Complexity-Theoretic Implications of Multicalibration

Sílvia Casacuberta
University of Oxford
Oxford, UK
scasacubertapuig@college.harvard.edu

Cynthia Dwork
Harvard University
Cambridge, MA, USA
dwork@seas.harvard.edu

Salil Vadhan
Harvard University
Cambridge, MA, USA
salil_vadhan@harvard.edu

ABSTRACT

We present connections between the recent literature on multigroup fairness for prediction algorithms and classical results in computational complexity. *Multiaccurate* predictors are correct in expectation on each member of an arbitrary collection of pre-specified sets. *Multicalibrated* predictors satisfy a stronger condition: they are calibrated on each set in the collection.

Multiaccuracy is equivalent to a regularity notion for functions defined by Trevisan, Tulsiani, and Vadhan (2009). They showed that, given a class $\mathcal{F}$ of (possibly simple) functions, an arbitrarily complex function g can be approximated by a low-complexity function h that makes a small number of oracle calls to members of $\mathcal{F}$, where the notion of approximation requires that h cannot be distinguished from g by members of $\mathcal{F}$. This complexity-theoretic Regularity Lemma is known to have implications in different areas, including in complexity theory, additive number theory, information theory, graph theory, and cryptography. Starting from the stronger notion of *multicalibration*, we obtain stronger and more general versions of a number of applications of the Regularity Lemma, including the Hardcore Lemma, the Dense Model Theorem, and the equivalence of conditional pseudo-min-entropy and unpredictability. For example, we show that *every* boolean function (regardless of its hardness) has a small collection of disjoint hardcore sets, where the sizes of those hardcore sets are related to how balanced the function is on corresponding pieces of an efficient partition of the domain.

CCS CONCEPTS

• **Theory of computation → Complexity theory and logic**.

KEYWORDS

Regularity Lemma, Indistinguishability, Multicalibration, Hardcore Lemma, Dense Model Theorem, Pseudo-min-entropy

ACM Reference Format:
Sílvia Casacuberta, Cynthia Dwork, and Salil Vadhan. 2024. Complexity-Theoretic Implications of Multicalibration. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24), June 24–28, 2024, Vancouver, BC, Canada.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3618260.3649748

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

ACM ISBN 979-8-4007-0383-6/24/06
https://doi.org/10.1145/3618260.3649748

1 INTRODUCTION

In this paper, we give novel complexity-theoretic consequences of recent results in the algorithmic fairness literature regarding "multicalibration." Before stating our results we review the concept of multicalibration as well as the complexity-theoretic Regularity Lemma that provides the context for our theorems.

1.1 Multicalibration in Algorithmic Fairness

Algorithms increasingly inform decisions that can deeply affect our lives, from hiring, to healthcare diagnoses, to granting of release on bail. A major concern that arises in this context is whether or not prediction algorithms are *fair* across different subpopulations and minority groups [4, 42, 34]. Part of the effort within the field of algorithmic fairness has focused on developing mathematical frameworks to formally define what it means for an algorithm to be "fair". The various definitions that have been proposed in the literature roughly fall into two main categories: individual fairness notions [10, 25] and group fairness notions [7, 3].

The *multigroup* framework was proposed as a way to bridge individual fairness, which requires similar treatment for individuals who are similar with respect to a given task, and group fairness, which requires that (typically disjoint) demographic groups receive similar treatment on average [23, 31]. The underlying principle is the following: we want to establish a group fairness notion that is to be satisfied, simultaneously, for every one of a pre-specified collection of *large, identifiable, subgroups*. This versatile framework allows us to consider the intersection of different subgroups, such as the intersection of gender, race, and socioeconomic status. Multicalibration in particular has proven to be a fruitful notion with wide applications including a new paradigm for loss minimization in machine learning. See, *e.g.*, [2, 11, 18, 32, 14, 19, 35].

Hébert-Johnson, Kim, Reingold, and Rothblum [23] introduced the notion of a *multicalibrated (MC) predictor*, which guarantees calibrated predictions across every subpopulation from a prespecified family $\mathcal{F}$ of potentially intersecting subsets of $\mathcal{X}$. More formally, let $\mathcal{X}$ be a domain of individuals. Consider a collection $\mathcal{F}$ of subpopulations, each described as a boolean indicator function $f : \mathcal{X} \to \{0, 1\}$. Then, given an arbitrary and unknown function g mapping individuals in $\mathcal{X}$ to $[0, 1]$, and a distribution $\mathcal{D}$ on $\mathcal{X}$, we say that a predictor $h : \mathcal{X} \to [0, 1]$ is a $(\mathcal{F}, \epsilon)$-*multicalibrated (MC)* predictor of g with respect to $\mathcal{D}$ if for all $f \in \mathcal{F}$ and for all $v \in \text{image}(h)$:

$$\left| \mathop{\mathbb{E}}_{x \sim \mathcal{D}} [f(x) \cdot (g(x) - h(x)) \mid h(x) = v] \right| \leq \epsilon. \tag{1.1}$$

Remarkably, Hébert-Johnson et al. proved that "low-complexity" multicalibrated predictors h exist, provided we slightly relax the definition to only require (1.1) on level sets $h(x) = v$ that are not too small. We will state this result in more detail below in Section 1.3, using a more convenient formulation in terms of partitions of $\mathcal{X}$.

Multiaccuracy is a relaxation of multicalibration in which the predictor is merely required to be accurate in expectation on each $f \in \mathcal{F}$ [23]. Formally, h is an $(\mathcal{F}, \epsilon)$-*multiaccurate (MA)* predictor of g with respect to distribution $\mathcal{D}$ if for all $f \in \mathcal{F}$:

$$\left| \mathop{\mathbb{E}}_{x \sim \mathcal{D}} [f(x) \cdot (g(x) - h(x))] \right| \leq \epsilon. \tag{1.2}$$

1.2 The Complexity-Theoretic Regularity Lemma

It turns out that the notion of $(\mathcal{F}, \epsilon)$-multiaccuracy is exactly equivalent to the notion of $(\mathcal{F}, \epsilon)$-*indistinguishability* defined and studied in 2009 by Trevisan, Tulsiani, and Vadhan [39] in the complexity theory literature. Specifically, they proved the following abstract, complexity-theoretic Regularity Lemma.

Theorem 1.1 (Regularity Lemma [39], informally stated). *For every finite domain $\mathcal{X}$, every function $g : \mathcal{X} \to [0,1]$, every distribution $\mathcal{D}$ on $\mathcal{X}$, and every $\epsilon > 0$, there exists a function $h : \mathcal{X} \to [0,1]$ such that:*

(1) h has "low complexity" relative to $\mathcal{F}$. Specifically, h can be computed by a boolean circuit that has $O(1/\epsilon^2)$ oracle gates instantiated with functions from $\mathcal{F}$ and has size $\mathrm{poly}(\log|\mathcal{X}|, 1/\epsilon)$.
(2) h is $(\mathcal{F}, \epsilon)$-indistinguishable *from g. That is, for all $f \in \mathcal{F}$, we have:*

$$\left| \mathop{\mathbb{E}}_{x \sim \mathcal{D}} [f(x) \cdot (g(x) - h(x))] \right| \leq \epsilon. \tag{1.3}$$

Notice that Condition (1.3) is identical to the Definition (1.2) of multiaccuracy.

Regularity and Complexity. This is referred to as a *Regularity Lemma* because it says that an arbitrarily complex function g can be 'simulated' by a low-complexity function h, in such a way that the family $\mathcal{F}$ of tests cannot distinguish them. This of a similar spirit to Szemerédi's Regularity Lemma [37], whereby an arbitrarily complex graph is shown to be, in a certain sense, indistinguishable from the union of a constant number of Erdős-Rényi bipartite graphs [37]. Indeed, in [39], it is shown that Theorem 1.1 implies the Frieze–Kannan Weak Regularity Lemma for graphs [13], which is a lower-complexity variant of Szemerédi's Regularity Lemma. In a typical complexity-theoretic application, $\mathcal{F}$ consists of Boolean circuits of some polynomial size in the length $n = \log|\mathcal{X}|$ of inputs and $\epsilon = 1/\mathrm{poly}(n)$. In this case, the Regularity Lemma says that the simulator h can also be computed by circuits of size $\mathrm{poly}(n)$.

In the fairness literature, as well as in uniform-complexity applications in complexity and cryptography (cf. [40]), it is important to consider the complexity of *learning* such a predictor h given samples (x, y) where $x \sim \mathcal{D}$ and $y \sim \mathrm{Bern}(g(x))$. This may be computationally hard even if g is "easy"; for example, if g is in the class $\mathcal{F}$. In this paper, however, we are only concerned with the oracle complexity of h; i.e., the complexity of evaluating h given oracle gates for functions $f \in \mathcal{F}$ (which is trivial if $g \in \mathcal{F}$). We remark that many of the learning algorithms in the fairness literature assume an agnostic learner for $\mathcal{F}$ as an oracle, which also trivializes the learning task if $g \in \mathcal{F}$.

Applications. In addition to the Frieze–Kannan Weak Regularity Lemma for graphs, Theorem 1.1 can be used to derive several other fundamental theorems in various areas of theoretical computer science. These include Impagliazzo's Hardcore Lemma [27], the Dense Model Theorem [22, 38, 36], Yao's XOR Lemma [39, 17], the Leakage Simulation Lemma in leakage-resilient cryptography [30, 6], characterizations of pseudo-entropy [41, 44], chain rules for computational entropy [15, 30], Chang's Inequality in Fourier analysis of boolean functions [29], and equivalences between weak notions of zero knowledge [8].

Fractional vs. Boolean Functions. Note that we allow all of the functions f, g, h to be $[0,1]$-valued rather than just Boolean. We think of a $[0,1]$-valued function $h : \mathcal{X} \to [0,1]$ as representing the randomized Boolean function $h^{\mathrm{R}} : \mathcal{X} \to \{0,1\}$ where for all $x \in \mathcal{X}$ we have that $\Pr_{\mathrm{coins}(h)}[h^{\mathrm{R}}(x) = 1] = h(x)$. Then Equation (1.3) is essentially equivalent to requiring that the distributions of $(X, g^{\mathrm{R}}(X))$ and $(X, h^{\mathrm{R}}(X))$, where $X \sim \mathcal{D}$, are computationally indistinguishable by the family $\mathcal{F}$, up to an advantage of ϵ.[1] In the Regularity Lemma, it does not matter much if we restrict f, g, h to be deterministic Boolean functions, but for multicalibrated predictors it is crucial that h is fractional (else there could only be two nonempty level sets, $h(x) = 0$ and $h(x) = 1$).

1.3 Multicalibrated Partitions

It will be more convenient for us to work with an equivalent formulation of multicalibration in terms of *partitions* $\mathcal{P} \subseteq 2^{\mathcal{X}}$ of the domain $\mathcal{X}$, as was done in [18, 21]. The pieces $P \in \mathcal{P}$ of the partition correspond to the level sets $h(x) = v$ in (1.1); furthermore, it can be shown that we can assume without loss of generality that on piece P, we can take the value v to be equal to $v_P = \mathbb{E}_{x \sim \mathcal{D}|_P}[g(x)]$, where $\mathcal{D}|_P$ denotes the distribution $\mathcal{D}$ conditioned on being in P. In this language, the Multicalibration Theorem can be stated as follows:

Theorem 1.2 (Multicalibration Theorem [23], informally stated). *Let $\mathcal{X}$ be a finite domain, $\mathcal{F}$ a class of functions $f : \mathcal{X} \to [0,1]$, $g : \mathcal{X} \to [0,1]$ an arbitrary function, $\mathcal{D}$ a probability distribution over $\mathcal{X}$, and $\epsilon, \gamma > 0$. There exists a partition $\mathcal{P}$ of $\mathcal{X}$ such that:*

(1) $\mathcal{P}$ has $k = O(1/\epsilon)$ parts.
(2) $\mathcal{P}$ has "low complexity" relative to $\mathcal{F}$. Specifically, there is a Boolean circuit[2] $C : \mathcal{X} \to [k]$ of size $\mathrm{poly}(1/\epsilon, 1/\gamma, \log|\mathcal{X}|)$ *with $O(1/\epsilon^2)$ oracle gates instantiated with functions from $\mathcal{F}$ such that $\mathcal{P} = \{C^{-1}(1), \ldots, C^{-1}(k)\}$.*
(3) $\mathcal{P}$ is $(\mathcal{F}, \epsilon, \gamma)$-approximately multicalibrated *(MC) for g on $\mathcal{D}$: that is, for all $f \in \mathcal{F}$ and all $P \in \mathcal{P}$ such that* $\Pr_{x \sim \mathcal{D}}[x \in P] \geq \gamma$, *we have*

$$\left| \mathop{\mathbb{E}}_{x \sim \mathcal{D}|_P} [f(x) \cdot (g(x) - v_P)] \right| \leq \epsilon \tag{1.4}$$

where $v_P := \mathbb{E}_{x \sim \mathcal{D}|_P}[g(x)]$ and $\mathcal{D}|_P$ denotes the conditional distribution $\mathcal{D}|_{h(x) \in P}$.

Note that the pieces P of probability mass smaller than γ (for which (1.4) doesn't apply), take up at most a $\gamma \cdot k = O(\gamma/\epsilon)$ fraction of $\mathcal{D}$, which can be made arbitrarily small by taking $\gamma \ll \epsilon$.

[1] This equivalence holds up to a small modification to the family $\mathcal{F}$, in particular to account from changing the domain of the functions from $\mathcal{X}$ to $\mathcal{X} \times \{0,1\}$.
[2] That is, C is a circuit with Boolean gates of fan-in at most 2 that has $\lceil \log|\mathcal{X}| \rceil$ input gates and $\lceil \log k \rceil$ output gates.

On each piece P of probability mass at least γ, (1.4) says that g is indistinguishable from the constant function v_P., i.e. the function $h_P : P \to [0,1]$ where for all $x \in P$, $h_P(x) = v_P$. Viewing h_P as representing a randomized function h_P^R, we have $h_P^R(x) \sim \mathrm{Bern}(v_P)$ for all $x \in P$, where $\mathrm{Bern}(v)$ is the Bernoulli distribution with expectation v. The key point is that the Bernoulli parameter v is the same value (namely v_P) for all $x \in P$. Thus we informally refer to h_P and h_P^R as a *constant-Bernoulli function.* As we will see in the proofs of our results, indistinguishability from a constant-Bernoulli function is a very powerful condition, and this is why we are able to get so much mileage out of the MC Theorem.

1.4 Our Contributions

Although the Multicalibration Theorem (Theorem 1.2) was introduced for the purpose of algorithmic fairness, we now see that it can be viewed as a strengthening of the complexity-theoretic Regularity Lemma (Theorem 1.1). Thus, in our work, we examine the complexity-theoretic implications of the MC Theorem. Doing so, we obtain stronger and more general versions of (1) Impagliazzo's Hardcore Lemma (IHCL), (2) Characterizations of pseudo-average-min-entropy (PAME), and (3) the Dense Model Theorem (DMT). In concurrent work to ours, Dwork, Lee, Lin, and Tankala [12] explore an intermediate notion of graph regularity that corresponds to multicalibration and lies between Frieze-Kannan weak regularity and Szemerédi Regularity, and show that Szemerédi Regularity corresponds to a stronger notion called *strict* multicalibration.

We elaborate on our results below, denoting the strengthened theorems as IHCL++, PAME++, and DMT++. For simplicity, here we will state some of the results for the special case where the initial distribution $\mathcal{D}$ on inputs is the uniform distribution on $\mathcal{X}$ and the functions g and $f \in \mathcal{F}$ are deterministic boolean functions; generalizations to arbitrary distributions and fractional/randomized functions can be found in the later technical sections.

Impagliazzo's Hardcore Lemma (IHCL) and IHCL++. Impagliazzo's Hardcore Lemma (IHCL) [27] is a fundamental result in complexity theory stating that if a function g is somewhat hard to compute on average by a family $\mathcal{F}$ of boolean functions, then there is a large-enough subset H of the inputs (called the "hardcore set") for which the function is very hard to compute, in the sense that g is indistinguishable from a random function by a family $\mathcal{F}'$ of distinguishers of complexity similar to that of $\mathcal{F}$. A stronger, and optimal, version was obtained by Holenstein [24]:

Theorem 1.3 (IHCL [27, 24], informally stated). *Let $\mathcal{F}$ be a family of boolean functions on $\mathcal{X}$, $\epsilon > 0$, and $g : \mathcal{X} \to \{0,1\}$ a function that is $(\mathcal{F}', \delta)$-*hard*, meaning that $\Pr[f(x) \neq g(x)] \geq \delta$ for all functions f that have "low complexity" relative to $\mathcal{F}$.[3] Then there exists a set $H \subseteq \mathcal{X}$ of size at least $2\delta|\mathcal{X}|$ such that g is $(\mathcal{F}, 1/2 - \epsilon)$-hard on H.*

In [39], it was shown that the Regularity Lemma implies IHCL, but with a hardcore density of δ (as in [27]), instead of the optimal 2δ from [24]. (2δ is optimal because erring with probability $1/2$ on a set of density 2δ yields a global error probability of δ.)

Using the Multicalibration Theorem, we prove:

[3] Specifically, take $\mathcal{F}'$ to consist of all functions computable by a boolean circuit of size $\mathrm{poly}(1/\epsilon, 1/\delta, \log|\mathcal{X}|)$ with oracle gates instantiated by functions from $\mathcal{F}$.

Theorem 1.4 (IHCL++, informal version). *Let $g : \mathcal{X} \to \{0,1\}$ be an arbitrary function, $\mathcal{F}$ a family of boolean functions, and $\epsilon > 0$. There exists partition $\mathcal{P}$ of $\mathcal{X}$ that has "low complexity" relative to $\mathcal{F}$ such that for every (large enough) $P \in \mathcal{P}$, there is a set $H_P \subseteq P$ of size at least $2b_P|P|$ such that g is $(\mathcal{F}, 1/2 - \epsilon_P)$-hard on H_P, where $b_P = \min\{\mathbb{E}_{x\sim P}[g(x)], 1 - \mathbb{E}_{x\sim P}[g(x)]\}$ and $\epsilon_P = \epsilon/b_P$.*

That is, instead of finding a single, globally dense hardcore set H, we find many "local" hardcore sets H_P, each of which is dense within its piece P of the partition. We illustrate the key differences between IHCL and IHCL++ in Figure 1.

Here (and in our other results), the "low complexity" of the partition $\mathcal{P}$ is formulated in the same way as in the MC Theorem (Theorem 1.2). Naturally we prove Theorem 1.4 by applying the MC Theorem to the function g; see Section 1.6 for more on our proof techniques.

The *balance parameter* b_P provides the moral equivalent of the hardness parameter δ in IHCL (since we make no hardness assumptions). We further show that our IHCL++ implies the original IHCL theorem with optimal density parameter 2δ. To do so, we observe that when we bring back the assumption that g is $(\mathcal{F}', \delta)$-hard, then $\mathbb{E}_P[b_P] \geq \delta$, where the expectation is taken over sampling piece P with probability $|P|/|\mathcal{X}|$. That is, if g is δ-hard, then g is not too imbalanced on average over the pieces of the partition. Otherwise, we would be able to predict $g(x)$ well on average by determining which piece $P \in \mathcal{P}$ contains x (because $\mathcal{P}$ has low complexity relative to $\mathcal{F}$) and then guessing the majority value on the piece P (which we can hardwire into our Boolean circuit for each of the $k = O(1/\epsilon)$ pieces). We are then able to "glue" together the hardcore sets H_P for the pieces $P \in \mathcal{P}$, yielding a hardcore set $H \subseteq \mathcal{X}$ that occupies at least a 2δ fraction of the domain $\mathcal{X}$. (Actually we get a hardcore *distribution* of density at least 2δ, but this can be converted to a hardcore *set* by a standard probabilistic argument [27].)

A partition-based variant of the IHCL was previously formulated and proved in the work of Reingold, Trevisan, Tulsiani, and Vadhan [36]. Their result is similar in spirit to IHCL++, but with two important differences. First, their partition has complexity *exponential* in $1/\epsilon$, in contrast to the polynomial complexity we obtain through the MC Theorem. The exponential complexity severely limits the complexity-theoretic applicability of their result. Second, they maintain the original assumption of IHCL that the function g is weakly hard on average, whereas in our IHCL++ we remove it and find local hardcore sets H_P for an arbitrary function g. However, their proof can be modified to also remove the hardness assumption and yield a similar conclusion to ours. Indeed, their proof, which proceeds by iteratively partitioning the domain, can be viewed in retrospect as constructing an MC partition with exponential complexity. At the time, exponential complexity seemed inherent in such iterative partitioning proofs [36, 39]. In this light, the power of the MC Theorem and successors [23, 18, 21] is that they give us the same kind of indistinguishability as provided by iterative partitioning but with polynomial complexity. Intuitively, the savings in complexity comes from using merging steps in addition to partitioning ones to avoid making too many, too small pieces.

Characterizations of pseudo-average min-entropy (PAME). A result of Vadhan and Zheng shows that we can characterize pseudoentropy, which is a computational analogue of Shannon

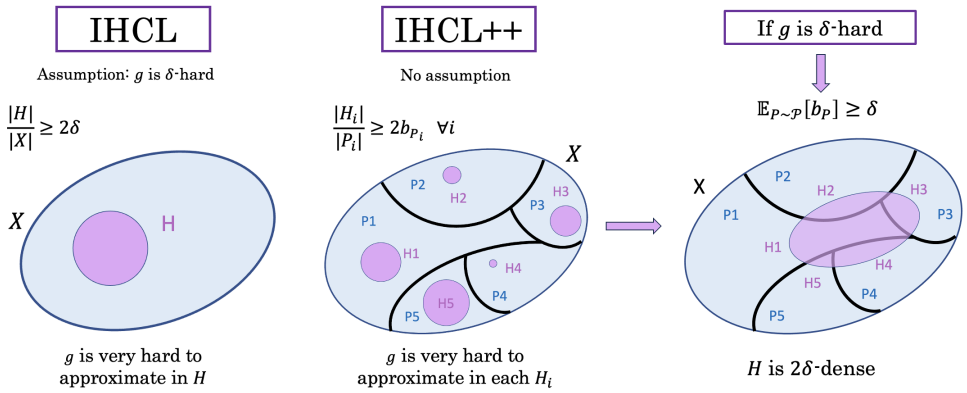

Figure 1: Illustration of the difference between IHCL and IHCL++, and how to recover IHCL from our ICHL++.

entropy, in terms of hardness of sampling [41].[4] In [44], they also provide a characterization for the related notion of *pseudo-average min-entopy (PAME)*, which is the computational analogue of average min-entropy [9].[5]

Informally, for a joint distribution (X, B) and a class of distinguishers $\mathcal{F}$, B has $(\mathcal{F}, \epsilon)$-*PAME at least k given X* if there exists a random variable C jointly distributed with X such that 1) the distributions (X, B) and (X, C) are $(\mathcal{F}, \epsilon)$-indistinguishable, and 2) $C|X$ has average min-entropy at least k. The PAME Theorem of Vadhan and Zheng shows that the PAME of B given X is precisely characterized by the hardness of predicting B given X:

Theorem 1.5 (PAME [41, 44], informally stated). *Let $\mathcal{F}$ be a family of boolean functions on $\mathcal{X} = \{0,1\}^n \times \{0,1\}^\ell$, $\epsilon > 0$, and let (X, B) be a joint distribution over $\{0,1\}^n \times \{0,1\}^\ell$, where $\ell = O(\log n)$. Suppose that B is $(\mathcal{F}', \delta)$-hard to predict from X, meaning that $\Pr[f(X) \neq B] \geq \delta$ for all functions $f : \{0,1\}^n \to \{0,1\}$ that have "low complexity" relative to $\mathcal{F}$. Then B has $(\mathcal{F}, \epsilon)$-PAME at least $\log(1/(1-\delta))$ given X.*

The reason for the setting $\ell = O(\log n)$ in Theorem 1.5 is that the "low complexity" parameter for $\mathcal{F}'$ includes a factor of 2^ℓ.

Like the MC Theorem and IHCL++, our PAME++ will involve partitions. For $P \subseteq \mathcal{X}$, let (X_P, B_P) denote (X, B) conditioned on $X \in P$, and denote the min-entropy of a random variable B by $\mathrm{H}_\infty(B) = \min_b \log(1/\Pr[B = b])$, where all logarithms in our paper are base 2 unless otherwise specified.

Theorem 1.6 (PAME++, informally stated). *Let (X, B) be a joint distribution over $\{0,1\}^n \times \{0,1\}^\ell$, where $B = g(X)$ and $\ell = O(\log\log n)$, $\epsilon > 0$, and let $\mathcal{F}$ be a family of boolean functions on $\{0,1\}^n \times \{0,1\}^\ell$. There exists a partition $\mathcal{P}$ of $\mathcal{X} = \{0,1\}^n$ that has "low complexity" relative to $\mathcal{F}$ such that B_P has $(\mathcal{F}, \epsilon)$-PAME at least $\mathrm{H}_\infty(B_P)$ given X_P for every (large enough) $P \in \mathcal{P}$.*

That is, on each (large-enough) piece P of the partition $\mathcal{P}$, the PAME of $B|_P$ given $X|_P$ is the same as the min-entropy of $B|_P$ without conditioning. That is, $(X|_P, B|_P)$ is indistinguishable from having as much average-min-entropy as it would if $X|_P$ and $B|_P$ were independent.

To prove this theorem, we apply a *Multiclass* Multicalibration Theorem from the literature [18, 20, 12] to the randomized function $g^{\mathrm{R}}(x) \sim B|_{X=x}$. These theorems incur a doubly-exponential dependence on ℓ in the complexity of the partition (recall that ℓ is the logarithm of the size of the output set), which is why our PAME++ Theorem restricts to $\ell = O(\log\log n)$ instead of the $\ell = O(\log n)$ in the original PAME Theorem. Note that $\mathrm{H}_\infty(B_P) = \log(1/m_P)$, where $m_P = \max_{y\in\{0,1\}^\ell} \Pr[B_P = y]$ provides a multiclass generalization of our previous notion of balance b_P. (Specifically, $m_P = 1 - b_P$ when $\ell = 1$.)

As in the case of IHCL++, the more balanced g is on a cell $P \in \mathcal{P}$, the better the corresponding PAME lower bound. Also as in the case of IHCL++, we are able to recover the original PAME statement (in the case that $\ell = O(\log\log n)$) from our PAME++ theorem. Specifically, when we bring back the assumption that g is δ-hard to predict, we have that $\mathbb{E}_P[m_P] \leq 1 - \delta$, where P is sampled according to its mass under X. Again, we are able to "glue" together the distributions C_P of high average min-entropy (which are indistinguishable from B_P) for all the pieces $P \in \mathcal{P}$ that have enough mass according to the distribution X and on which g is balanced enough. This yields a conditional distribution $C|X$ over $\{0,1\}^\ell$ that has PAME at least $\log(1/(1-\delta))$. We remark that Vadhan and Zheng [41] also formulate and prove an analogue of the PAME Theorem for distinguishers and predictors f that are given by uniform probabilistic algorithms. Our results only apply to the nonuniform case (e.g., Boolean circuits); a uniform treatment of multicalibration is an interesting problem for future work.

The Dense Model Theorem (DMT). This result originated from additive number theory [22, 38] and states that if R is a pseudorandom set, whereby we mean that the uniform distribution on R is indistinguishable from the uniform distribution on the entire domain $\mathcal{X}$ by some family of tests, and S is a dense subset of R, then there exists a truly dense set $M \subseteq \mathcal{X}$ (called the *model* for S) that is indistinguishable from S by a related family of tests. The DMT was a crucial proof component used in Green and Tao's celebrated result

[4] Vadhan and Zheng prove their results for both uniform and non-uniform families of distinguishers. We focus on the non-uniform case.

[5] For a joint distribution (X, C), the *average min-entropy* of C given X is defined as $\tilde{\mathrm{H}}_\infty(C|X) = \log\left(1/\mathbb{E}_{x\sim X}\left[2^{-\mathrm{H}_\infty(C|_{X=x})}\right]\right)$, where all logarithms in the paper are taken base 2 unless otherwise specified.

that there exist arbitrarily long arithmetic progressions among the prime numbers [22].

A more general, complexity-theoretic version of the Dense Model Theorem was given by Reingold, Trevisan, Tulsiani, and Vadhan [36]. Following Impagliazzo [28, 26], the idea of a dense subset S of pseudorandom set R can be generalized to that of a *pseudodense* set S, which is defined in the theorem statement below.

Theorem 1.7 (DMT [22, 38, 36, 28, 26], informally stated). *Let $\mathcal{F}$ be a family of boolean functions on domain $\mathcal{X}$. For every $\epsilon, \delta > 0$, let $S \subseteq \mathcal{X}$ be $(\mathcal{F}', \epsilon, \delta)$-pseudodense, meaning that $\Pr_{x \in \mathcal{X}}[f(x) = 1] \geq \delta \cdot \Pr_{x \in S}[f(x) = 1] - \epsilon$, for all functions f of "low complexity" relative to $\mathcal{F}$. Then there exists a set M of density $|M|/|\mathcal{X}| \geq \delta - O(\epsilon)$ such that the uniform distribution on M is $(\mathcal{F}, \epsilon/\delta)$-indistinguishable from the uniform distribution on S.*

Using the MC Theorem, we instead obtain the following:

Theorem 1.8 (DMT++, informally stated). *Let S, V be two disjoint sets, let $\mathcal{F}$ be a family of boolean functions on $S \cup V$, and let $\epsilon > 0$. There exists a "low-complexity" partition $\mathcal{P}$ of $S \cup V$ such that for every (large enough) P, the uniform distribution over $P \cap V$ is $(\mathcal{F}, \epsilon_P)$-indistinguishable from $P \cap S$, for an appropriate choice of ϵ_P.*

To see the connection between this statement and the original DMT, think of V as a disjoint copy of the entire domain $\mathcal{X}$ that S lives in. Then we think of $P \cap V$ as a model for $P \cap S$, one that has true density $|P \cap V|/|V|$. If we bring back the assumption that S is δ-pseudodense, then we can take an appropriate convex combination of the models $P \cap V$ to obtain a distribution that is truly δ-dense in $\mathcal{X}$ and is indistinguishable from S. Specifically, our model distribution will select P with probability $|P \cap S|/|S|$ and then return a uniformly random element of $P \cap V$. As usual, we will only use the pieces P that are both large enough and not too imbalanced; i.e., ones where neither $|P \cap S|/|S|$ nor $|P \cap V|/|V|$ is too small. We remark that, given a dense model distribution, it is possible to obtain a dense set as in Theorem 1.7 by a standard probabilistic argument, similarly to the IHCL [27].

To prove Theorem 1.8, we apply the Multicalibration Theorem to characteristic function of the set S on the distribution that selects a uniformly random element of S with probability 1/2 and selects a uniformly random element of V with probability 1/2.

1.5 Common Themes

There are several common themes that recur across all of our ++ theorems.

Reproducing the original theorems locally. Because multicalibration yields a low-complexity partition of the domain, all of our stronger and more general ++ theorems find a low-complexity partition of the domain such that the original theorem is reproduced "locally" in each of the pieces of the partition simultaneously. In the case of IHCL, we find a hardcore set within each piece of the partition. In the case of PAME, we construct a distribution within in each piece of the partition. For the DMT, we construct a model within each piece of the partition.

Our theorems remove the original assumptions. All of the IHCL, PAME, and DMT theorems assume some kind of computational hardness in their hypotheses: In IHCL, the function g is assumed to be weakly hard on average. In the case of PAME, B is assumed to be unpredictable given X. In the case of the DMT, the set S is assumed to be pseudodense in $\mathcal{X}$. In all of the three original theorems, given the hardness assumption, an object is then constructed achieving: 1) A stronger hardness condition, and 2) Some density guarantee.

In the case of IHCL (Theorem 1.3), the original theorem finds a subset H of the domain such that 1) the input function is maximally unpredictable inside H, and 2) H occupies at least a certain fraction of the domain. For PAME (Theorem 1.5), the original theorem builds a conditional distribution $C|X$ that is 1) indistinguishable from the input distribution, and 2) that has at least some amount of average min-entropy. For the DMT (Theorem 1.7), the original theorem constructs a set that is 1) indistinguishable from uniform distribution on the pseudodense set S and 2) satisfies a lower bound on its density.

When we use the Multicalibration theorem instead of the Regularity/Multiaccuracy theorem, we find that we can maintain 1) the strong hardness conclusion in each piece of the partition, but *without requiring the hardness assumption in the hypotheses*, and so our stronger theorems hold for an arbitrary function (in the case of IHCL++), an arbitrary conditional distribution (in the case of PAME++), and an arbitrary set (in the case of DMT++). That is, we show that *every* input object admits a certain regular partition/decomposition, where this "regularity" captures the required hardness/unpredictability/indistinguishability in the sense of the original theorems.

Hardness vs Density. In the original theorems, the density lower bounds depend on a parameter that is given by the hardness assumption. For example, in the case of IHCL and PAME, the input function/distribution is assumed to be δ-hard for some δ, and then the hardcore set given by IHCL is shown to have density at least 2δ and the indistinguishable distribution given by PAME is shown to have average min-entropy at least $\log(1/(1-\delta))$. In DMT, the input set S is assumed to be δ-pseudodense, and the conclusion gives a model of true density at least δ. Given that our ++ theorems remove the hardness assumption, we no longer have such a hardness parameter. Instead, the density lower bounds that we show in our stronger ++ theorems relate to how balanced our function/object is on each piece of the partition.

Recovering the original theorems. Moreover, our ++ theorems are also stronger than the original ones because we are able to derive the original theorems as a corollary of our new results. Indeed, we show that if we bring back the assumption from the original theorems, we can "stitch" together the objects that we have built within each piece P in our ++ theorems (i.e., hardcore sets in the case of IHCL++, high-entropy distributions in the case of PAME++, and dense models in the case of DMT++) so that the resulting object satisfies the conclusion of the original theorems. We are able to lower-bound the weighted average of the balance parameters in terms of the hardness parameter from our assumption, crucially using the fact that the partition is of low complexity.

These themes are summarized in Table 1.

1.6 Proof Techniques

As discussed in Section 1.3, the power of the Multicalibration Theorem (Theorem 1.2) is that it partitions the domain $\mathcal{X}$ into pieces

Table 1: Summary of how the results in this paper generalize the original IHCL, PAME, and DMT theorems. The notation $\mathcal{U}_S$ denotes the uniform distribution over the set S, and $\approx_{(\mathcal{F},\epsilon)}$ denotes (F,ϵ)-indistinguishability.

	IHCL	IHCL++	PAME	PAME++	DMT	DMT++
Input object	Function f		Joint distribution (X,B)		Sets S,V	
Assumption on the input	f is $(\mathcal{F}',\delta)$-hard	–	B is $(\mathcal{F}',\delta)$-hard given X	–	S is $(\mathcal{F}',\epsilon,\delta)$-pseudodense	–
Indist. guarantee	$\exists$ hardcore set H	$\exists$ low-complexity partition w/ $O(1/\epsilon)$ hardcore sets H_P	$\exists$ dist. C s.t. $(X,C)\approx_{(\mathcal{F},\epsilon)}(X,B)$	$\exists\, O(1/\epsilon)$ dists. C_P s.t. $(X_P,B_P)\approx_{(\mathcal{F},\epsilon)}(X_P,C_P)$	$\exists$ set M s.t. $\mathcal{U}_M\approx_{(\mathcal{F},\epsilon/\delta)}\mathcal{U}_S$	$\exists\, O(1/\epsilon)$ sets S_P,V_P s.t. $\mathcal{U}_{S_P}\approx_{(\mathcal{F},\epsilon_P)}\mathcal{U}_{V_P}$
Density guarantee	$\lvert H\rvert/\lvert\mathcal{X}\rvert\geq 2\delta$	$\lvert H_P\rvert/\lvert\mathcal{X}\rvert\geq 2b_P$	$\tilde{\mathrm{H}}_\infty(C\lvert X)\geq\log(1/(1-\delta))$	$\tilde{\mathrm{H}}_\infty(C_P\lvert X_P)\geq\mathrm{H}(B_P)$	$\lvert M\rvert/\lvert\mathcal{X}\rvert\geq\delta-O(\epsilon)$	$\lvert P\cap V\rvert/\lvert V\rvert$
Recovery	–	Recovers IHCL	–	Recovers PAME	–	Recovers DMT

P such that on each piece, g is indistinguishable from a constant-Bernoulli function. This is a powerful condition, as we see through the following lemma.

Lemma 1.9 (Characterizing Indistinguishability from Constant-Bernoulli Functions, informally stated). *Let $\mathcal{F}$ be a family of boolean functions on $\mathcal{X}$, and let $\epsilon>0$. Suppose that $g:\mathcal{X}\to\{0,1\}$ is $(\mathcal{F},\epsilon)$-indistinguishable from a constant-Bernoulli function with expectation $v=\mathbb{E}_{x\in\mathcal{X}}[g(x)]$ and balance $b=\min\{v,1-v\}$. Then, each of the following hold up to small changes in the family $\mathcal{F}$ and/or the parameter ϵ (denoted as $\mathcal{F}'$ and ϵ'):*

1. *There is a set H of density $2b$ in $\mathcal{X}$ such that g is $(\mathcal{F}',1/2-\epsilon'/b)$-hard on H.*
2. *The distribution $(X,g(X))$ is $(\mathcal{F}',\epsilon)$-indistinguishable from the distribution $(X,\mathrm{Bern}(v))$, where X is sampled uniformly from $\mathcal{X}$.*
3. *The uniform distributions on $g^{-1}(1)$ and $g^{-1}(0)$ are $(\mathcal{F}',\epsilon'/b)$-indistinguishable from each other.*

The three parts of the lemma are what we use in the proofs of the IHCL++, PAME++, and IHCL++, respectively. Indeed, H is our local hardcore set, $\mathrm{Bern}(v)$ is our local distribution of high average min-entropy, and $g^{-1}(0)$ is our local dense model of $g^{-1}(1)$.

2 NOTATION AND PRELIMINARIES

We denote the domain by $\mathcal{X}$, the class of distinguishers by $\mathcal{F}=\{f\}$, the function to which we apply the multicalibration theorem by g, and the MA/MC predictor by h. We use $\mathcal{P}=\{P\}\subseteq 2^{\mathcal{X}}$ to denote a partition of the domain. The notation $\Pr_{x\in\mathcal{X}}$ means that x is sampled uniformly from $\mathcal{X}$, whereas $x\sim\mathcal{D}$ denotes that x is sampled according to distribution $\mathcal{D}$. We denote the constant 0 and 1 functions by $\mathbf{0}$ and $\mathbf{1}$, respectively. We say that a class of functions $\mathcal{F}$ is *closed under negation* if for all $f\in\mathcal{F}$, the function $-f$ is also in $\mathcal{F}$. All the logarithms in this paper are assumed to be in base 2.

Following the intuition that we provided in the introduction, the formal definition of multiaccuracy is as follows:

Definition 2.1 (Multiaccuracy [23, 33])**.** Let $\mathcal{X}$ be a finite domain, $\mathcal{F}$ a collection of functions $f:\mathcal{X}\to[0,1]$, $g:\mathcal{X}\to[0,1]$ an arbitrary function, $\mathcal{D}$ a probability distribution over $\mathcal{X}$, and $\epsilon>0$. We say that $h:\mathcal{X}\to[0,1]$ is an $(\mathcal{F},\epsilon)$-*multiaccurate* (MA) predictor for g on $\mathcal{D}$ if, for all $f\in\mathcal{F}$,

$$\left|\mathop{\mathbb{E}}_{x\sim\mathcal{D}}[f(x)\cdot(g(x)-h(x))]\right|\leq\epsilon.$$

We think of a $[0,1]$-valued function f as describing a randomized $\{0,1\}$-valued function f^{R} where $\Pr_{\mathrm{coins}(f^{\mathrm{R}})}[f^{\mathrm{R}}(x)=1]=f(x)$. Then,

$$\mathop{\mathbb{E}}_{x\sim\mathcal{D}}[f(x)\cdot(g(x)-h(x))]=\mathop{\mathbb{E}}_{\substack{x\sim\mathcal{D},\,\mathrm{coins}(f^{\mathrm{R}}),\\ \mathrm{coins}(g^{\mathrm{R}}),\,\mathrm{coins}(h^{\mathrm{R}})}}\left[f^{\mathrm{R}}(x)\cdot(g^{\mathrm{R}}(x)-h^{\mathrm{R}}(x))\right].$$

Therefore, the notion of multiaccuracy applies to the randomized functions as well.

The starting point of this work is the observation that multiaccuracy corresponds *exactly* to the classical notion of indistinguishability with respect to a class of functions:

Definition 2.2 ($(\mathcal{F},\epsilon)$-indistinguishability [39])**.** Let $\mathcal{X}$ be a finite domain, $\mathcal{F}$ a class of functions $f:\mathcal{X}\to\{0,1\}$, $g:\mathcal{X}\to[0,1]$, $\mathcal{D}$ a distribution on $\mathcal{X}$ and $\epsilon>0$. We say that a function $h:\mathcal{X}\to[0,1]$ is $(\mathcal{F},\epsilon)$-*indistinguishable* from g on $\mathcal{D}$ if, for all $f\in\mathcal{F}$,

$$\left|\mathop{\mathbb{E}}_{x\sim\mathcal{D}}[f(x)\cdot(g(x)-h(x))]\right|\leq\epsilon.$$

Multicalibration is a stronger notion than multiaccuracy, where the predictor h satisfies that $\mathbb{E}_{x\sim\mathcal{D}|_{h(x)=v}}[f(x)\cdot(g(x)-v)]\leq\epsilon$ for every $v\in\mathrm{range}(h)$ and every $f\in\mathcal{F}$ [23], where $\mathcal{D}|_{h(x)=v}$ denotes the conditional distribution. Thus the level sets of h induce a partition $\mathcal{P}$ of the domain, and the value of h in each piece $P\in\mathcal{P}$ can be made to be equal to the expected value of g over P, which we denote by v_P:

Definition 2.3 (Balance of g)**.** Given an arbitrary function $g:\mathcal{X}\to[0,1]$ and a partition $\mathcal{P}=\{P\}$ of $\mathcal{X}$, we let $v_P=\mathbb{E}_{x\sim\mathcal{D}|_P}[g(x)]$ for each $P\in\mathcal{P}$ and $b_P=\min\{v_P,1-v_P\}\leq 1/2$, where $\mathcal{D}|_P$ denotes the conditional distribution $\mathcal{D}|_{h(x)\in P}$ We call b_P the *balance* of g on P.

In particular, $b_P = 1/2$ corresponds to g^R being perfectly balanced; i.e.,

$$\Pr_{\substack{x\sim\mathcal{D},\\ \text{coins}(g^R)}}[g^R(x)=1] = \Pr_{\substack{x\sim\mathcal{D},\\ \text{coins}(g^R)}}[g^R(x)=0] = 1/2,$$

whereas $b_P = 0$ corresponds to g being completely imbalanced; i.e., $g^R(x)$ is always 0 or 1.

Moreover, as explained in the introduction, we need to relax the notion of multicalibration to *approximate multicalibration* by introducing a lower bound γ on the size of each $P \in \mathcal{P}$ (according to distribution $\mathcal{D}$). This yields the definition of an *approximate MC partition*:

Definition 2.4 (Approximate MC partition)**.** Let $\mathcal{X}$ be a finite domain, $\mathcal{F}$ a class of functions $f\colon \mathcal{X} \to [0,1]$, $g\colon \mathcal{X} \to [0,1]$ an arbitrary function, $\mathcal{D}$ a probability distribution over $\mathcal{X}$, and $\epsilon, \gamma > 0$. We say that a partition $\mathcal{P}$ of $\mathcal{X}$ is $(\mathcal{F}, \epsilon, \gamma)$-*approximately multicalibrated* (MC) for g on $\mathcal{D}$ if for all $f \in \mathcal{F}$ and all $P \in \mathcal{P}$ such that $\Pr_{x\sim\mathcal{D}}[x \in P] \geq \gamma$,

$$\left| \mathop{\mathbb{E}}_{x\sim\mathcal{D}|_P}[f(x)\cdot(g(x)-v_P)] \right| \leq \epsilon$$

where $v_P := \mathbb{E}_{x\sim\mathcal{D}|_P} P[g(x)]$ and $\mathcal{D}|_P$ denotes the conditional distribution $\mathcal{D}|_{h(x)\in P}$.

Note that achieving Definition 2.4 is trivial if we allow $O(1/\gamma)$ pieces P. We will want to achieve a partition that is much smaller; the goal is to satisfy approximate multcalibration with only $O(1/\epsilon)$ pieces, where $\epsilon \gg \gamma$. This turns out to be possible, as demonstrated by Theorem 2.10.

In the case where $\mathcal{D}$ corresponds to the uniform distribution over $\mathcal{X}$, then $\Pr_{x\sim\mathcal{D}}[x \in P] = |P|/|\mathcal{X}|$. That is, in this case, Definition 2.4 should be understood as saying that we do not make any guarantees about sets that are too small (namely, about sets that occupy less than a γ fraction of the space). Additionally, in order to keep track of the impact of the size of each $P \in \mathcal{P}$, we introduce the following notation:

Definition 2.5. Given $P \subseteq \mathcal{X}$, we let $\eta_P = \Pr_{x\sim\mathcal{D}}[x \in P]$ denote the *size* parameter of P in $\mathcal{X}$. If $\mathcal{D}$ corresponds to the uniform distribution over $\mathcal{X}$, then $\eta_P := |P|/|\mathcal{X}|$.

Definition 2.6. Given a partition $\mathcal{P}$ of $\mathcal{X}$ and a distribution $\mathcal{D}$ over $\mathcal{X}$, $\mathcal{P}(\mathcal{D})$ denotes the distribution on $\mathcal{P}$ that selects each $P \in \mathcal{P}$ with probability $\sum_{x\in P}\mathcal{D}(x)$.

Next, we study the notion of complexity of a partition. We use the number of wires of a circuit as the circuit size measure.

Complexity of a partition. As we developed in the introduction, a key property of a multicalibrated partition is that it is a *low-complexity* partition of the domain $\mathcal{X}$. We now formalize this idea.

Definition 2.7 (Relative complexity of a function [30, Definition 6])**.** Let $\mathcal{F}$ be a family of functions $f\colon \mathcal{X} \to [0,1]$. A function h has *complexity* (t,q) *relative to* $\mathcal{F}$ if it can be computed by an oracle-aided circuit of size t with q oracle gates, where each oracle gate is instantiated with a function from $\mathcal{F}$.

The notion of relative complexity captures the idea that we can make oracle calls to the functions in $\mathcal{F}$ without these factoring into the complexity. The algorithms to construct MA and MC predictors h use an oracle for a weak agnostic learner for the family $\mathcal{F}$ [39, 23, 18, 16]. In such a case, the parameter q corresponds to the number of oracle calls to the weak agnostic learner q.

Definition 2.8. Given an arbitrary class of functions $\mathcal{F}$, we denote by $\mathcal{F}_{t,q}$ the class of functions that have complexity at most (t,q) relative to $\mathcal{F}$.

Definition 2.9. Given a set of functions $\mathcal{F} = \{f\}$ on a finite domain $\mathcal{X}$, $\mathcal{F}_{t,q,k}$ denotes the class of partitions $\mathcal{P}$ of $\mathcal{X}$ such that there exists $\hat{f} \in \mathcal{F}_{t,q}$, $\hat{f} : \mathcal{X} \to [k]$, satisfying $\mathcal{P} = \{\hat{f}^{-1}(1), \ldots, \hat{f}^{-1}(k)\}$.

The condition $P_i = \hat{f}^{-1}(i)$ stated in Definition 2.9 ensures that we can always know to which level set each $x \in \mathcal{X}$ belongs to by performing an oracle call to a function in $\mathcal{F}_{t,q}$. Intuitively, we are associating each $P \in \mathcal{P}$ with an integer in $[k]$, and then Definition 2.9 requires the existence of a function in $F_{t,q}$ that we use to query to which P each $x \in \mathcal{X}$ belongs to. We call this function $\hat{f}$ the *partition membership function*. Naturally, this $\hat{f}$ is constant on each $P \in \mathcal{P}$.

Having formalized the complexity class $\mathcal{F}_{t,q,k}$ of partitions, we can now state the theorem that is the backbone of all our results in this paper:

Theorem 2.10 (Multicalibration Theorem [23]). *Let $\mathcal{X}$ be a finite domain, $\mathcal{F}$ a class of functions $f\colon \mathcal{X} \to [0,1]$, $g\colon \mathcal{X} \to [0,1]$ an arbitrary function, $\mathcal{D}$ a probability distribution over $\mathcal{X}$, and $\epsilon, \gamma > 0$. There exists an $(\mathcal{F}, \epsilon, \gamma)$-approximately multicalibrated partition $\mathcal{P}$ of $\mathcal{X}$ for g on $\mathcal{D}$ such that $\mathcal{P} \in \mathcal{F}_{t,q,k}$, where*

1. $t = O(1/(\epsilon^4\gamma)\cdot\log(|\mathcal{X}|/\epsilon))$,
2. $q = O(1/\epsilon^2)$,
3. $k = O(1/\epsilon)$.

We defer the proof of Theorem 2.10 to the full version of the paper.

2.1 Hardness and Indistinguishability Notions

One of the contributions of our paper is to provide a complexity-theoretic perspective on the power of a multicalibrated partition. Indeed, as we developed in the introduction, given the definition of $(\mathcal{F}, \epsilon)$-indistinguishability, the MC partition theorem is equivalent to stating that, given an arbitrary function g, we can find a low-complexity partition of the domain such that g is indistinguishable from the constant function v_P on each piece P of the partition.

In order to illustrate why this is a powerful statement, in this section we relate indistinguishability from a constant function to Yao's lemma on the equivalence between pseudorandomness and unpredictability [43]. To do so, we formally define what it means for a function to be *hard* with respect to a class of functions $\mathcal{F}$:

Definition 2.11 (Hardness of a function)**.** Given a class $\mathcal{F}$ of randomized functions $f\colon \mathcal{X} \to \{0,1\}^\ell$, a distribution $\mathcal{D}$ on $\mathcal{X}$, an arbitrary randomized function $g\colon \mathcal{X} \to \{0,1\}^\ell$, and $\delta > 0$, we say that g is $(\mathcal{F}, \delta)$-*hard* on $\mathcal{D}$ if, for all $f \in \mathcal{F}$,

$$\Pr_{\substack{x\sim\mathcal{D},\,\text{coins}(f)\\ \text{coins}(g)}}[f(x) = g(x)] \leq 1-\delta.$$

Note that here we consider randomized functions with discrete range rather than deterministic functions with range $[0,1]$. When $\ell = 1$, the maximal possible hardness occurs when $\delta = 1/2 - \epsilon$, given that being $(\mathcal{F}, 1/2-\epsilon)$-hard corresponds to stating that

$$\Pr_{\substack{x\sim\mathcal{D},\,\mathrm{coins}(f)\\ \mathrm{coins}(g)}}[f(x) = g(x)] \le 1/2 + \epsilon.$$

That is, no distinguisher in $\mathcal{F}$ can guess g noticeably better than a random bit. This is why we sometimes refer to being $(\mathcal{F}, 1/2-\epsilon)$-hard as being ϵ-*strongly hard*, whereas being $(\mathcal{F}, \delta)$-hard is sometimes referred to as being δ-*weakly* hard. In the case of Impagliazzo's Hardcore Lemma, the task is precisely to find a subset H of the domain on which g is $(\mathcal{F}, 1/2-\epsilon)$-hard, and hence maximally unpredictable.

Yao first showed a relationship between pseudorandomness (i.e., indistinguishability from a constant $1/2$ function h, for which $\Pr_{\mathrm{coins}(h)}[h^{\mathrm{R}}(x) = 1] = \Pr_{\mathrm{coins}(h)}[h^{\mathrm{R}}(x) = 0] = 1/2$) and unpredictability:

Lemma 2.12 (Equivalence between indistinguishability and pseudorandomness [43]). *Given a class of functions $\mathcal{F}$, a distribution $\mathcal{D}$ on $\mathcal{X}$ and $\epsilon > 0$, a function $g : \mathcal{X} \to [0,1]$ such that $\mathbb{E}_{x\sim\mathcal{D}}[g(x) = 1/2]$ is $(\mathcal{F}, \epsilon)$-indistinguishable on $\mathcal{D}$ from the constant $1/2$ function if and only if g^{R} is $(\mathcal{F}^{\mathrm{R}}, 1/2 - 2\epsilon)$-hard.*

In other words, stating that g is strongly hard corresponds exactly to stating that g is indistinguishable from a uniform random bit. Lemma 2.12 follows from the identity

$$\Pr_{\substack{x\sim\mathcal{D},\\ \mathrm{coins}(f^{\mathrm{R}}),\\ \mathrm{coins}(g^{\mathrm{R}})}}[f^{\mathrm{R}}(x) = g^{\mathrm{R}}(x)] = 2\,\mathop{\mathbb{E}}_{x\sim\mathcal{D}}[(f(x)-1/2)(g(x)-1/2)] + 1/2. \tag{2.1}$$

Next, we relate Yao's Lemma to the MC Theorem. For it, we need to introduce the definition of the density of a distribution and extend the notion of $(\mathcal{F}, \epsilon)$-indistinguishability to distributions:

Definition 2.13 (δ-dense distribution)**.** A distribution A is δ-*dense* in a distribution B if for all $x \in \mathcal{X}$,

$$\delta \cdot \Pr[A = x] \le \Pr[B = x].$$

If B is the uniform distribution on X, then this becomes $\Pr[A = x] \le 1/(\delta|\mathcal{X}|)$; i.e., a condition that is satisfied by the uniform distribution on any set of size at least $\delta|\mathcal{X}|$.

Definition 2.14 (Indistinguishable distributions)**.** Given a class $\mathcal{F}$ of functions $f : \mathcal{X} \to [0,1]$ and two distributions $\mathcal{D}_1, \mathcal{D}_2$ on $\mathcal{X}$, we say that $\mathcal{D}_1$ and $\mathcal{D}_2$ are $(\mathcal{F}, \epsilon)$-*indistinguishable* if, for all $f \in \mathcal{F}$,

$$\left| \mathop{\mathbb{E}}_{x\sim\mathcal{D}_1}[f(x)] - \mathop{\mathbb{E}}_{x\sim\mathcal{D}_2}[f(x)] \right| \le \epsilon.$$

2.2 Characterizations of Constant-Bernoulli Functions

While Yao's lemma characterizes indistinguishability from the constant $1/2$ function, in a multicalibrated partition (Definition 2.4) we have indistinguishability from the constant $v_P = \mathbb{E}_{x\sim\mathcal{D}|_P}[g(x)]$ function on each piece P, where v_P can take any value in $[0,1]$. Thus we seek to characterize functions $g : \mathcal{X} \to [0,1]$ that are indistinguishable from constant functions $h : \mathcal{X} \to [0,1]$ (i.e., there is a $v \in [0,1]$ such that $h(x) = v$ for all $x \in \mathcal{X}$). Note that such an h represents a randomized function h^{R} such that $h^{\mathrm{R}}(x)$ is identically distributed to $\mathrm{Bern}(v)$ for all $x \in \mathcal{X}$, which we refer to as a *constant-Bernoulli* function. If g is $(\mathcal{F}, \epsilon)$ indistinguishable from a constant-Bernoulli function $h(x) = v$ on distribution $\mathcal{D}$ and $\mathcal{F}$ contains the constant 1 function, then we can assume that $v = \mathbb{E}_{x\sim\mathcal{D}}[g(x)]$ with a factor 2 change in ϵ. Thus we state the following lemma with this assumption, which is anyhow guaranteed to us by the MC Theorem as formulated in Theorem 2.10.

Lemma 2.15 (Characterizing indistinguishability from constant-Bernoulli functions). *Let $\mathcal{F}$ be a class of functions $f : \mathcal{X} \to [0,1]$ closed under negation and such that $\mathbf{0}, \mathbf{1} \in \mathcal{F}$, let $\mathcal{D}$ be a distribution over the domain $\mathcal{X}$, and let $\epsilon > 0$. Let $g : \mathcal{X} \to [0,1]$ be $(\mathcal{F}, \epsilon)$-indistinguishable from the constant function v on $\mathcal{D}$, where $v = \mathbb{E}_{x\sim\mathcal{D}}[g(x)]$ and $b = \min\{v, 1-v\}$. Then, the following statements hold:*

(1) The distribution $(X, g^{\mathrm{R}}(X))$ is $(\mathcal{F}', \epsilon)$-indistinguishable from the distribution $(X, \mathrm{Bern}(v))$, where $X \sim \mathcal{D}$ and $\mathcal{F}'$ is any class such that $\mathcal{F}'_{O(n),O(1)} \subseteq \mathcal{F}$.

(2) The function g^{R} is $(\mathcal{F}^{\mathrm{R}}, b - 2\epsilon)$-hard on $\mathcal{D}$.

(3) The distribution $\mathcal{D}|_{g^{\mathrm{R}}(x)=1}$ is $\left(\mathcal{F}', \frac{\epsilon}{v(1-v)}\right)$-indistinguishable from $\mathcal{D}|_{g^{\mathrm{R}}(x)=0}$ for any class $\mathcal{F}'$ such that $\mathcal{F}'_{c\log|\mathcal{X}|,c} \subseteq \mathcal{F}$, where c is a universal constant.

(4) Assuming g is Boolean (i.e., $g(x) \in \{0,1\}$ for all $x \in \mathcal{X}$), there exists a distribution of density $2b$ in $\mathcal{D}$ such that g is $\left(\mathcal{F}^{\mathrm{R}}, 1/2 - \frac{\epsilon}{2v(1-v)}\right)$-hard on it.

Why Lemma 2.15 is central to our ++ theorems. As we mentioned in the introduction, we prove our ++ theorems as consequences of Lemma 2.15. First, Lemma 2.15 is applicable to our setting because of our complexity-theoretic recasting of the definition of a multicalibrated partition. Namely, a multicalibrated partition $\mathcal{P} = \{P\}$ for $g, \mathcal{F}, \mathcal{D}, \epsilon > 0$ is such that for each (large enough) $P \in \mathcal{P}$, the function g is $(\mathcal{F}, \epsilon)$-indistinguishable from the constant function $v = \mathbb{E}_{\mathcal{D}}[g(x)]$. Hence, each such piece $P \in \mathcal{P}$ satisfies the assumption of Lemma 2.15. We then use the statements of Lemma 2.15 as follows:

(1) The proof of IHCL++ follows from statement (2.15). Essentially, the distribution given by statement (2.15) corresponds to a "small" hardcore set contained within each piece of $\mathcal{P}$.

(2) The proof of PAME++ follows from statement (2.15). Roughly, when showing the existence of a distribution that has high average min-entropy and that is indistinguishable from $B|_P$ (as required by the definition of PAME), we use a Bernoulli distribution with parameter v_P.

(3) The proof of DMT++ follows from statement (2.15). While there is no function g in the statement of the DMT, we define g precisely as the characteristic function of the set S. Then, the fact that $\mathcal{D}|_{g^{-1}(0)}$ and $\mathcal{D}|_{g^{-1}(1)}$ are indistinguishable allows us to argue that the sets $S \cap P$ and $U \cap P$ are indistinguishable with respect to $\mathcal{F}$, which in turn shows that $U \cap P$ is a model for the corresponding set $S \cap P$.

Proof of Lemma 2.15. (1.) Given a class of functions $\mathcal{F}' = \{f' : \mathcal{X} \times \{0,1\} \to [0,1]\}$, we construct a class of functions $\mathcal{F}'' = \{f'' : \mathcal{X} \to [0,1]\}$ as follows. For each $f' \in \mathcal{F}'$, we add the

function $f''(x) = f'(x,1) - f'(x,0)$ and its negation to $\mathcal{F}''$. Note that $\mathcal{F}'' \subseteq \mathcal{F}'_{cn,c} \subseteq \mathcal{F}$. Then,

$$\mathop{\mathbb{E}}_{\substack{x\sim\mathcal{D},\\ b\sim\mathrm{Bern}(v),\\ \mathrm{coins}(g^{\mathrm{R}})}}[f'(x,g^{\mathrm{R}}(x)) - f'(x,b)] =$$

$$= \mathop{\mathbb{E}}_{x\sim\mathcal{D}}[g(x)f'(x,1) + (1-g(x))f'(x,0) - vf'(x,1) - (1-v)f'(x,0)]$$
$$= \mathop{\mathbb{E}}_{x\sim\mathcal{D}}[f''(x)(g(x)-v)].$$

(2.) Assume without loss of generality that $b = v \leq 1/2$; otherwise, we can replace f and g by $-f$ and $-g$ respectively in the argument below. Then,

$$\Pr_{\substack{x\sim\mathcal{D},\,\mathrm{coins}(f^{\mathrm{R}}),\\ \mathrm{coins}(g^{\mathrm{R}})}}[f^{\mathrm{R}}(x) = g^{\mathrm{R}}(x)] =$$
$$= \mathop{\mathbb{E}}_{x\sim\mathcal{D}}[f(x)g(x) + (1-f(x))(1-g(x))]$$
$$= \mathop{\mathbb{E}}_{x\sim\mathcal{D}}[(2f(x)-1)(g(x)-v) + 1 - v - (1-2v)f(x)] \leq 1 - v + 2\epsilon,$$

given that $v \leq 1/2$ implies $(1-2v)f(x) \geq 0$ and that, by assumption, $\left|\mathbb{E}_{\mathcal{D}}[f(x)\cdot(g(x)-v)]\right| \leq \epsilon$.

(3.) By the assumption on g, it follows that

$$\left|\mathop{\mathbb{E}}_{x\sim\mathcal{D}}[f(x)\cdot(g(x)-v)]\right| = \left|\mathop{\mathbb{E}}_{\substack{x\sim\mathcal{D},\\ \mathrm{coins}(g^{\mathrm{R}})}}[f(x)\cdot(g^{\mathrm{R}}(x)-v)]\right| =$$

$$= \left|v\cdot\mathop{\mathbb{E}}_{\substack{x\sim\mathcal{D}|_{g^{\mathrm{R}}(x)=1},\\ \mathrm{coins}(g^{\mathrm{R}})}}[f(x)\cdot(1-v)] + (1-v)\cdot\mathop{\mathbb{E}}_{\substack{x\sim\mathcal{D}|_{g^{\mathrm{R}}(x)=0},\\ \mathrm{coins}(g^{\mathrm{R}})}}[f(x)\cdot(-v)]\right|$$
$$= v\cdot(1-v)\cdot\left|\mathop{\mathbb{E}}_{\substack{x\sim\mathcal{D}|_{g^{\mathrm{R}}(x)=1},\\ \mathrm{coins}(g^{\mathrm{R}})}}[f(x)] - \mathop{\mathbb{E}}_{\substack{x\sim\mathcal{D}|_{g^{\mathrm{R}}(x)=0},\\ \mathrm{coins}(g^{\mathrm{R}})}}[f(x)]\right| \leq \epsilon.$$

Therefore, the distributions $\mathcal{D}|_{g^{\mathrm{R}}(x)=1}$ and $\mathcal{D}|_{g^{\mathrm{R}}(x)=0}$ are $\left(\mathcal{F}, \frac{\epsilon}{v(1-v)}\right)$-indistinguishable.

(4.) In light of Yao's Lemma, the idea for showing this implication is to define a probability distribution μ such that $\mathbb{E}_{x\sim\mu}[g(x)] = 1/2$. That is, given that g has expected value v when sampling according to distribution $\mathcal{D}$, we want to "shift" v back to $1/2$ when sampling according to distribution μ. Intuitively, if $v > 1/2$, then we should add more weight to the points x in the domain such that $g(x) = 0$, and viceversa if $v \leq 1/2$.

We can do this boosting of the minority values by defining μ as follows:

$$\mu(x) = \begin{cases} \frac{1}{2v}\cdot\mathcal{D}(x) & \text{if } g(x) = 1, \\ \frac{1}{2(1-v)}\cdot\mathcal{D}(x) & \text{if } g(x) = 0. \end{cases}$$

It is direct to check that this is indeed a probability distribution. Next, we show that the expected value of g when sampling according to μ is indeed $1/2$. Let $G^0 = \{x \in \mathcal{X} \mid g(x) = 0\}$ and $G^1 = \{x \in \mathcal{X} \mid g(x) = 1\}$. Then, given that $\mathbb{E}_{\mathcal{D}}[g(x)] = \sum_{x\in G^1}\mathcal{D}(x) = v$, it follows that

$$\mathop{\mathbb{E}}_{x\sim\mu}[g(x)] = \sum_{x\in\mathcal{X}}\mu(x)\cdot g(x) = \sum_{x\in G^1}\mu(x) = \frac{1}{2v}\sum_{x\in G^1}\mathcal{D}(x) = 1/2.$$

Next, we show that μ has density $2b$ in $\mathcal{D}$. This follows directly by our construction of μ, the definition of density for distributions (Definition 2.13) and the definition of b as $b = \min\{v, 1-v\}$: If $x \in g^{-1}(1)$, then $2b\cdot\mu(x) \leq \mathcal{D}(x)$ because $\mu(x) = \frac{1}{2v}\cdot\mathcal{D}(x)$ and $b \leq v$. If $x \in g^{-1}(0)$, then $2b\cdot\mu(x) \leq \mathcal{D}(x)$ as well because $\mu(x) = \frac{1}{2(1-v)}\cdot\mathcal{D}(x)$ and $b \leq 1-v$.

Lastly, we show that g is $(\mathcal{F}, 1/2 - \epsilon'/2)$-hard on μ, where $\epsilon' = \epsilon/(v\cdot(1-v))$. By statement 2.15, we know that the distributions $\mathcal{D}|_{G^1}$ and $\mathcal{D}|_{G^0}$ are $(\mathcal{F}, \epsilon')$-indistinguishable. Let μ_0 correspond to the restriction of μ on the domain G^0, and let μ_1 correspond to the restriction of μ on the domain G^1. By the definition of μ, it follows that

$$\Pr_{\substack{x\sim\mu,\mathrm{coins}(f^{\mathrm{R}}),\\ \mathrm{coins}(g^{\mathrm{R}})}}[f(x) = g(x)] =$$
$$= \frac{1}{2}\Pr_{\substack{x\sim\mu_1,\\ \mathrm{coins}(f^{\mathrm{R}})}}[f(x) = 1] + \frac{1}{2}\Pr_{\substack{x\sim\mu_0,\\ \mathrm{coins}(f^{\mathrm{R}})}}[f(x) = 0]$$
$$= \frac{1}{2} + \frac{1}{2}\cdot\left(\mathop{\mathbb{E}}_{x\sim\mu_1}[f(x)] - \mathop{\mathbb{E}}_{x\sim\mu_0}[f(x)]\right) \leq \frac{1}{2} + \frac{\epsilon'}{2},$$

and hence g is $(\mathcal{F}^{\mathrm{R}}, 1/2 - \epsilon'/2)$-hard on μ, as we wanted to show. We remark that statement (2.15) need not be restricted to boolean functions; in the case of a non-boolean function g we can show it using joint distributions instead. However, the boolean case suffices for our applications of statement (2.15) in this paper. □

Remark 2.16. By the definition of b, it follows that $b/2 \leq v(1-v) \leq b$. Therefore, statement (2.15) in Lemma 2.15 implies that g is $(\mathcal{F}^{\mathrm{R}}, 1/2 - \epsilon/b)$-hard.

Indistinguishability from a constant-Bernoulli function is stronger than average-case hardness. Lemma 2.15 states unidirectional relationships, namely that statements (2.15)-(2.15) are *implied* by g being indistinguishable from a constant-Bernoulli function. It is natural to ask whether the converses also hold, as it is the case in Yao's Lemma (Lemma 2.12). As it can be seen from the proofs of statements (2.15) and (2.15) in Lemma 2.15, these two statements do imply that g is indistinguishable from a constant-Bernoulli function. However, this is *not* the case for statements (2.15), and (2.15), which are *weaker* than being indistinguishable from a constant-Bernoulli function.

The fact that statement (2.15) does not imply the assumption of Lemma 2.15, even when we allow large changes in $\mathcal{F}$ and ϵ can be shown with the following counter-example. Let $\mathcal{D}$ be the uniform distribution on $\mathcal{X} = \{0,1\}^n$ and let $\mathcal{F}$ be all circuits of size n^c. Let g be a random function where $\Pr[g(x) = 1] = 3/4$ if $x_1 = 0$ and $\Pr[g(x) = 1] = 1$ if $x_1 = 1$, where x_1 denotes the first bit of x. Then, by Chernoff and union bounds it can be shown that with high probability, $v \geq 7/8 - 2^{-\Omega(n)}$ and g^{R} is $\left(v - 2^{-\Omega(n)}\right)$-hard against circuits of size $2^{\Omega(n)}$. On the other hand, the distinguisher $f(x) = x_1$ has

$$\mathop{\mathbb{E}}_{x\sim\mathcal{D}}[f(x)(g(x) - 7/8)] = 1/2 - 7/16 = 1/16,$$

so g is not ϵ-indistinguishable from the constant function v for any $\epsilon < 1/16$ and circuits of size $O(1)$. We note that the fact that

statement (2.15) does not imply indistinguishability from a constant-Bernoulli function is in contrast to Yao's Lemma, which states that (2.15) $\iff$ (2.15) in the special case where $v = 1/2$.

Lastly, the reason why (2.15) is weaker than the assumption of Lemma 2.15 is because (2.15) $\implies$ (2.15) by IHCL (given that statement (2.15) corresponds to the assumption of IHCL and statement (2.15) to the conclusion of IHCL), and so it follows that (2.15) does *not* imply that g is a constant-Bernoulli function.

3 THE HARDCORE LEMMA

Impagliazzo's Hardcore Lemma (IHCL) is a fundamental result in complexity theory that dates to 1995 [27]. Informally, it states that if a function is somewhat hard to compute on average by a family $\mathcal{F}$ of boolean functions, then there is a fairly large subset of the inputs (called the "hardcore set") for which the function is very hard to compute, in the sense that g is maximally unpredictable to the family $\mathcal{F}$. That is, no distinguisher can do better than random guessing.

3.1 The Original IHCL Statement

Impagliazzo's Hardcore Lemma can be formally stated as follows:

Theorem 3.1 (IHCL, [27, 24]). *Let $\mathcal{F}$ be a family of functions $f : \mathcal{X} \to \{0,1\}$, let $\mathcal{D}$ be a probability distribution over $\mathcal{X}$, and let $\epsilon, \delta > 0$. There exists $t = \text{poly}(\log|\mathcal{X}|, 1/\epsilon, 1/\delta)$ and $q = \text{poly}(1/\epsilon, 1/\delta)$ such that the following holds: If $g : \mathcal{X} \to \{0,1\}$ is $(\mathcal{F}_{t,q}, \delta)$-hard on $\mathcal{D}$, then there is a distribution $\mathcal{H}$ that is 2δ-dense in $\mathcal{D}$ and for which g is $(\mathcal{F}, 1/2 - \epsilon)$-hard on $\mathcal{H}$.*

3.2 The IHCL++ Theorem

For the ease of notation, recall the definitions of v_P and b_P (which we call the *balance* of function g) from Definition 2.3. From Section 2, recall that we also need to consider the size parameter $\eta_P = \Pr_{x\sim\mathcal{D}}[x \in P]$ of each $P \in \mathcal{P}$ (Definition 2.5). Because we are using the notion of approximate multicalibration, we will only be considering the sets $P \in \mathcal{P}$ such that $\eta_P \geq \gamma$.

We can now introduce our IHCL++ statement:

Theorem 3.2 (IHCL++). *Let $\mathcal{X}$ be a finite domain, let $\mathcal{F}$ be a family of functions $f\colon \mathcal{X} \to [0,1]$, let $g\colon \mathcal{X} \to [0,1]$ be an arbitrary function, $\mathcal{D}$ a probability distribution over $\mathcal{X}$, and let $\epsilon, \gamma > 0$. There exists a partition $\mathcal{P} \in \mathcal{F}_{t,q,k}$ of $\mathcal{X}$ with $t = O(1/(\epsilon^4\gamma) \cdot \log(|\mathcal{X}|/\epsilon))$, $q = O(1/\epsilon^2)$, $k = O(1/\epsilon)$ which satisfies that for all $P \in \mathcal{P}$ such that $\eta_P \geq \gamma$, there exists a distribution $\mathcal{H}_P$ in P of density $2b_P$ in $\mathcal{D}|_P$ such that g^R is $\left(\mathcal{F}^R, 1/2 - \frac{\epsilon}{2b_P(1-b_P)}\right)$-hard on $\mathcal{H}_P$.*

First, we present a proof of our proposed IHCL++ as a direct corollary of our theorem characterizing maximal hardness (Lemma 2.15). Next, we summarize a second proof of IHCL++ by adapting the proof of Trevisan et al. [39]. We include this second proof because it helps in understanding how multicalibration relates to the Regularity Lemma and because we are able to improve the density parameter from δ to 2δ.

Proof of Theorem 3.2. The proof is a combination of the MC Theorem (Theorem 2.10) and our lemma characterizing indistinguishability from constant-Bernoulli functions (Lemma 2.15). We first apply the MC Theorem (Theorem 2.10) to $\mathcal{F}, g, \mathcal{D}$ with the same parameters ϵ, γ. This yields a partition $\mathcal{P} \in \mathcal{F}_{t,q,k}$ of $\mathcal{X}$ with $t = O(1/(\epsilon^4\gamma) \cdot \log(|\mathcal{X}|/\epsilon))$, $q = O(1/\epsilon^2)$, $k = O(1/\epsilon)$ satisfying

$$\left| \mathop{\mathbb{E}}_{x\sim\mathcal{D}|_P} [f(x) \cdot (g(x) - v_P)] \right| \leq \epsilon$$

for all $f \in \mathcal{F}$ and for all $P \in \mathcal{P}$ such that $\eta_P \geq \gamma$. Next, we apply Lemma 2.15 to each piece $P \in \mathcal{P}$ such that $\eta_P \geq \gamma$: Given that g is $(\mathcal{F}, \epsilon)$-indistinguishable from the constant function v_P, by statement (2.15) in Lemma 2.15 it follows that there exists a distribution $\mathcal{H}_P$ of density $2b_P$ in $\mathcal{D}|_P$ such that g^R is $(\mathcal{F}^R, 1/2 - \epsilon_P)$-hard on it, where $\epsilon_P = \frac{\epsilon}{2b_P(1-b_P)}$. Hence $\mathcal{H}_P$ is a hardcore distribution for g on P, as required. □

Another proof of IHCL++ as a modification from [39]. Another way to prove IHCL++ is as a modification of Trevisan et al.'s proof that the Regularity Lemma implies IHCL [39]. We summarize the approach of this proof here; full details appear in [5]. The Trevisan et al. proof of IHCL begins by invoking the Regularity Lemma / Multiaccuracy Theorem (Theorem 1.1) on g and where $\mathcal{D}$ corresponds to the uniform distribution; we begin by invoking the MC Theorem instead (Theorem 2.10).[6] In [39], they use the resulting MA predictor h to define the following distribution $\mathcal{H}$ over the domain $\mathcal{X}$:

$$\mathcal{H}(x) := \frac{|g(x) - h(x)|}{\sum_{y\in\mathcal{X}} |g(y) - h(y)|}.$$

The intuition behind this choice of distribution is to put more mass where h and g disagree. Trevisan et al. then show that 1) $\mathcal{H}$ is δ-dense in $\mathcal{D}$, and that 2) g is strongly hard on $\mathcal{H}$. Inspired by their proof, we define the following probability distribution on each (large enough) set $P \in \mathcal{P}$:

$$\mathcal{H}_P(x) := \frac{|g(x) - v_p|}{\sum_{y\in P} |g(y) - v_p|}.$$

We remark that, unlike in the multiaccuracy case of [39], the denominator in the expression for $\mathcal{H}_P$ sums over the set P instead of over the entirety of the domain $\mathcal{X}$.

We then can show that 1) $\mathcal{H}_P$ is $2b_P$-dense in $\mathcal{D}|_P$, and that 2) g is strongly hard on $\mathcal{H}_P$. To show 1), we analyze the quantity $\sum_{x\in P} |g(x) - v_P|$; the fact that v_P is a constant is what allows us to recover the optimal $2b_P$ density parameter, whereas the same analysis carried out in [39] for IHLC does not. That is, a multiaccurate predictor does not seem to imply IHCL with optimal density parameters, but a multicalibrated predictor can. To show 2), we relate the probability that $f^R(x) = g^R(x)$ to the expected value in the definition of multicalibration using a similar expression to the one used in the proof of Yao's Lemma (i.e., Equation 2.1).

Another proof of IHCL++ using the original IHCL. A third approach to proving our IHCL++ theorem is by using the original IHCL coupled with our theorem characterizing maximal hardness (Theorem 2.15). Namely, we begin by applying the MC theorem to obtain a partition $\mathcal{P}$, thus satisfying the assumption of Lemma 2.15 on each $P \in \mathcal{P}$ such that $\eta_P \geq \gamma$. By statement (2.15) in Lemma 2.15, it follows that the assumption for IHCL is satisfied on each such piece with weak hardness $\delta := b_P - 2\epsilon$, and hence we can apply IHCL

[6] Trevisan et al. consider the restricted version of IHCL where $\mathcal{D}$ corresponds to the uniform distribution on $\mathcal{X}$, whereas we consider the general version of IHCL with an arbitrary distribution.

to each such P to obtain a hardcore distribution within each piece of the partition, which corresponds to the conclusion in IHCL++. However, this yields worse indistinguishability parameters for the hardcore distribution than those in Theorem 3.2.

3.3 Recovering IHCL from IHCL++

Having proved IHCL++, we now show how to recover the original IHCL theorem from it. The key idea is to "glue together" the hardcore distributions $\mathcal{H}_P$ within each $P \in \mathcal{P}$, where in this gluing each $P \in \mathcal{P}$ is weighted according to its size parameter η_P of the set P. When we bring back the assumption that g is δ-weakly hard (as in the original IHCL statement), by using the fact that the multicalibrated partition is of low-complexity, it follows that the glued hardcore distribution $\mathcal{H}$ has density at least 2δ in $\mathcal{D}$, which corresponds to the optimal density parameter in IHCL.

We begin by showing that if g is δ-hard, then the g cannot be too imbalanced on average over the pieces of the partition.

Proposition 3.3. *Let $\mathcal{X}, \mathcal{D}, \mathcal{F}, g, \epsilon, \gamma, \mathcal{P}, t, q, k$ as in Theorem 3.2. Moreover, assume that g is $(\mathcal{F}_{t+k,q}, \delta)$-hard, and suppose that $\eta_P \geq \gamma$ for all $P \in \mathcal{P}$. Then,*

$$\mathop{\mathbb{E}}_{P \sim \mathcal{P}(\mathcal{D})}[b_P] \geq \delta.$$

Proof. Assume by contradiction $\mathbb{E}_{P \sim \mathcal{P}(\mathcal{D})}[b_P] < \delta$. We show that this contradicts the fact that g is δ-hard on $\mathcal{D}$. More specifically, we show that we can construct a function $f \in \mathcal{F}_{t+k,q}$ such that

$$\Pr_{\substack{x \sim \mathcal{D}, \mathrm{coins}(g^{\mathrm{R}}), \\ \mathrm{coins}(f^{\mathrm{R}})}}[f^{\mathrm{R}}(x) = g^{\mathrm{R}}(x)] > 1 - \delta.$$

Let $\hat{f} \in \mathcal{F}_{t,q}$, where $t = O(1/(\epsilon^4\gamma) \cdot \log(|\mathcal{X}|/\epsilon))$, $q = O(1/\epsilon^2)$, be the partition membership function for $\mathcal{P}$ as given by Definition 2.9. That is, $\mathcal{P} = \{\hat{f}^{-1}(1), \ldots, \hat{f}^{-1}(k)\}$. We define our f as $f = f_{\mathrm{post}} \circ \hat{f}$, where $f_{\mathrm{post}} : [k] \to \{0,1\}$ is the indicator function $f_{\mathrm{post}}(i) = \mathbb{1}[v_{\hat{f}^{-1}(i)} \geq 1/2]$. Thus, f_{post} can be by a circuit of size k (see [1, §9.1.1.]), so $f \in \mathcal{F}_{t+k,q}$.

The intuitive meaning of the indicator function $f_{\mathrm{post}}(i)$ is the following: we want to show that f approximates g "quite well", in the sense that $\Pr[f^{\mathrm{R}}(x) = g^{\mathrm{R}}(x)] > 1 - \delta$. The above construction is saying that f is equal to 0 in all of the $P \in \mathcal{P}$ such that $\mathbb{E}_{\mathcal{D}|_P}[g(x)] \leq 1/2$, and equal to 1 otherwise. We now show that this is indeed a good approximation of g; good enough that it contradicts the assumption that g is $(\mathcal{F}_{t+k,q}, \delta)$-hard.

Fix some $P \in \mathcal{P}$, and as usual let $v_P = \mathbb{E}_{x \sim \mathcal{D}|_P}[g(x)]$. Since g is $\{0,1\}$-valued and f equals the majority value of g on P by construction, it follows that

$$\begin{aligned}\Pr_{\substack{x \sim \mathcal{D}|_P, \mathrm{coins}(g^{\mathrm{R}}), \\ \mathrm{coins}(f^{\mathrm{R}})}}[f^{\mathrm{R}}(x) = g^{\mathrm{R}}(x)] &= \Pr_{\substack{x \sim \mathcal{D}|_P, \\ \mathrm{coins}(g^{\mathrm{R}})}}[f(x) = g^{\mathrm{R}}(x)] \\ &= 1 - \min\{v_P, 1 - v_P\} = 1 - b_P,\end{aligned}$$

since $f = \mathbf{0}$ when $v_P < 1/2$ and $f = \mathbf{1}$ when $v_P \geq 1/2$. Because this expression holds for every $P \in \mathcal{P}$, when we consider the probability that $f(x) = g(x)$ over $\mathcal{X}$ it follows that

$$\Pr_{\substack{x \sim \mathcal{D}, \\ \mathrm{coins}(g^{\mathrm{R}})}}[f(x) = g^{\mathrm{R}}(x)] = 1 - \mathop{\mathbb{E}}_{P \sim \mathcal{P}(\mathcal{D})}[b_P].$$

Since by assumption $\mathbb{E}_{P \sim \mathcal{P}(\mathcal{D})}[b_P] < \delta$, it follows that

$$\Pr_{\substack{x \sim \mathcal{D}, \\ \mathrm{coins}(g^{\mathrm{R}})}}[f(x) = g^{\mathrm{R}}(x)] > 1 - \delta,$$

which contradicts the $(\mathcal{F}_{t+k,q}, \delta)$-hardness of g^{R}. □

In Proposition 3.3, we are assuming that $\eta_P \geq \gamma$ for all $P \in \mathcal{P}$ in order to make its proof cleaner. In the full version of the paper, we only "glue" together the pieces $P \in \mathcal{P}$ that have enough size and mass. We then show that, when we bring back the assumption that g is δ-weakly hard, our IHCL++ theorem implies that the glued hardcore set has density at least 2δ.

In the full version of the paper, we also include a set version of IHCL++. The formal statements of PAME++, DMT++, and their corresponding proofs are also deferred to the full version of the paper, which can be found at https://arxiv.org/abs/2312.17223.

ACKNOWLEDGMENTS

We are indebted to Parikshit Gopalan, Fabian Gundlach, Daniel Lee, Huijia (Rachel) Lin, Omer Reingold, Jessica Sorrell, Pranay Tankala, and Benjy Firester for helpful conversations throughout the development of this work. We are also thankful to the Simons' workshop on Multigroup Fairness and the Validity of Statistical Judgement that took place in April 2023, where we presented this work and received insightful feedback. Lastly, we are grateful to the anonymous STOC reviewers for their additional suggestions.

Sílvia Casacuberta was supported by the Herchel Smith Undergraduate Science Research Program. Cynthia Dwork was supported by grant G-2020-13941 of the Alfred P. Sloan Foundation and the Simons Foundation collaboration project 733782. Salil Vadhan was supported by a Simons Investigator Award.

REFERENCES

[1] Boaz Barak. 2022. *Introduction to Theoretical Computer Science*. Creative Commons. https://files.boazbarak.org/introtcs/lnotes_book.pdf.

[2] Noam Barda, Gal Yona, Guy N Rothblum, Philip Greenland, Morton Leibowitz, Ran Balicer, Eitan Bachmat, and Noa Dagan. 2021. Addressing bias in prediction models by improving subpopulation calibration. *Journal of the American Medical Informatics Association*, 28, 3, 549–558.

[3] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2023. *Fairness and machine learning: Limitations and opportunities*. MIT Press.

[4] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*. PMLR, 77–91.

[5] Sílvia Casacuberta Puig. 2023. *Finding Simple Models of Complex Objects: From Regularity Lemmas to Algorithmic Fairness*. Bachelor's thesis. Harvard University. https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37376430.

[6] Yi-Hsiu Chen, Kai-Min Chung, and Jyun-Jie Liao. 2018. On the complexity of simulating auxiliary input. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III* (Lecture Notes in Computer Science). Jesper Buus Nielsen and Vincent Rijmen, (Eds.) Vol. 10822. Springer, 371–390.

[7] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5, 2, 153–163. https://doi.org/10.1089/big.2016.0047.

[8] Kai-Min Chung, Edward Lui, and Rafael Pass. 2015. From weak to strong zero-knowledge and applications. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part I* (Lecture Notes in Computer Science). Yevgeniy Dodis and Jesper Buus Nielsen, (Eds.) Vol. 9014. Springer, 66–92.

[9] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. 2008. Fuzzy extractors: how to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38, 1, 97–139. https://doi.org/10.1137/060651380.

[10] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. 2012. Fairness through awareness. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*. Shafi Goldwasser, (Ed.) ACM, 214–226. https://doi.org/10.1145/2090236.2090255.

[11] Cynthia Dwork, Michael P. Kim, Omer Reingold, Guy N. Rothblum, and Gal Yona. 2021. Outcome indistinguishability. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*. Samir Khuller and Virginia Vassilevska Williams, (Eds.) ACM, 1095–1108. https://doi.org/10.1145/3406325.3451064.

[12] Cynthia Dwork, Daniel Lee, Huijia Lin, and Pranay Tankala. 2023. From pseudorandomness to multi-group fairness and back. In *The Thirty Sixth Annual Conference on Learning Theory, COLT 2023, 12-15 July 2023, Bangalore, India* (Proceedings of Machine Learning Research). Gergely Neu and Lorenzo Rosasco, (Eds.) Vol. 195. PMLR, 3566–3614. https://proceedings.mlr.press/v195/dwork23a.html.

[13] Alan M. Frieze and Ravi Kannan. 1999. Quick approximation to matrices and applications. *Comb.*, 19, 2, 175–220. https://doi.org/10.1007/s004930050052.

[14] Sumegha Garg, Christopher Jung, Omer Reingold, and Aaron Roth. 2023. Oracle efficient online multicalibration and omniprediction. *CoRR*, abs/2307.08999. arXiv: 2307.08999.

[15] Craig Gentry and Daniel Wichs. 2011. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*. Lance Fortnow and Salil P. Vadhan, (Eds.) ACM, 99–108. https://doi.org/10.1145/1993636.1993651.

[16] Ira Globus-Harris, Declan Harrison, Michael Kearns, Aaron Roth, and Jessica Sorrell. 2023. Multicalibration as boosting for regression. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA* (Proceedings of Machine Learning Research). Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, (Eds.) Vol. 202. PMLR, 11459–11492. https://proceedings.mlr.press/v202/globus-harris23a.html.

[17] Oded Goldreich, Noam Nisan, and Avi Wigderson. 2011. On Yao's XOR-Lemma. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation - In Collaboration with Lidor Avigad, Mihir Bellare, Zvika Brakerski, Shafi Goldwasser, Shai Halevi, Tali Kaufman, Leonid Levin, Noam Nisan, Dana Ron, Madhu Sudan, Luca Trevisan, Salil Vadhan, Avi Wigderson, David Zuckerman*. Lecture Notes in Computer Science. Vol. 6650. Oded Goldreich, (Ed.) Springer, 273–301. https://doi.org/10.1007/978-3-642-22670-0%5C_23.

[18] Parikshit Gopalan, Adam Tauman Kalai, Omer Reingold, Vatsal Sharan, and Udi Wieder. 2022. Omnipredictors. In *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA* (LIPIcs). Mark Braverman, (Ed.) Vol. 215. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 79:1–79:21. https://doi.org/10.4230/LIPIcs.ITCS.2022.79.

[19] Parikshit Gopalan, Michael P. Kim, and Omer Reingold. 2023. Characterizing notions of omniprediction via multicalibration. *CoRR*, abs/2302.06726. arXiv: 2302.06726.

[20] Parikshit Gopalan, Michael P. Kim, Mihir Singhal, and Shengjia Zhao. 2022. Low-degree multicalibration. In *Conference on Learning Theory, 2-5 July 2022, London, UK* (Proceedings of Machine Learning Research). Po-Ling Loh and Maxim Raginsky, (Eds.) Vol. 178. PMLR, 3193–3234. https://proceedings.mlr.press/v178/gopalan22a.html.

[21] Parikshit Gopalan, Omer Reingold, Vatsal Sharan, and Udi Wieder. 2022. Multicalibrated partitions for importance weights. In *International Conference on Algorithmic Learning Theory, 29 March - 1 April 2022, Paris, France* (Proceedings of Machine Learning Research). Sanjoy Dasgupta and Nika Haghtalab, (Eds.) Vol. 167. PMLR, 408–435. https://proceedings.mlr.press/v167/gopalan22a.html.

[22] Ben Green and Terence Tao. 2008. The primes contain arbitrarily long arithmetic progressions. *Ann. of Math. (2)*, 167, 2, 481–547. https://doi.org/10.4007/annals.2008.167.481.

[23] Úrsula Hébert-Johnson, Michael P. Kim, Omer Reingold, and Guy N. Rothblum. 2018. Multicalibration: calibration for the (computationally-identifiable) masses. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018* (Proceedings of Machine Learning Research). Jennifer G. Dy and Andreas Krause, (Eds.) Vol. 80. PMLR, 1944–1953. http://proceedings.mlr.press/v80/hebert-johnson18a.html.

[24] Thomas Holenstein. 2005. Key agreement from weak bit agreement. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*. Harold N. Gabow and Ronald Fagin, (Eds.) ACM, 664–673. https://doi.org/10.1145/1060590.1060689.

[25] Christina Ilvento. 2020. Metric learning for individual fairness. In *1st Symposium on Foundations of Responsible Computing, FORC 2020, June 1-3, 2020, Harvard University, Cambridge, MA, USA (virtual conference)* (LIPIcs). Aaron Roth, (Ed.) Vol. 156. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2:1–2:11. https://doi.org/10.4230/LIPIcs.FORC.2020.2.

[26] Russell Impagliazzo. 2009. Algorithmic dense model theorems and weak regularity. In Unpublished manuscript. https://simons-institute.github.io/pseudorandomness/pdfs/nov09.pdf.

[27] Russell Impagliazzo. 1995. Hard-core distributions for somewhat hard problems. In *36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, USA, 23-25 October 1995*. IEEE Computer Society, 538–545. https://doi.org/10.1109/SFCS.1995.492584.

[28] Russell Impagliazzo. 2008. When do sparse sets have dense models. In Talk at the Pseudorandomness in Mathematics and Computer Science Miniworkshop. Institute for Advanced Study. https://www.ias.edu/sites/default/files/math/russell_scribe.pdf.

[29] Russell Impagliazzo, Cristopher Moore, and Alexander Russell. 2014. An entropic proof of chang's inequality. *SIAM J. Discret. Math.*, 28, 1, 173–176.

[30] Dimitar Jetchev and Krzysztof Pietrzak. 2014. How to fake auxiliary input. In *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings* (Lecture Notes in Computer Science). Yehuda Lindell, (Ed.) Vol. 8349. Springer, 566–590.

[31] Michael J. Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing fairness gerrymandering: auditing and learning for subgroup fairness. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018* (Proceedings of Machine Learning Research). Jennifer G. Dy and Andreas Krause, (Eds.) Vol. 80. PMLR, 2569–2577. http://proceedings.mlr.press/v80/kearns18a.html.

[32] Michael P Kim, Christoph Kern, Shafi Goldwasser, Frauke Kreuter, and Omer Reingold. 2022. Universal adaptability: target-independent inference that competes with propensity scoring. *Proceedings of the National Academy of Sciences*, 119, 4.

[33] Michael P. Kim, Amirata Ghorbani, and James Y. Zou. 2019. Multiaccuracy: black-box post-processing for fairness in classification. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society, AIES 2019, Honolulu, HI, USA, January 27-28, 2019*. Vincent Conitzer, Gillian K. Hadfield, and Shannon Vallor, (Eds.) ACM, 247–254. https://doi.org/10.1145/3306618.3314287.

[34] Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. 2016. How we analyzed the compas recidivism algorithm. *ProPublica (5 2016)*, 9, 1.

[35] Georgy Noarov, Ramya Ramalingam, Aaron Roth, and Stephan Xie. 2023. High-dimensional prediction for sequential decision making. *CoRR*, abs/2310.17651. arXiv: 2310.17651.

[36] Omer Reingold, Luca Trevisan, Madhur Tulsiani, and Salil P. Vadhan. 2008. Dense subsets of pseudorandom sets. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*. IEEE Computer Society, 76–85. https://doi.org/10.1109/FOCS.2008.38.

[37] Endre Szemerédi. 1978. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*. Colloq. Internat. CNRS. Vol. 260. CNRS, Paris, 399–401. ISBN: 2-222-02070-0.

[38] Terence Tao and Tamar Ziegler. 2008. The primes contain arbitrarily long polynomial progressions. *Acta Math.*, 201, 2, 213–305. https://doi.org/10.1007/s11511-008-0032-5.

[39] Luca Trevisan, Madhur Tulsiani, and Salil P. Vadhan. 2009. Regularity, boosting, and efficiently simulating every high-entropy distribution. In *Proceedings of the 24th Annual IEEE Conference on Computational Complexity, CCC 2009, Paris, France, 15-18 July 2009*. IEEE Computer Society, 126–136. https://doi.org/10.1109/CCC.2009.41.

[40] Salil P. Vadhan and Colin Jia Zheng. 2013. A uniform min-max theorem with applications in cryptography. In *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I* (Lecture Notes in Computer Science). Ran Canetti and Juan A. Garay, (Eds.) Vol. 8042. Springer, 93–110.

[41] Salil P. Vadhan and Colin Jia Zheng. 2012. Characterizing pseudoentropy and simplifying pseudorandom generator constructions. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*. Howard J. Karloff and Toniann Pitassi, (Eds.) ACM, 817–836. https://doi.org/10.1145/2213977.2214051.

[42] James Vincent. 2018. Amazon reportedly scraps internal ai recruiting tool that was biased against women. *The Verge*, 10.

[43] Andrew Chi-Chih Yao. 1982. Theory and applications of trapdoor functions (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science, Chicago, Illinois, USA, 3-5 November 1982*. IEEE Computer Society, 80–91. https://doi.org/10.1109/SFCS.1982.45.

[44] Colin Jia Zheng. 2014. *A uniform min-max theorem and characterizations of computational randomness*. Ph.D. Dissertation. Harvard University. http://nrs.harvard.edu/urn-3:HUL.InstRepos:11745716.

Received 13-NOV-2023; accepted 2024-02-11