text{CRS})) = (\text{CRS}, X, \Pi)$ and $(\text{CRS}, A_n(1^n, \text{CRS})) = (\text{CRS}, U, F_n(\text{CRS}, U))$, violating (\star) .)

This completes the proof provided that $2^{-O(\ell)} \leq 1/\text{poly}(k)$, which follows if Red does not make queries that are too short. If instead $2^{-O(\ell)} > 1/\text{poly}(k)$, then we construct a simulator

B_n differently — simply by letting B_n be such that $B_n(1^n, \text{crs}) = (U, F_n(\text{crs}, U))$ where F_n is the random output of R . Then with probability at least $\Omega(\epsilon^2/\ell) \geq 1/\text{poly}(k)$ over F_n , (\star) holds for F_n , hence we can define the adversary A_n from F_n (defined to be $A_n(1^n, \text{crs}) = (U, F_n(\text{crs}, U))$, as explained above) to obtain a perfect simulator $B_n = A_n$. (Gentry and Wichs [GW] handle short queries using nonuniformity — by hardcoding the answers to all short queries.)

Lemma 8.7 (Existence of adversary simulator). *Let L be an **NP** language associated with relation R that has a subset membership problem $((X, W), U)$, and (G, P, V) is a non-interactive proof system for R that satisfies the completeness property. Let n be a security parameter, $((X, W), U) = ((X, W), U)(n)$, $(\text{PRIV}, \text{CRS}) = G(1^n)$, $\Pi = P(\text{CRS}, X, W)$. Let $\ell = \ell(n) \geq n$ be a polynomial bound on the running time of $G(1^n)$ as well as the proof size $|\Pi|$, and c be a constant such that $|X| + |\Pi| \leq n^c$.*

Suppose X and U are ϵ -indistinguishable for all t -time randomized algorithms, for appropriate $\epsilon = 2^{-O(\ell)}$ and $t = 2^{O(\ell)}$. Let Red be any length-mapping PPT oracle algorithm where $\text{query}_{\text{Red}}(k) = \omega(1)$. Then there is a PPT algorithm Sim such that for all polynomials $q(\cdot)$, for all sufficiently large k , and for $n = \text{query}_{\text{Red}}(k)$, w.p. at least $1/\text{poly}(k)$, $\text{Sim}(1^k)$ outputs a randomized circuit B_n such that there is a randomized function A_n satisfying:

- (A_n, c) is a (G, P, V) -adversary on the security parameter n ;
- $\text{Break}_{\text{Red}^{A_n}}(k) - \text{Break}_{\text{Red}^{B_n}}(k) < 1/q(k)$. (WLOG B_n only takes inputs $(1^n, \cdot)$.)

Proof. Let S be the PPT algorithm that on input $(1^n, \text{crs})$ samples $(x, w) \leftarrow (X, W)$ and outputs $(x, P(\text{crs}, x, w))$, so that $S(1^n, \text{CRS}) = (X, \Pi)$. For technical convenience we assume $|\text{CRS}| = \ell/2$. To construct Sim , we shall apply Theorem 6.12 to the following oracle algorithm D :

Claim 8.8. Let $Q = (U_{\ell/2}, U)$ (where $U_{\ell/2}$ is uniform on $\{0, 1\}^{\ell/2}$ and independent from U). There is a $t' = 2^{O(\ell)} \cdot \text{poly}(1/\epsilon)$ -time oracle algorithm D such that the following holds for all polynomials $q(\cdot)$, all sufficiently large n , and all randomized functions $F_n : \text{supp}(Q) \rightarrow \{0, 1\}^\ell$ satisfying

$$\mathbb{E}[D^{O_{Q, F_n(Q)}}(\text{CRS}, X, \Pi)] - \mathbb{E}[D^{O_{Q, F_n(Q)}}(\text{CRS}, U, F_n(\text{CRS}, U))] < \epsilon' = 2\epsilon.$$

Define

$$A_n(1^n, \text{crs}) = \begin{cases} (U, F_n(\text{crs}, U)), & \text{crs} \in \text{supp}(\text{CRS}) \\ S(1^n, \text{crs}), & \text{crs} \notin \text{supp}(\text{CRS}) \end{cases}.$$

Then

- A_n break soundness of (G, P, V) on security parameter n ; and
- For all $k \leq 2^\ell$ such that $\text{query}_{\text{Red}}(k) = n$,

$$\text{Break}_{\text{Red}^{A_n}}(k) - \text{Break}_{\text{Red}^S}(k) < 1/q(k).$$

Proof of Claim. We will prove the contrapositive. Suppose that either

Case 1. A_n does not break soundness of (G, P, V) on security parameter n , or

Case 2. For some $k \leq 2^\ell$ such that $\text{query}_{\text{Red}}(k) = n$,

$$\text{Break}_{\text{Red}^{A_n}}(k) - \text{Break}_{\text{Red}^S}(k) \geq 1/q(k).$$

We show how to construct a t' -time oracle algorithm D with

$$\mathbb{E}[D^{O_{Q, F_n(Q)}}(\text{CRS}, X, \Pi)] - \mathbb{E}[D^{O_{Q, F_n(Q)}}(\text{CRS}, U, F_n(\text{CRS}, U))] \geq \epsilon'.$$

To do so, we will show how to construct D with distinguishing advantage at least $3\epsilon'$, both in Case 1 and in Case 2 *where we assume k is known*. This suffices, because then we can test the distinguisher in Case 1 as well as the distinguisher in Case 2 *for all* choices of $k = 1, \dots, 2^\ell$, and output the best performing one. (More specifically, we run these $1 + 2^\ell$ distinguishers on $O((1/\epsilon'^2) \log(1/\epsilon'))$ independent samples of (CRS, X, Π) and $(\text{CRS}, U, F_n(\text{CRS}, U))$ as well as their coin tosses and oracle answers, and output the one with the highest average distinguishing advantage, and it follows from a Chernoff bound that this yields an ϵ' -distinguisher.)

Case 1. A_n does not break soundness on security parameter n . Recall that soundness says $\Pr[V(\text{PRIV}, U, \Pi') = 1 \wedge U \notin L] \leq n^{-c}$ if we choose $(\text{CRS}, \text{PRIV}) \leftarrow G(1^n)$, $(U, \Pi') \leftarrow A_n(1^n, \text{CRS})$ (thus $\Pi' = F_n(\text{CRS}, U)$). By union bound,

$$\Pr[V(\text{PRIV}, U, \Pi') = 1] \leq [V(\text{PRIV}, U, \Pi') = 1 \wedge U \notin L] + \Pr[U \in L] = 1 - n^{-O(1)}.$$

On the other hand, the completeness property says

$$\Pr[V(\text{PRIV}, X, \Pi) = 1] = 1 - \text{negl}(n).$$

Thus V is an $n^{-O(1)}$ -distinguisher between (PRIV, X, Π) and (PRIV, U, Π') . Note that conditioned on $\text{CRS} = \text{crs}$ for any crs , PRIV is independent of (X, Π) , and that $\text{PRIV}|_{\text{CRS}=\text{crs}}$ can be sampled in $2^{O(\ell)}$ time given crs (by running $G(1^n; z)$ on all sequences $z \in \{0, 1\}^\ell$ of coin tosses). Thus from V we also get a $2^{O(\ell)}$ time $n^{-O(1)}$ -distinguisher D for (CRS, X, Π) and (CRS, U, Π') . Specifically, $D(\text{crs}, x, \pi)$ samples $\text{priv} \leftarrow \text{PRIV}|_{\text{CRS}=\text{crs}}$ and outputs $V(\text{priv}, x, \pi)$, so

$$\begin{aligned} \mathbb{E}[D(\text{CRS}, X, \Pi)] - \mathbb{E}[D(\text{CRS}, U, F_n(\text{CRS}, U))] &= \mathbb{E}[D(\text{CRS}, X, \Pi)] - \mathbb{E}[D(\text{CRS}, U, \Pi')] \\ &= \Pr[V(\text{PRIV}, X, \Pi) = 1] - \Pr[V(\text{PRIV}, U, \Pi') = 1] \\ &= n^{-O(1)} \geq 3\epsilon'. \end{aligned}$$

Case 2. For some $k \leq 2^\ell$ such that $\text{query}_{\text{Red}}(k) = n$, we have

$$\text{Break}_{\text{Red}^{A_n}}(k) - \text{Break}_{\text{Red}^S}(k) \geq 1/q(k).$$

Assuming k is given, we use the hybrid argument to construct a distinguisher D between $(\text{CRS}, X, \Pi) = (\text{CRS}, S(1^n, \text{CRS}))$ and $(\text{CRS}, U, F_n(\text{CRS}, U)) = (\text{CRS}, A_n(1^n, \text{CRS}))$. Suppose $\text{Red}(1^k)$ runs in time $p(k)$ for some polynomial p . Let H_i be the stateful oracle that behaves like A_n for the first i queries and S for all rest of the queries, so that $H_q = A_n$ and $H_0 = S$. By the hybrid argument, $\mathbb{E}[\text{Red}^{H_{I+1}}(1^k)] - \mathbb{E}[\text{Red}^{H_I}(1^k)] \geq 1/p(k)q(k)$ for a randomly chosen $I \in_R \{1, \dots, p(k)\}$. This immediately gives us a distinguisher D' for $(Z, \text{CRS}', S(1^n, \text{CRS}'))$ and $(Z, \text{CRS}', A_n(1^n, \text{CRS}'))$ where Z is the internal state of the interaction $(\text{Red}^{H_I}(1^k), \text{Chal}(1^k))$ after $I \in_R \{1, \dots, p(k)\}$ queries, and CRS' is the I -th query (which is determined by Z). Specifically: $D'(z, \text{crs}, x, \pi)$ sets

the internal state of $\text{Red}(1^k)$ and $\text{Chal}(1^k)$ to z , runs the interaction $(\text{Red}^S(1^k), \text{Chal}(1^k))$ starting from state z using (x, π) as the answer to the I -th query $(1^n, \text{crs})$, and finally outputs 0 or 1 depending on whether Chal outputs win. Thus

$$\begin{aligned} & \mathbb{E} [D'(Z, \text{CRS}', S(1^n, \text{CRS}'))] - \mathbb{E} [D'(Z, \text{CRS}', A_n(1^n, \text{CRS}'))] \\ & \geq \mathbb{E} [\text{Red}^{H_{I-1}}(1^k)] - \mathbb{E} [\text{Red}^{H_I}(1^k)] \\ & \geq \frac{1}{p(k)q(k)} = 2^{-O(\ell)}. \end{aligned}$$

To obtain a desired distinguisher D'' for $(\text{CRS}, A_n(1^n, \text{CRS}))$ and $(\text{CRS}, S(1^n, \text{CRS}))$, we simply let D'' sample $(z, \text{crs}') \leftarrow (Z, \text{CRS}')$ and output

$$D''(\text{crs}, x, \pi) = \begin{cases} \frac{D'(z, \text{crs}, x, \pi)}{2^\ell \cdot \Pr[\text{CRS} = \text{crs}]}, & (\text{crs} = \text{crs}') \\ 0, & (\text{crs} \neq \text{crs}') \end{cases}.$$

Note that D'' is $[0, 1]$ -bounded since CRS is sampled by G using ℓ coin tosses (so $\Pr[\text{CRS} = \text{crs}] \geq 2^{-\ell}$ for all $\text{crs} \in \text{supp}(\text{CRS})$). We are dividing by $\Pr[\text{CRS} = \text{crs}]$ in order to “uniformize” CRS , so that

$$\begin{aligned} \mathbb{E}[D''(\text{CRS}, A(1^n, \text{CRS}))] &= \sum_{\text{crs} \in \text{supp}(\text{CRS})} \Pr[\text{CRS} = \text{crs}] \cdot \mathbb{E} \left[\frac{D'(Z, \text{CRS}', A_n(1^n, \text{CRS}'))}{2^\ell \cdot \Pr[\text{CRS} = \text{crs}]} \cdot I(\text{CRS}' = \text{crs}) \right] \\ &= \mathbb{E} [D'(Z, \text{CRS}', A_n(1^n, \text{CRS}')) \cdot I(\text{CRS}' \in \text{supp}(\text{CRS}))] \cdot 2^{-\ell} \end{aligned}$$

(where $I(\cdot)$ is the indicator function), and similarly for S . Thus D'' has distinguishing advantage

$$\begin{aligned} & \mathbb{E}[D''(\text{CRS}, S(1^n, \text{CRS}))] - \mathbb{E}[D''(\text{CRS}, A_n(1^n, \text{CRS}))] \\ &= \mathbb{E} [D'(Z, \text{CRS}', S(1^n, \text{CRS}')) \cdot I(\text{CRS}' \in \text{supp}(\text{CRS}))] \cdot 2^{-\ell} \\ & \quad - \mathbb{E} [D'(Z, \text{CRS}', A_n(1^n, \text{CRS}')) \cdot I(\text{CRS}' \in \text{supp}(\text{CRS}))] \cdot 2^{-\ell} \\ &= (\mathbb{E}[D'(Z, \text{CRS}', S(1^n, \text{CRS}'))] - \mathbb{E}[D'(Z, \text{CRS}', A_n(1^n, \text{CRS}'))]) \cdot 2^{-\ell} \\ &\geq 2^{-O(\ell)} \cdot 2^{-\ell} = 4\epsilon', \end{aligned}$$

where the second equality holds because $S(1^n, \text{CRS}')$ and $A_n(1^n, \text{CRS}')$ are identical whenever $\text{CRS}' \notin \text{supp}(\text{CRS})$.

To conclude Case 2, it remains to show that D'' can be implemented in time $2^{O(\ell)} \cdot \text{poly}(1/\epsilon)$. First, $\Pr[\text{CRS} = \text{crs}]$ can be computed in time $2^{O(\ell)}$ by enumerating coin tosses of $G(1^n)$. A query $(1^n, \text{crs})$ to S can be answered in $\text{poly}(n)$ time. A query $(1^n, \text{crs})$ to A_n with $\text{crs} \in \text{CRS}$ can be answered by sampling $(Q, F_n(Q)) = (U_{\ell/2}, U, F_n(U_{\ell/2}, U))$ for up to $O(2^\ell \cdot \log(1/(\epsilon \cdot p(k))))$ times until $U_{\ell/2} = \text{crs}$ (recall that we assume $|\text{crs}| = \ell/2$ in the setup of Lemma 8.6). Thus we can sample $(z, \text{crs}) \leftarrow (Z, \text{CRS}')$ and run $D'(z, \text{crs}, x, \pi)$ in $p(k) \cdot \max(\text{poly}(p(k)), 2^{O(\ell)}) = 2^{O(\ell)}$ time. It follows from a union bound that

$$\mathbb{E}[D^{O_{Q, F_n(Q)}}(\text{CRS}, X, \Pi)] - \mathbb{E}[D^{O_{Q, F_n(Q)}}(\text{CRS}, U, F_n(\text{CRS}, U))] \geq 3\epsilon'.$$

□

Given Claim 8.8, we now apply Theorem 6.12 to the oracle algorithm D we constructed in Claim 8.8. Since X and U are ϵ -indistinguishable, Theorem 6.12 yields a $t'' = \text{poly}(2^\ell, t', 1/\epsilon)$ -time randomized algorithm $R(1^n)$ that w.p. at least $\Omega(\epsilon^2/\ell)$ outputs a randomized circuit F_n satisfying

$$\mathbb{E}[D^{O_{Q, F_n(Q)}}(\text{CRS}, X, \Pi)] - \mathbb{E}[D^{O_{Q, F_n(Q)}}(\text{CRS}, U, F_n(\text{CRS}, U))] < 2\epsilon. \quad (3)$$

We define the simulator $\text{Sim}(1^k)$ to be the following algorithm:

1. Let $n = \text{query}_{\text{Red}}(k)$;
2. If $\ell(n) \geq \log k$, then output a circuit B_n where $B_n(1^n, \text{crs})$ runs $S(1^n, \text{crs})$;
3. Else, $\ell(n) < \log k$. We run $R(1^n)$ to obtain a randomized circuit F'_n , and output the randomized circuit B_n where

$$B_n(1^n, \text{crs}) = \begin{cases} (U, F'_n(\text{crs}, U)), & \text{crs} \in \text{supp}(\text{CRS}) \\ S(1^n, \text{crs}), & \text{crs} \notin \text{supp}(\text{CRS}) \end{cases}.$$

Note that Sim is a PPT algorithm since it runs in time $2^{O(\ell(n))} = \text{poly}(k)$ if $\ell(n) < \log k$, and in time $\text{poly}(k)$ if $\ell(n) \geq \log k$. To prove that Sim is indeed an adversary simulator, we define the adversary A_n (which depends on the coins of Sim) to be

$$A_n(1^n, \text{crs}) = \begin{cases} (U, F_n^*(\text{crs}, U)), & \text{crs} \in \text{supp}(\text{CRS}) \\ S(1^n, \text{crs}), & \text{crs} \notin \text{supp}(\text{CRS}) \end{cases}$$

where F_n^* is defined as follows:

- If $\ell(n) \geq \log k$, we let F_n^* be any randomized circuit such that Eq. 3 holds for $F_n = F_n^*$;
- If $\ell(n) < \log k$, we let F_n^* be F'_n generated by R in Step 3 of Sim , so that Eq. 3 holds for $F_n = F_n^*$ w.p. at least $\Omega(\epsilon(n)^2/\ell(n)) = 2^{-O(\ell(n))} \geq 1/\text{poly}(k)$ over the coins of R (hence coins of Sim).

We now apply Claim 8.8 to $F_n = F_n^*$. Note that Claim 8.8 holds “for all sufficiently large n ”, but since $\text{query}_{\text{Red}}(k) = \omega(1)$ it must also hold for all sufficiently large k and $n = \text{query}_{\text{Red}}(k)$. Thus Claim 8.8 implies that for all polynomials $q(\cdot)$, for all sufficiently large k and $n = \text{query}_{\text{Red}}(k)$, w.p. at least $1/\text{poly}(k)$ over B_n , A_n satisfies

1. (A_n, c) is a (G, P, V) -adversary on the security parameter n ; and
2. If $\ell(n) \geq \log k$, then $\text{Break}_{\text{Red}^{A_n}}(k) - \text{Break}_{\text{Red}^S}(k) < 1/q(k)$.

To conclude the proof it remains to show that

$$\text{Break}_{\text{Red}^{A_n}}(k) - \text{Break}_{\text{Red}^{B_n}}(k) < 1/q(k). \quad (4)$$

Indeed, if $\ell(n) \geq \log k$, then B_n runs S , so Eq. 4 follows from Item 2 above. If $\ell(n) < \log k$, then Eq. 4 holds because $A_n = B_n$ (since $F_n^* = F'_n$) thus $\text{Break}_{\text{Red}^{A_n}}(k) = \text{Break}_{\text{Red}^{B_n}}(k)$. \square

8.2 Proof of Main Theorem (Theorem 8.5)

The next two lemmas show that we can “convert” a generic black-box reduction into a length-mapping reduction, which in addition does not make very short queries. To do so, we first convert a generic black-box reduction into one that does not make very short queries (Lemma 8.9), by guessing “optimal” oracle answers for these very short queries. We then convert it to a length-mapping reduction (Lemma 8.10) by a “sparsification” trick, due to Chung, Mahmoody, and Pass [CMP]. As a consequence of “sparsification” the resulting length-mapping reduction no longer works with an arbitrary SNARG adversary. However, it still suffices for proving the main theorem using Lemma 8.6.

Lemma 8.9. *Let $c \in \mathbb{N}$ be a constant. Suppose there is a PPT oracle algorithm Red with the property that for every randomized function A where (A, c) is a (G, P, V) -adversary, Red^A breaks the falsifiable assumption. Then there is another PPT oracle algorithm $\widehat{\text{Red}}$ satisfying the same property, and in addition every query of $\widehat{\text{Red}}$ is of length at least $s = s(n) = (\log \log n)^{\Omega(1)}$.*

Proof. Suppose $G(1^m)$ outputs a crs of length m^d and let $s = s(n) = (\log \log n)^{1/(d+1)}$. We define $\widehat{\text{Red}}(1^n)$ as follows:

1. For each $m < s$, select a random function $B_m : \{0, 1\}^{m^d} \rightarrow \{0, 1\}^{m^c}$;
2. Run Red , using $B_m(\text{crs})$ to answer every query $(1^m, \text{crs})$ where $m < s$.

To see that $\widehat{\text{Red}}$ satisfies the same property as Red , consider any (G, P, V) -adversary (A, c) . By averaging the coins of A , for each $m < s$ we can fix some (deterministic) function $A_m : \{0, 1\}^{m^d} \rightarrow \{0, 1\}^{m^c}$ and define

$$\widehat{A}(1^m, \text{crs}) = \begin{cases} A_m(\text{crs}), & m < s \\ A(1^m, \text{crs}) & m \geq s \end{cases},$$

such that $\text{Red}^{\widehat{A}}$ breaks the falsifiable assumption. Note that $\{B_m : m < s\}$ can be encoded as an $s \cdot (s^c)^{s^d} = O(\log n)$ bit string. Thus w.p. at least $1/\text{poly}(n)$, $\widehat{\text{Red}}^A(1^n)$ sets $B_m = A_m$ for all $m < s$ and behaves identically to $\text{Red}^{\widehat{A}}$. Therefore $\widehat{\text{Red}}^A$ also breaks the falsifiable assumption. \square

Lemma 8.10 (Chung, Mahmoody, and Pass [CMP]). *Let $c \in \mathbb{N}$ be a constant. Suppose there is a PPT oracle algorithm Red with the property that for every randomized function A where (A, c) is a (G, P, V) -adversary, Red^A breaks the falsifiable assumption, and every query of Red is of length at least $s(n) = (\log \log n)^{\Omega(1)}$. Then there is a length-mapping PPT oracle algorithm $\widehat{\text{Red}}$ where $\text{query}_{\widehat{\text{Red}}}(n) \geq s(n)$, such that for infinitely many n and $m = \text{query}_{\widehat{\text{Red}}}(n)$, for every randomized function A_m where (A_m, c) is a (G, P, V) -adversary on security parameter m (of SNARG), $\widehat{\text{Red}}^{A_m}(1^n)$ breaks the assumption on security parameter n (of the assumption).*

Proof. We construct $\widehat{\text{Red}}$ from Red as follows. Fix a sparse sequence h_1, h_2, \dots where $h_1 = 1$ and $h_{m+1} = 2^{2^{h_m}}$ for $m \geq 1$. Note that the interval $[s(n), \text{poly}(n)]$ contains at most one element of the sequence h_1, h_2, \dots , for some n_c and all $n \geq n_c$. Let $\widehat{\text{Red}}^A(1^n)$ run $\text{Red}^A(1^n)$, where a query $(1^m, \text{crs})$ is answered as follows:

1. If $n < n_c$ or $m \notin \{h_1, h_2, \dots\}$, then answer the query with a special symbol \perp ;

2. Otherwise, answer the query using oracle A .

$\widehat{\text{Red}}$ is length-mapping, because every query of $\text{Red}^A(1^n)$ has length in the interval $[s(n), \text{poly}(n)]$ (since Red runs in time $\text{poly}(n)$), and for all $n \geq n_c$, at most one of h_1, h_2, \dots lies in that interval.

Suppose for contradiction that the $\widehat{\text{Red}}$ we construct does not satisfy the desired properties. That is, for all sufficiently large n and $m = \text{query}_{\widehat{\text{Red}}}(n)$, there exists some randomized function A_m where (A_m, c) is a (G, P, V) -adversary on security parameter m , but $\widehat{\text{Red}}^{A_m}(1^n)$ does not break the assumption on security parameter n .

Let A be any randomized function such that $A(1^m, \text{crs}) = A_m(1^m, \text{crs})$ where $m = \text{query}_{\widehat{\text{Red}}}(n)$ and for all sufficiently large n . Thus (A, c) is an a.e. (G, P, V) -adversary. Let \widehat{A} be a “sparsification” of A : $\widehat{A}(1^m, \text{crs}) := A(1^m, \text{crs})$ whenever $m \in \{h_1, h_2, \dots\}$ and $\widehat{A}(1^m, \text{crs}) := \perp$ for all other m . Thus (\widehat{A}, c) is a (G, P, V) -adversary.

Since (\widehat{A}, c) is a (G, P, V) -adversary $\text{Red}^{\widehat{A}}$ breaks the falsifiable assumption. On the other hand, $\widehat{\text{Red}}^A(1^n)$ behaves like $\widehat{\text{Red}}^{A_m}$ for all sufficiently large n , hence does not break the assumption. This yields a contradiction, because by construction $\widehat{\text{Red}}^A(1^n) = \text{Red}^{\widehat{A}}(1^n)$ for all $n \geq n_c$. \square

Finally, we use Lemma 8.6 to deduce Theorem 8.5. Note that the length-mapping reduction we obtain from Lemma 8.10 is slightly weaker, as it requires that the adversary break soundness on a *fixed* infinite sequence of security parameters (rather than *any* infinite sequence of security parameters). However, it suffices because Lemma 8.6 provides adversaries that break soundness on *almost all* security parameters.

Proof of Theorem 8.5 (Main Theorem). Suppose there is a generic uniform black-box reduction showing the soundness of (G, P, V) based on a uniform falsifiable assumption. We will show that the falsifiable assumption is already false, by constructing a PPT algorithm that breaks it.

Fix c to be the constant given by Lemma 8.6. By Lemma 8.9 and Lemma 8.10, there is a length-mapping PPT oracle algorithm Red where $\text{query}_{\text{Red}}(k) = \omega(1)$, and for infinitely many k , for $n = \text{query}_{\text{Red}}(k)$, and for every randomized function A_n where (A_n, c) is a (G, P, V) -adversary on security parameter n , $\text{Red}^{A_n}(1^k)$ breaks the assumption on security parameter k .

We now apply Lemma 8.6 to Red to obtain a PPT algorithm Sim such that for all polynomials $q(\cdot)$, all sufficiently large k and $n = \text{query}_{\text{Red}}(k)$, w.p. at least $1/\text{poly}(k)$, $\text{Sim}(1^k)$ outputs a randomized circuit B_n such that there is a randomized function A_n satisfying:

1. (A_n, c) is a (G, P, V) -adversary on the security parameter n ;
2. $\text{Break}_{\text{Red}^{A_n}}(k) - \text{Break}_{\text{Red}^{B_n}}(k) < 1/q(k)$.

By the previous discussion, for infinitely many k , Item 1 implies that $\text{Red}^{A_n}(1^k)$ breaks the assumption on security parameter k . Thus by Item 2, for infinitely many k , $\text{Red}^{B_n}(1^k)$ also breaks the assumption on security parameter k . Hence we obtain a PPT algorithm breaking the assumption for infinitely many k : first generate the circuit B_n by running $\text{Sim}(1^k)$, then run $\text{Red}^{B_n}(1^k)$. \square

Acknowledgments

We thank Kai-Min Chung for many helpful discussions (especially on the SNARGs application), Jacob Fox for pointing out our weak regularity lemma is optimal, and anonymous reviewers for their comments.

References

- [BHK] Boaz Barak, Moritz Hardt, and Satyen Kale. The uniform hardcore lemma via approximate bregman projections. In *SODA '09: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1193–1200, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.
- [BSW] Boaz Barak, Ronen Shaltiel, and Avi Wigderson. Computational analogues of entropy. In *RANDOM-APPROX*, pages 200–215, 2003.
- [CF] David Conlon and Jacob Fox. Bounds for graph regularity and removal lemmas. *Geom. Funct. Anal.*, 22(5):1191–1256, 2012.
- [CLP] Kai-Min Chung, Edward Lui, and Rafael Pass. From weak to strong zero-knowledge and applications. Cryptology ePrint Archive, Report 2013/260, 2013. <http://eprint.iacr.org/>.
- [CMP] Kai-Min Chung, Mohammad Mahmoody, and Rafael Pass. Personal communication, 2012/12.
- [CT] Thomas M. Cover and Joy A. Thomas. *Elements of information theory (2. ed.)*. Wiley, 2006.
- [DP1] Yevgeniy Dodis and Krzysztof Pietrzak. Leakage-resilient pseudorandom functions and side-channel attacks on Feistel networks. In *Advances in cryptology—CRYPTO 2010*, volume 6223 of *Lecture Notes in Comput. Sci.*, pages 21–40. Springer, Berlin, 2010.
- [DP2] Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *FOCS*, pages 293–302. IEEE Computer Society, 2008.
- [FK] Alan Frieze and Ravi Kannan. Quick approximation to matrices and applications. *Combinatorica*, 19(2):175–220, 1999.
- [FOR] Benjamin Fuller, Adam O’Neill, and Leonid Reyzin. A unified approach to deterministic encryption: New constructions and a connection to computational entropy. In Ronald Cramer, editor, *TCC*, volume 7194 of *Lecture Notes in Computer Science*, pages 582–599. Springer, 2012.
- [FR] Benjamin Fuller and Leonid Reyzin. Computational entropy and information leakage. 2011. (available at <http://www.cs.bu.edu/fac/reyzin>).
- [FS] Yoav Freund and Robert E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29:79–103, 1999.

- [Gow] W. T. Gowers. Decompositions, approximate structure, transference, and the Hahn-Banach theorem. *Bull. Lond. Math. Soc.*, 42(4):573–606, 2010.
- [GT] Ben Green and Terence Tao. The primes contain arbitrarily long arithmetic progressions. *Ann. of Math. (2)*, 167(2):481–547, 2008.
- [GW] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *STOC*, pages 99–108. ACM, 2011.
- [HH] Iftach Haitner and Thomas Holenstein. On the (im)possibility of key dependent encryption. In *TCC*, pages 202–219, 2009.
- [HHR] Iftach Haitner, Danny Harnik, and Omer Reingold. Efficient pseudorandom generators from exponentially hard one-way functions. In *Automata, Languages and Programming, 24th International Colloquium, ICALP*, 2006.
- [Hol1] Thomas Holenstein. Key agreement from weak bit agreement. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 664–673, 2005.
- [Hol2] Thomas Holenstein. Pseudorandom generators from one-way functions: A simple construction for any hardness. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006*, 2006.
- [HRV] Iftach Haitner, Omer Reingold, and Salil Vadhan. Efficiency improvements in constructing pseudorandom generators from one-way functions. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 437–446, 2010.
- [HW] M. Herbster and M. Warmuth. Tracking the best linear predictor. *Journal of Machine Learning Research*, 1:281–309, 2001.
- [Imp] Russell Impagliazzo. Hard-core distributions for somewhat hard problems. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 538–545, 1995.
- [KS] Adam R. Klivans and Rocco A. Servedio. Boosting and hard-core set construction. *Machine Learning*, 51(3):217–238, 2003.
- [LTW] Chi-Jen Lu, Shi-Chun Tsai, and Hsin-Lung Wu. Complexity of hard-core set proofs. *Computational Complexity*, 20(1):145–171, 2011.
- [MPRV] Ilya Mironov, Omkant Pandey, Omer Reingold, and Salil Vadhan. Computational differential privacy. In *Advances in cryptology—CRYPTO 2009*, volume 5677 of *Lecture Notes in Comput. Sci.*, pages 126–142. Springer, Berlin, 2009.
- [PJ] Krzysztof Pietrzak and Dimitar Jetchev. How to fake auxiliary input. *ICITS 2012 Invited Talk*.
- [RTTV] Omer Reingold, Luca Trevisan, Madhur Tulsiani, and Salil Vadhan. Dense subsets of pseudorandom sets. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '08)*, pages 76–85. IEEE, 26–28 October 2008.

- [TTV] Luca Trevisan, Madhur Tulsiani, and Salil Vadhan. Regularity, boosting, and efficiently simulating every high-entropy distribution. In *Proceedings of the 24th Annual IEEE Conference on Computational Complexity (CCC '09)*, pages 126–136, 15–18 July 2009. Preliminary version posted as *ECCC* TR08-103.
- [TZ] Terence Tao and Tamar Ziegler. The primes contain arbitrarily long polynomial progressions. *Acta Math.*, 201(2):213–305, 2008.
- [VZ] Salil Vadhan and Colin Jia Zheng. Characterizing pseudoentropy and simplifying pseudorandom generator constructions. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC '12)*, pages 817–836, 19–22 May 2012.
- [Zha] Jiapeng Zhang. On the query complexity for showing dense model. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:141, 2011.

A Information-Theoretic Facts, Efficient Approximations, Etc

Lemma A.1 (Multiplicative weight update decreases KL). *Let A, B be distributions over $[N]$ and $f : [N] \rightarrow [0, 1]$ any function. Define random variable A' such that*

$$\Pr[A' = x] \propto e^{\epsilon f(x)} \Pr[A = x]$$

for $0 \leq \epsilon \leq 1$. Then $\text{KL}(B \parallel A') \leq \text{KL}(B \parallel A) - (\log e)\epsilon(\mathbb{E}[f(B)] - \mathbb{E}[f(A)] - \epsilon)$.

Proof. By definition,

$$\begin{aligned} \text{KL}(B \parallel A) - \text{KL}(B \parallel A') &= \sum_x \Pr[B = x] \left(\log \frac{\Pr[B = x]}{\Pr[A = x]} - \log \frac{\Pr[B = x]}{\Pr[A' = x]} \right) \\ &= \sum_x \Pr[B = x] \log \frac{\Pr[A' = x]}{\Pr[A = x]} \\ &= \sum_x \Pr[B = x] \left(\log \frac{e^{\epsilon f(x)}}{\sum_y e^{\epsilon f(y)} \Pr[A = y]} \right) \\ &= (\log e) \left(\epsilon \mathbb{E}[f(B)] - \ln \left(\sum_y e^{\epsilon f(y)} \Pr[A = y] \right) \right) \end{aligned}$$

Applying the inequalities $1 + z \leq e^z$, $e^z \leq 1 + z + z^2$ for $0 \leq z \leq 1$, and using $0 \leq f(x) \leq 1$, we have

$$\begin{aligned} \text{KL}(B \parallel A) - \text{KL}(B \parallel A') &\geq (\log e) \left(\epsilon \mathbb{E}[f(B)] - \ln \left(\sum_y (1 + \epsilon f(y) + \epsilon^2) \Pr[A = y] \right) \right) \\ &= (\log e) \left(\epsilon \mathbb{E}[f(B)] - \ln (1 + \epsilon \mathbb{E}[f(A)] + \epsilon^2) \right) \\ &\geq (\log e) \left(\epsilon \mathbb{E}[f(B)] - (\epsilon \mathbb{E}[f(A)] + \epsilon^2) \right) \\ &= (\log e)\epsilon (\mathbb{E}[f(B)] - \mathbb{E}[f(A)] - \epsilon) \end{aligned}$$

□

Lemma A.2 (Sampling from a high density measure). *Let n be a security parameter, $\delta = \delta(n)$, $\sigma = \sigma(n)$. Then for $k = O((1/\delta) \log(1/\sigma))$, there is a randomized algorithm that, given k and oracle access to a measure $M \in \mathcal{M}_{n,\delta}$, w.p. at least $1 - \sigma$ outputs a random sample of Φ_M . The algorithm runs in $O(k(s+n))$ time and makes k oracle queries, where s is a bound on the bit length of $M(x)$.*

Proof. Use rejection sampling. Select a random $z \in_R \{0, 1\}^n$ and output z w.p. $M(z)$. Repeat up to $k = O((1/\delta) \log(1/\sigma))$ times until some z is outputted. Thus with all but $(1 - \mathbb{E}_z [M(z)])^k = (1 - \delta)^k \leq \sigma$ probability we output some $z \leftarrow \Phi_M$. \square

Lemma A.3 (Approximating KL projection on high min-entropy distributions). *Let \mathcal{C} be the set of distributions over $\{0, 1\}^n$ with min-entropy at least $n - \log(1/\delta)$. Then there is a probabilistic algorithm which, given any n , $\delta > 0$, $\epsilon > 0$, $\eta > 0$, achieves the following in $\text{poly}(n, 1/\delta, 1/\epsilon, \log(1/\eta))$ time. Given oracle access to a measure N with $\Phi_N \in \mathcal{C}^\epsilon$ (where \mathcal{C}^ϵ denotes the ϵ -neighborhood of \mathcal{C} ; see Definition 6.5), the algorithm w.p. at least $1 - \eta$ computes a measure M where Φ_M is an ϵ^2 -approximate KL projection of Φ_N on \mathcal{C} .*

Specifically, $M(x) = \min(1, c \cdot N(x))$ for some constant $c \in [1, 1 + e^\epsilon]$ as a multiple of $\Omega(\epsilon^2)$.

This follows immediately from Lemma 2.3 of Barak et al. [BHK], where they show how to approximate the KL projection on the set of high density measures (rather than high min-entropy distributions), which is equivalent to KL projection on high density distributions.

Proof. For measures M and N , we define the *KL divergence from M to N* to be

$$\text{KL}(M||N) = \sum_x \left(M(x) \log \frac{M(x)}{N(x)} - M(x) + N(x) \right).$$

Note that

$$\text{KL}(\Phi_M||\Phi_N) = \frac{\text{KL}(M||N)}{|M|} + 1 - \frac{|N|}{|M|} + \log \frac{|N|}{|M|}.$$

Barak et al. [BHK] show how to compute \widetilde{M}^* , a $\sigma \cdot (\delta 2^n)$ -approximate KL projection of N on the set of high density measures \mathcal{M}_δ . Let M^* be the (exact) KL projection of N on \mathcal{M}_δ , with $|M^*| = \delta 2^m$ (WLOG the KL projection is always on the boundary; see Lemma A.4). Thus by the above equality, Φ_{M^*} is the (exact) KL projection of N on \mathcal{C}_δ . Furthermore, for every $M \in \mathcal{M}_\delta$,

$$\begin{aligned} & \text{KL}(\Phi_M||\Phi_{\widetilde{M}^*}) - \text{KL}(\Phi_M||\Phi_{M^*}) \\ &= \frac{\text{KL}(M||\widetilde{M}^*) - \text{KL}(M||M^*)}{|M|} - \left(\frac{|\widetilde{M}^*|}{|M|} - \frac{|M^*|}{|M|} \right) + \left(\log \frac{|\widetilde{M}^*|}{|M|} - \log \frac{|M^*|}{|M|} \right) \\ &\leq \frac{\text{KL}(M||\widetilde{M}^*) - \text{KL}(M||M^*)}{|M|} \end{aligned}$$

where the inequality holds because $|\widetilde{M}^*| \geq |M^*|$. Thus $\Phi_{\widetilde{M}^*}$ is a σ -approximate KL projection of Φ_N on \mathcal{C}_δ . The parameters follow from Lemma 2.3 of [BHK]. \square

Lemma A.4. *The KL projection of any measure N on any convex set $\mathcal{M} \ni N$ is on the boundary of \mathcal{M} .*

Proof. Follows since the KL projection minimizes the convex function $\text{KL}(\cdot || N)$. \square