

From Pseudorandomness to Multi-Group Fairness and Back

Cynthia Dwork
Harvard University

DWORK@SEAS.HARVARD.EDU

Daniel Lee
Massachusetts Institute of Technology

LEE.D@MIT.EDU

Huijia Lin
University of Washington

RACHEL@CS.WASHINGTON.EDU

Pranay Tankala
Harvard University

PRANAY_TANKALA@G.HARVARD.EDU

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

We identify and explore connections between the recent literature on *multi-group fairness* for prediction algorithms and the pseudorandomness notions of *leakage-resilience* and *graph regularity*. We frame our investigation using new variants of multicalibration based on statistical distance and closely related to the concept of *outcome indistinguishability*. Adopting this perspective leads us not only to new, more efficient algorithms for multicalibration, but also to our graph theoretic results and a proof of a novel *hardcore lemma* for real-valued functions.

Keywords: Fairness, calibration, pseudorandomness, regularity, weak regularity, leakage resilience

1. Introduction

A central question in the field of algorithmic fairness concerns the extent to which prediction algorithms, which assign numeric “probabilities” to individuals in a population, systematically mistreat members of large demographic subpopulations. Although interest in *group fairness* is far from new [Turini \(2008\)](#); [Kamiran and Calders \(2009\)](#), the study of *multi-group fairness*, originally conceived as a bridge between *individual fairness* [Dwork et al. \(2012\)](#) and group fairness notions and in which the subpopulations of interest are numerous and overlapping, is relatively young, initiated by the seminal works of [Hébert-Johnson et al. \(2018\)](#); [Kearns et al. \(2018\)](#). In the past few years, a fruitful line of research has investigated how to achieve various notions of multi-group fairness and their applications to learning e.g., [Gopalan et al. \(2022a\)](#); [Kim et al. \(2022\)](#); [Deng et al. \(2023\)](#); [Gopalan et al. \(2022b\)](#); [Jung et al. \(2021\)](#); [Gupta et al. \(2022\)](#). In this work, complementing the work of [Casacuberta et al. \(2023\)](#), we excavate a deep connection between multi-group fairness and pseudorandomness, and exhibit a productive relationship between the key concept of multicalibration introduced in [Hébert-Johnson et al. \(2018\)](#), and notions of leakage simulation [Jetchev and Pietrzak \(2014\)](#), graph regularity [Szemerédi \(1975\)](#); [Frieze and Kannan \(1996\)](#), and hardcore lemmas [Impagliazzo \(1995\)](#), drawing particular inspiration from [Trevisan et al. \(2009\)](#); [Chen et al. \(2018\)](#).

A New Definition. Speaking informally, multicalibration requires that predictions be calibrated simultaneously on each member c of a pre-specified collection \mathcal{C} of arbitrarily intersecting population groups, the intuition being that a score of v “means the same thing” independent of one’s group membership(s) Kleinberg et al. (2016). We begin with a new variant of the definition of multicalibration that draws on the indistinguishability-based point of view of Dwork et al. (2021) and generalizes prior definitions along several axes. Consider a distribution \mathcal{D} on individual-outcome pairs (i, o_i^*) for individuals $i \in \mathcal{X}$ and their associated real-world outcomes o_i^* , and a collection \mathcal{C} of functions capturing intersecting subpopulations in the fairness-based view and *distinguishers* in the outcome-indistinguishability framework. Letting p^* denote the real-world outcome distribution, where $o_i^* \sim \text{Ber}(p_i^*)$, the goal of a predictor \tilde{p} is to provide outcome distributions \tilde{p}_i , for $i \in \mathcal{X}$, in a way that cannot be distinguished from p^* by the distinguishers. Unlike in previous work, our definition of multicalibration is in terms of *statistical distance*, requiring that for all $c \in \mathcal{C}$,

$$(c(i), \tilde{o}_i, \tilde{p}_i) \text{ and } (c(i), o_i^*, \tilde{p}_i) \text{ are } \varepsilon\text{-statistically close, where } (i, o_i^*) \leftarrow \mathcal{D} \text{ and } \tilde{o}_i \sim \text{Ber}(\tilde{p}_i).$$

(See Section 3 for a formal treatment.) Our definition naturally accommodates outcomes $o_i^* \in \mathcal{O}$ for an arbitrary set \mathcal{O} of possible outcomes, as well as functions $c : \mathcal{X} \rightarrow \mathcal{Y}$ with arbitrary ranges. A weaker variant of our definition corresponds to *multiaccuracy* Hébert-Johnson et al. (2018); Kim et al. (2019) and equivalently *no-access* outcome indistinguishability Dwork et al. (2021) and the approximation notion in Theorem 1.1 of Trevisan et al. (2009), which in the group-fairness view only requires the predictor be accurate in expectation (rather than calibrated) on each group simultaneously. We also describe a stronger variant called *strict multicalibration*, which is closely related to the notion of “swap” multicalibration that was independently proposed by the concurrent work of Gopalan et al. (2023b). The strong statistical distance condition in our definitions lends itself naturally to applications and gives a strikingly simple proof of the observation, due to Gopalan et al. (2023a,b), that *omniprediction* Gopalan et al. (2022a) can be achieved from multiaccuracy and overall calibration¹, as well as a generalization of omniprediction to new settings. These new settings include the *multi-objective learning* problem studied in Haghtalab et al. (2023), where the goal is to achieve strong performance not only across multiple loss functions but also across shifts in the underlying distribution over \mathcal{X} .

From Pseudorandomness to Fairness. The leakage simulation lemma Jetchev and Pietrzak (2014) is a cryptographic result concerning when a few bits of auxiliary input regarding a secret can be “faked,” and therefore pose no threat to secrecy. Translating to our setting, in the simplest case there is a single bit of “auxiliary” input and this corresponds to $o_i^* \sim \text{Ber}(p_i^*)$ or $\tilde{o}_i \sim \text{Ber}(\tilde{p}_i)$, where again p_i^* is the true distribution from which individual i ’s outcome is chosen and \tilde{p}_i is the distribution proposed by the predictor. The lemma provides a construction for creating a simulator that outputs “fake” bits that fool any function in a family \mathcal{F} of distinguishers that receive (i, o_i^*) or (i, \tilde{o}_i) , and is a strengthening of a conceptually similar result in Trevisan et al. (2009). The simulator is a simple combination of only a small number of functions in \mathcal{F} .

Upon inspection, multiaccuracy is a “moral equivalent” to leakage simulation. Armed with this observation, we leverage a lower bound on the size of leakage simulators Chen et al. (2018) to obtain the first relative lower bound on the size of *multiaccurate* predictors, refuting the possibility of having “dream” predictors that are more efficient than functions in \mathcal{C} . In other words, there is

1. An omnipredictor allows post-processing to obtain best-in-class (with respect to \mathcal{C}) loss with respect to any loss function in a rich set.

nothing analogous to a pseudorandom generator, *i.e.*, no small predictor \tilde{p} can fool all polynomial-sized distinguishers.

Inspired by the specific leakage simulation algorithm of [Chen et al. \(2018\)](#), we next construct a general framework for multicalibration algorithms using no-regret learning. A particular instantiation of the framework results in a set of new algorithms with improved sample complexity in the *multiclass* and *low-degree* settings recently introduced by [Gopalan et al. \(2022b\)](#)².

Specifically, when $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$ is the collection of demographic subpopulations and there are $\ell \gg 2$ possible outcomes for each member of the population, our algorithm in Section 5 needs only $\log |\mathcal{C}| + (1/\varepsilon)^\ell$ samples to achieve $(\mathcal{C}, O(\varepsilon))$ -multicalibration, ignoring factors of $(\ell/\varepsilon)^{O(1)}$. In contrast, the previous best upper bound of [Gopalan et al. \(2022b\)](#) required $\log |\mathcal{C}| \times (\ell/\varepsilon)^{4\ell}$, again ignoring factors of $(\ell/\varepsilon)^{O(1)}$.³ In particular, in our sample complexity, the $\log |\mathcal{C}|$ term is additive and the dependence on ℓ is simply $\exp(\ell)$. Similarly, for degree- k multicalibration our algorithm uses $\Theta((\log |\mathcal{C}| + k \log(\ell/\varepsilon))\varepsilon^{-4} \times \log(\ell))$ samples, yielding an exponential improvement on the dependence in ℓ over the bound in [Gopalan et al. \(2022b\)](#) of $\Theta((\log |\mathcal{C}| + k \log(\ell/\varepsilon))\varepsilon^{-4} \times \ell)$.

From Fairness Back to Pseudorandomness We find a tantalizing parallel between multicalibration and the *Szemerédi regularity lemma* from extremal graph theory, which decomposes large dense graphs into parts that behave pseudorandomly [Szemerédi \(1975\)](#). We show in Theorem 19 that for an appropriate graph-based instantiation of the multi-group fairness framework, there is a tight correspondence between predictors satisfying *strict* multicalibration and Szemerédi regularity partitions of the underlying graph. Our theorem can be viewed as an extension of a result of [Trevisan et al. \(2009\)](#) that established an analogous link between *multiaccuracy* and the *Frieze-Kannan weak regularity lemma* [Frieze and Kannan \(1996\)](#). It also builds on the work of [Skorski \(2017\)](#), which showed that the criteria of Szemerédi regularity can be phrased in terms of distinguishers—our work further shows that these distinguishers have the structure of tests for strict multicalibration with respect to an appropriate collection \mathcal{C} . Finally, by considering the standard notion of multicalibration, which is more demanding than multiaccuracy but less so than strict multicalibration, our analogy naturally leads us at a new notion of graph regularity situated between Frieze-Kannan regularity and Szemerédi regularity that we call *intermediate regularity*, which may be of independent interest.

The regularity lemma of Trevisan, Tulsiani, and Vadhan yields important implications in different areas, including the weak Szemerédi regularity lemma in graph theory, Impagliazzo’s Hardcore Lemma in complexity theory [Impagliazzo \(1995\)](#), the Dense Model Theorem in additive combinatorics [Impagliazzo \(2009\)](#); [Reingold et al. \(2008\)](#), computational analogues of entropy in information theory [Vadhan and Zheng \(2012, 2013\)](#); [Zheng \(2014\)](#), and weaker notions of zero-knowledge in cryptography [Chung et al. \(2015\)](#). Capitalizing on the increased strength of multicalibration over multiaccuracy, applying our multicalibration algorithm we derive a version of the hardcore lemma for bounded *real-valued* functions with respect to natural notions of hardness and pseudorandomness⁴.

2. [Gopalan et al. \(2022b\)](#) defines a hierarchy of relaxations of multicalibration in which the k th level (“degree k ”) constrains the first k moments of the predictor, conditioned on subpopulations in \mathcal{C} , and demonstrates that some properties of multicalibration related to fairness and accuracy manifest as low-degree properties.

3. [Gopalan et al. \(2022b\)](#) reports sample complexity in terms of $\text{VC}(\mathcal{C})$, which could be as large as $\log(|\mathcal{C}|)$.

4. In concurrent work also capitalizing on the strength of multicalibration as a starting point, [Casacuberta et al. \(2023\)](#) obtains stronger and more general versions of the Hardcore Lemma, Dense Model Theorem, and characterizations of pseudo-average min-entropy.

Organization In Section 2, we formally state the problem setup for multi-group fairness, with an emphasis on the case of multiclass prediction. In Section 3, we present our new notions of multicalibration, relate them to prior definitions in the literature, and discuss their applications. In Section 4, we discuss outcome indistinguishability and its relationship to leakage simulation, a connection that motivates several of our results. In Section 5, we give an algorithm template that unifies prior algorithms for achieving outcome indistinguishability, and also derive our improved sample complexity upper bounds. In Section 6, we detail the connection to graph regularity. In Section 7, we prove our novel variant of the hardcore lemma.

General Notation For sets A and B , we let B^A denote the set of functions $f : A \rightarrow B$. For $f \in B^A$, we let $f(a)$ and f_a both denote the output of f on input $a \in A$.

2. The Multi-Group Fairness Framework

Individuals and Outcomes Building on the framework introduced by Dwork et al. (2021), we consider a pair (i, o_i^*) of jointly distributed random variables, where i is an *individual* drawn from some fixed distribution over a finite *population* \mathcal{X} and o_i^* is an *outcome* of individual i that belongs to a finite set \mathcal{O} consisting of $\ell = |\mathcal{O}|$ possible outcomes.

Modeled Outcomes A *predictor* associates a probability distribution over possible outcomes to each member of the population. In other words, a predictor is a function $\tilde{p} : \mathcal{X} \rightarrow \Delta\mathcal{O}$, where $\Delta\mathcal{O} = \{f \in [0, 1]^{\mathcal{O}} : \sum_{o \in \mathcal{O}} f(o) = 1\}$. Let \tilde{p}_j denote the output $\tilde{p}(j)$ of the function \tilde{p} on input $j \in \mathcal{X}$, and let $\tilde{o}_i \in \mathcal{O}$ be a random variable whose conditional distribution given i is specified by \tilde{p}_i . In other words, $\Pr[\tilde{o}_i = o \mid i] = \tilde{p}_i(o)$ for each possible outcome $o \in \mathcal{O}$. We call \tilde{o}_i the *modeled outcome* of individual i .

Binary Outcomes We say that outcomes are *binary* if $\mathcal{O} = \{0, 1\}$. In this case, we can naturally identify $\Delta\mathcal{O}$ with the unit interval $[0, 1]$ by mapping the distribution $\tilde{p}_j \in \Delta\mathcal{O}$ to the probability $\tilde{p}_j(1) \in [0, 1]$ that \tilde{p}_j assigns to a positive outcome. With this convention, the conditional distribution of \tilde{o}_i given i is $\text{Ber}(\tilde{p}_i)$, which is how \tilde{o}_i was originally defined in Dwork et al. (2021).

Demographic Subpopulations Multi-group fairness examines the ways that a predictor \tilde{p} might mistreat members of large, possibly overlapping *subpopulations* $S \subseteq \mathcal{X}$ in a prespecified collection \mathcal{C} . Each such subpopulation has an associated indicator function $\mathbf{1}_S : \mathcal{X} \rightarrow \{0, 1\}$, and, following Gopalan et al. (2022a), it will be notationally convenient for us to represent \mathcal{C} directly as a collection of such functions. Concretely, we allow \mathcal{C} to be any collection of functions $c : \mathcal{X} \rightarrow \mathcal{Y}$ for some set \mathcal{Y} . For consistency, we will write the output $c(j)$ of the function c on input $j \in \mathcal{X}$ as c_j .

Discretization We will sometimes round predictions to the nearest point in a finite set $\mathcal{G} \subseteq \Delta\mathcal{O}$, which we assume to be an η -covering of $\Delta\mathcal{O}$ with respect to the *statistical distance metric* $\delta(f, g) = \frac{1}{2} \sum_{o \in \mathcal{O}} |f(o) - g(o)|$, meaning that for all $f \in \Delta\mathcal{O}$, there exists $g \in \mathcal{G}$ such that $\delta(f, g) \leq \eta$. Formally, we say that $\hat{p} : \mathcal{X} \rightarrow \mathcal{G}$ is the *discretization* of \tilde{p} to \mathcal{G} if $\hat{p}_j = \operatorname{argmin}_{g \in \mathcal{G}} \delta(\tilde{p}_j, g)$ for each $j \in \mathcal{X}$, breaking ties arbitrarily. We define $\hat{o}_i \in \mathcal{O}$ to be the modeled outcome of an individual i with respect to the predictor \hat{p} , meaning that $\Pr[\hat{o}_i = o \mid i] = \hat{p}_i(o)$ for each $o \in \mathcal{O}$.

The size of \mathcal{G} will also play an important role in deriving our improved complexity upper bounds in Section 5. We show in Section 3 that by taking \mathcal{G} to be the intersection of $\Delta\mathcal{O}$ with the grid $\{0, \frac{1}{m}, \frac{2}{m}, \dots, \frac{m-1}{m}, 1\}^\ell$ for an appropriate integer m , we achieve $|\mathcal{G}| < (3/\eta)^{\ell-1}$. In this case, discretization to \mathcal{G} amounts to coordinate-wise rounding.

Finally, we remark that discretization merges a predictor’s level sets, and this process may destroy any special structure these level sets possess. This observation will become important in our discussion of graph regularity in Section 6.

3. Multicalibration via Statistical Closeness

Multiaccuracy and *multicalibration* are two essential multi-group fairness notions introduced by Hébert-Johnson et al. (2018). In this section, we state new, natural variants of these definitions in terms of *statistical distance*, which, for random variables X and Y with finite support, is measured by the formula

$$\delta(X, Y) = \max_A |\Pr[X \in A] - \Pr[Y \in A]|.$$

The maximum is taken over $A \subseteq \text{supp}(X) \cup \text{supp}(Y)$, and we write $X \approx_\varepsilon Y$ if $\delta(X, Y) \leq \varepsilon$.

Definition 1 A predictor \tilde{p} is statistically $(\mathcal{C}, \varepsilon)$ -multiaccurate if $(c_i, \tilde{o}_i) \approx_\varepsilon (c_i, o_i^*)$ for all $c \in \mathcal{C}$.

Definition 2 A predictor \tilde{p} is statistically $(\mathcal{C}, \varepsilon)$ -multicalibrated if $(c_i, \tilde{o}_i, \tilde{p}_i) \approx_\varepsilon (c_i, o_i^*, \tilde{p}_i)$ for all $c \in \mathcal{C}$.

A notable way in which Definitions 1 and 2 differ from previous definitions in the multi-group fairness literature is that in the case of $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$, our definitions concern the behavior of \tilde{p} on *both* the 0-level set and 1-level set of each $c \in \mathcal{C}$, as opposed to merely the 1-level sets. We will demonstrate shortly that this distinction is unimportant if $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$ is closed under complement, meaning that $\mathbf{1}_S \in \mathcal{C}$ if and only if $\mathbf{1}_{\mathcal{X} \setminus S} \in \mathcal{C}$. We also emphasize that our definition generalizes to the case where c has an arbitrary range \mathcal{Y} ; we do not require that \mathcal{Y} is either $\{0, 1\}$ or $[0, 1]$.

Henceforth, we will write “multiaccuracy” and “multicalibration” to refer to statistical multiaccuracy and statistical multicalibration, respectively, unless otherwise specified. However, when discussing results that crucially rely on these new, statistical distance-based definitions, we will refrain from using these abbreviations.

In our discussion of graph regularity in Section 6, we will also need a stronger variant of the multicalibration definition, which we call *strict multicalibration*, that has appeared only implicitly in prior works on algorithmic fairness and has a close connection to the notion of “swap” multicalibration independently proposed by Gopalan et al. (2023b). To state the definition succinctly, we introduce the shorthand

$$\delta(X, Y \mid Z) = \max_A |\Pr[X \in A \mid Z] - \Pr[Y \in A \mid Z]|,$$

which is a function of a random variable Z distributed jointly with X and with Y . Specifically, if the value of Z is z , then the value of $\delta(X, Y \mid Z)$ is

$$\delta(X, Y \mid Z = z) = \max_A |\Pr[X \in A \mid Z = z] - \Pr[Y \in A \mid Z = z]|.$$

Definition 3 A predictor \tilde{p} is strictly $(\mathcal{C}, \varepsilon)$ -multicalibrated if

$$\mathbf{E} \left[\max_{c \in \mathcal{C}} \delta((c_i, \tilde{o}_i), (c_i, o_i^*) \mid \tilde{p}_i) \right] \leq \varepsilon.$$




Figure 1: Implications (solid arrow: rounding not required, dashed arrow: rounding required)

Intuitively, strict multicalibration asks that a predictor be multiaccurate on most of its level sets. As we will see shortly, Definition 2, which is closest to the original definition of multicalibration and suffices for some applications, does *not* require multiaccuracy on even a single level set.

In Appendix A, we prove that strict $(\mathcal{C}, \varepsilon)$ -multicalibration implies $(\mathcal{C}, \varepsilon)$ -multicalibration, which in turn implies $(\mathcal{C}, \varepsilon)$ -multiaccuracy. In Appendix B, we prove that the following theorem, which shows that the converse implications do not hold.

Theorem 4 *For all $\varepsilon > 0$, there exist finite sets \mathcal{X} and $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$, a pair of joint random variables $(i, o_i^*) \in \mathcal{X} \times \{0, 1\}$, and predictors $\tilde{p}, \tilde{p}' : \mathcal{X} \rightarrow [0, 1]$ such that:*

- (a) \tilde{p} is $(\mathcal{C}, 0)$ -multiaccurate but not $(\mathcal{C}, 1/3)$ -multicalibrated.
- (b) \tilde{p}' is $(\mathcal{C}, \varepsilon)$ -multicalibrated but not strictly $(\mathcal{C}, 1/4)$ -multicalibrated.

The separation of multicalibration and strict multicalibration comes with an important caveat: any multicalibrated predictor \tilde{p} can be *discretized* to achieve strict multicalibration with respect to the same collection \mathcal{C} but a significantly worse parameter ε . To state this result, proved in Appendix C, recall that \hat{p}_i denotes the discretization of \tilde{p}_i to a finite η -covering \mathcal{G} of $\Delta_{\mathcal{O}}$.

Theorem 5 *If \tilde{p} is $(\mathcal{C}, \varepsilon)$ -multicalibrated, then \hat{p} is strictly $(\mathcal{C}, |\mathcal{G}|\varepsilon + \eta)$ -multicalibrated.*

Corollary 6 *Let $\ell = |\mathcal{O}|$. For sufficiently small $\varepsilon > 0$, any $(\mathcal{C}, \varepsilon^\ell)$ -multicalibrated predictor can be made strictly $(\mathcal{C}, 4\varepsilon)$ -multicalibrated by coordinate-wise rounding to a precision depending only on ε and ℓ .*

The relationships among these definitions are depicted in Figure 1. In Appendix D, we explain the relationships of these definitions to existing notions of multiaccuracy and multicalibration that appear in the algorithmic fairness literature. Strict multicalibration was not been explicitly defined in the original works on multicalibration. Nevertheless, their algorithms actually achieve this stronger notion. In Appendix E, we demonstrate the usefulness of our new definitions by (re)proving the recent interesting results on omniprediction by Gopalan et al. (2022a, 2023a). The statistical distance condition of our definitions yields a very simple and concise proof and extends the state-of-the-art to the multiclass prediction setting. We further extend the concept of omniprediction to consider loss functions that may depend on information of individuals.

A few remarks are in order. Although our statistical distance-based definitions handle real-valued functions c , our algorithms for achieving them only work with discretized ranges. When c has continuous range, c_i can completely describe i , and statistical closeness would then force \tilde{p} to be essentially equal to p^* . At the same time, we *can* achieve covariance-based multicalibration for continuous functions c with range $[0, 1]$ by viewing c_i as a probability distribution and replacing c_i with a random instantiation $b_i \sim \text{Ber}(c_i)$. This gives rise to a weaker statistical closeness condition, namely $(b_i, \tilde{o}_i, \tilde{p}_i) \approx_\varepsilon (b_i, o_i^*, \tilde{p}_i)$ for all $c \in \mathcal{C}$, that is nonetheless sufficient for some

applications (e.g. Lemma 64 in Appendix J). In particular, if an application is only concerned with expectations of the form $\mathbf{E}[f(o, \tilde{p}_i)c_i]$ for an arbitrary function f , where c_i is used in a linear fashion, the weaker statistical closeness condition suffices because $\mathbf{E}[f(o_i^*, \tilde{p}_i)c_i] = \mathbf{E}[f(o_i^*, \tilde{p}_i)b_i]$ and $\mathbf{E}[f(\tilde{o}_i, \tilde{p}_i)c_i] = \mathbf{E}[f(\tilde{o}_i, \tilde{p}_i)b_i]$.

4. Leakage Simulation and Outcome Indistinguishability

The interesting relationship between multiaccuracy and the *leakage simulation lemma*, or *simulating auxiliary inputs* problem, in cryptography on the one hand allows us to obtain the first lower bound on the complexity of multiaccurate predictors. On the other hand, it inspires us to ask whether the stronger notion of multicalibration yields stronger consequences. We show this is the case, deriving a multicalibration-based proof of a hardcore lemma for *real-valued* functions.

Originating in the field of *leakage-resilient cryptography* Dziembowski and Pietrzak (2008), the problem of leakage simulation defined by Jetchev and Pietrzak (2014) is as follows. Given correlated random variables (X, O) on a set $\mathcal{X} \times \mathcal{O}$ and a collection of *distinguisher functions* $\mathcal{A} = \{A : \mathcal{X} \times \mathcal{O} \rightarrow \{0, 1\}\}$, the objective is to construct a low-complexity (w.r.t. \mathcal{A}) *simulator* $h : \mathcal{X} \rightarrow \Delta\mathcal{O}$ such that no function in \mathcal{A} can distinguish a sample (X, O) from the true joint distribution from a simulated sample (X, \tilde{O}) , where X is sampled from the true marginal distribution over \mathcal{X} and \tilde{O} is sampled according to the simulated distribution $h(X)$.

Observe that the leakage simulation problem can also be viewed as an equivalent reformulation of the problem of constructing a predictor satisfying *no-access outcome indistinguishability* proposed in Dwork et al. (2021), which they showed is equivalent to *multiaccuracy*. In the most general form, outcome indistinguishability studies a family \mathcal{A} of *distinguishers*, which are functions $A : \mathcal{X} \times \mathcal{O} \times (\Delta\mathcal{O})^{\mathcal{X}} \rightarrow [0, 1]$ that take as input an individual, an outcome, and a predictor and attempts to distinguish genuine outcomes o_i^* from modeled outcomes \tilde{o}_i . The definition of outcome indistinguishability requires that the distinguishing advantage shall be small. Formally,

Definition 7 A predictor $\tilde{p} : \mathcal{X} \rightarrow \Delta\mathcal{O}$ is $(\mathcal{A}, \varepsilon)$ -outcome-indistinguishable if for all $A \in \mathcal{A}$,

$$\left| \mathbf{E}[A(i, \tilde{o}_i, \tilde{p})] - \mathbf{E}[A(i, o_i^*, \tilde{p})] \right| \leq \varepsilon.$$

In the above definition, the distinguisher A has white-box access to the predictor \tilde{p} , which gives the most information of \tilde{p} . By restricting the access of A to the predictor \tilde{p} in different manners, Dwork et al. (2021) obtained a hierarchy of definitions of outcome indistinguishability. Of particular interest to us will be their notions of *no-access* outcome indistinguishability and *sample-access* outcome indistinguishability, which are shown to be equivalent to multiaccuracy and multicalibration respectively. In *no-access* outcome indistinguishability, the distinguisher cannot access \tilde{p} at all, that is, every distinguisher $A \in \mathcal{A}$ takes the form $A(j, o, \tilde{p}) = A'(j, o)$. In *sample-access* outcome indistinguishability, the distinguisher is only provided the output of \tilde{p} on the individual under consideration—equivalently, each distinguisher $A \in \mathcal{A}$ takes the form $A(j, o, \tilde{p}) = A'(j, o, \tilde{p}_j)$ for some function $A' : \mathcal{X} \times \mathcal{O} \times (\Delta\mathcal{O}) \rightarrow [0, 1]$.

We observe that *no access* outcome indistinguishability is equivalent to leakage simulation: (i, o_i^*) is the analogue of (X, O) ; the predictor \tilde{p} is analogous to the simulator h while \tilde{o}_i is analogous to \tilde{O} (sampled according to $h(X)$ and \tilde{p}_i respectively). The goals are also identical: no distinguisher A in the class considered can tell apart (i, o_i^*) from (i, \tilde{o}_i) , or (X, O) from (X, \tilde{O}) . In addition,

algorithms in [Trevisan et al. \(2009\)](#); [Jetchev and Pietrzak \(2014\)](#) for achieving leakage simulation are very similar to the multiaccuracy algorithm in [Hébert-Johnson et al. \(2018\)](#).

Leveraging this equivalence immediately yields the first lower bound for the complexity of no-access outcome indistinguishability predictors relative to \mathcal{A} . This follows from the result in [Chen et al. \(2018\)](#) that, relative to \mathcal{A} , the complexity of a simulator is at least $\Omega(\ell\varepsilon^{-2})$, namely, the simulator makes at least $\Omega(\ell\varepsilon^{-2})$ black-box calls to some distinguishers in \mathcal{A} . Note that the lower bound only holds for simulators that are restricted to black-box use of the distinguishers and satisfy a restriction that, when invoked on input X , they only make black-box calls to the distinguishers on the same input X . All the leakage simulation, multiaccuracy, and multicalibration algorithms in the literature satisfy this restriction. Therefore, we conclude that predictors satisfying no-access outcome indistinguishability (under same constraints) also have relative complexity $\Omega(\ell\varepsilon^{-2})$ w.r.t. distinguishers in \mathcal{A} . The equivalence between no-access outcome indistinguishability and multiaccuracy further tells us that the same relative complexity lower bound holds for multiaccurate predictors w.r.t. \mathcal{C} (the analogue of \mathcal{A}). Finally, since multicalibration is stronger, the same lower bound extends to multicalibrated predictors.

This forecloses the existence of predictors that are smaller than the distinguisher yet fools them all (subject to the above restriction). In other words, there is no predictor analogous to a pseudo-random generator that fools all polynomial-time tests.

Beyond the lower bound, the connection between no-access outcome indistinguishability and leakage simulation suggests two additional directions. First, the work of [Chen et al. \(2018\)](#) presented a leakage simulation algorithm via no-regret learning. Inspired by their algorithm, we present in Section 5 a general algorithmic framework for achieving sample-access outcome indistinguishability, equivalently multi-calibration, also via no-regret learning. Our framework unifies algorithms in prior works. Second, inspired by the connection between leakage simulation and the hardcore lemma for Boolean functions, we ask whether the stronger notion of multicalibration yields stronger consequences. Indeed, we present in Section 7 a multicalibration-based proof of a hardcore lemma for real-valued functions.

5. Sample Complexity and No-Regret Learning

Various notions of multi-group fairness, including multiaccuracy, multicalibration, strict multicalibration (Definitions 1, 2, and 3), and low-degree multicalibration, a hierarchy of tractable relaxations of the standard multicalibration definition introduced in [Gopalan et al. \(2022b\)](#), are implied by the notion of *outcome indistinguishability* of [Dwork et al. \(2021\)](#) with respect to different classes of adversaries. Thus, to achieve these notions it suffices to design algorithms for *outcome indistinguishability*. Here, our contributions are twofold. First, we present an algorithmic template that unifies prior algorithms for achieving outcome indistinguishability through the lens of no-regret learning (see Appendix F). The template in Algorithm 2 can be viewed as running $|\mathcal{X}|$ instances of a no-regret algorithm in parallel, where each instance corresponds to one member j of the population \mathcal{X} , and the distribution \tilde{p}_j corresponds to a mixed strategy over $o \in \mathcal{O}$. The predicted probabilities $\tilde{p}_j(o)$ are refined over multiple rounds. The loss function, however, is chosen based on the entire predictor \tilde{p} at the current round; see Appendix F for details. Second, we show that Algorithm 1, an instantiation of our algorithmic template, yields an improved upper bound on sample complexity for achieving multicalibration in the multiclass setting and for low-degree multicalibration.

Algorithm 1: Multiclass Outcome Indistinguishability via Multiplicative Weight Updates

Data: distinguisher family \mathcal{A} , parameter ε , iteration count t , samples per iteration n

Result: predictor $\tilde{p} : \mathcal{X} \rightarrow \Delta\mathcal{O}$

```

 $\tilde{p}_j^{(0)}(o) \leftarrow 1/|\mathcal{O}|$  for all  $j \in \mathcal{X}$  and  $o \in \mathcal{O}$            // initialize  $\tilde{p}^{(0)}$  to constant
for  $s = 0, 1, \dots, t - 1$  do
     $(i_{s1}, o_{i_{s1}}^*), \dots, (i_{sn}, o_{i_{sn}}^*) \sim (i, o_i^*)$            // draw  $n$  fresh iid samples
    changed  $\leftarrow$  false
    for  $A \in \mathcal{A}$  do
         $\alpha \leftarrow \sum_{m=1}^n A(i_{sm}, o_{i_{sm}}^*, \tilde{p}^{(s)})$            // genuine outcomes
         $\beta \leftarrow \sum_{m=1}^n \sum_{o \in \mathcal{O}} \tilde{p}_{i_{sm}}(o) A(i_{sm}, o, \tilde{p}^{(s)})$            // modeled outcomes
        if not changed and  $|\beta - \alpha| > \varepsilon n / 2$  then           // is  $A$ 's advantage large?
             $\gamma \leftarrow \text{sign}(\beta - \alpha) \cdot \varepsilon / 3$            // step size
            for  $j \in \mathcal{X}$  do
                for  $o \in \mathcal{O}$  do
                     $f_j(o) \leftarrow \tilde{p}_j^{(s)}(o) \exp(-\gamma A(j, o, \tilde{p}^{(s)}))$            // update weights
                end
                 $\tilde{p}_j^{(s+1)}(o) \leftarrow f_j(o) / \sum_{o' \in \mathcal{O}} f_j(o')$ 
                changed  $\leftarrow$  true
            end
        end
    end
    if not changed then return  $\tilde{p}^{(s)}$ 
end
return  $\perp$ 
    
```

Lemma 8 Running Algorithm 1 with appropriately chosen $t \lesssim \log(|\mathcal{O}|)/\varepsilon^2$ and $n \lesssim \log(|\mathcal{A}|t)/\varepsilon^2$ yields an $(\mathcal{A}, \varepsilon)$ -outcome-indistinguishable predictor with probability at least 99%.⁵

By choosing the distinguisher family \mathcal{A} judiciously, we can achieve multicalibration in a more sample-efficient manner than existing algorithms. In fact, the construction of the family \mathcal{A} follows naturally from our statistical distance-based definition of multicalibration:

Definition 9 Let $\mathcal{A}_{\mathcal{C}, \mathcal{G}}^{\text{MC}} = \{A_{c,E} \mid c \in \mathcal{C}, E \subseteq \mathcal{Y} \times \mathcal{O} \times \mathcal{G}\}$, where

$$A_{c,E}(j, o, \hat{p}) = \mathbf{1}[(c_j, o, \hat{p}_j) \in E]$$

for each member $j \in \mathcal{X}$, possible outcome $o \in \mathcal{O}$, and predictor $\hat{p} : \mathcal{X} \rightarrow \Delta\mathcal{O}$. (Recall from Section 2 that \mathcal{C} is a collection of functions $c : \mathcal{X} \rightarrow \mathcal{Y}$ and that \hat{p} denotes the discretization of \tilde{p} to a “grid” \mathcal{G} , which we assume to be an η -covering of $\Delta\mathcal{O}$.)

Theorem 10 Running Algorithm 1 on $\mathcal{A} = \mathcal{A}_{\mathcal{C}, \mathcal{G}}^{\text{MC}}$ and appropriately chosen \mathcal{G}, t, n yields a predictor \tilde{p} such that \hat{p} is $(\mathcal{C}, \varepsilon)$ -multicalibrated with probability at least 99%. The algorithm samples at

5. By $f(x) \lesssim g(x)$, we mean that $f(x) = O(g(x))$.

most

$$tn \lesssim \left(\log(|\mathcal{C}|) + |\mathcal{Y}||\mathcal{O}| \left(\frac{4}{\varepsilon} \right)^{|\mathcal{O}|-1} \right) \cdot \frac{\log(|\mathcal{O}|)}{\varepsilon^4}$$

i.i.d. individual-outcome pairs.

Letting $\mathcal{Y} = \{0, 1\}$ in Theorem 10, so that $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$, we recover the first sample complexity upper bound that we stated in Section 1. Intuitively, our savings compared to prior works comes from the fact that our algorithm *directly* targets the milder requirements of ordinary multicalibration, while prior algorithms typically “go through” strict multicalibration by aiming for stringent per-level-set guarantees.

To justify the second upper bound we stated in Section 1, we now turn our attention to the notion of *low-degree* multicalibration from Gopalan et al. (2022b). A rephrased statement of the definition is as follows:

Definition 11 *A function $f : [0, 1]^\ell \rightarrow [0, 1]$ is a monomial of degree less than k if it takes the form $f(v) = v_{t_1} \cdots v_{t_j}$ for some $j < k$ indices $t_1, \dots, t_j \in [\ell]$. For $\mathcal{O} = \{0, 1\}$ and $\mathcal{C} \subseteq [0, 1]^{\mathcal{X}}$ and $k \in \mathbb{N}$, let $\mathcal{A}_{\mathcal{C},k}^{\text{MC}}$ to be the family of all distinguishers of the form*

$$A_{c,f,o'}(j, o, \tilde{p}) = c(j)f(\tilde{p}_j)\mathbf{1}[o = o'],$$

where $c \in \mathcal{C}$ and $o' \in \mathcal{O}$ and f is a monomial of degree less than k . We say that a predictor \tilde{p} is $(\mathcal{C}, \varepsilon)$ -degree- k multicalibrated if \tilde{p} is $(\mathcal{A}_{\mathcal{C},k}^{\text{MC}}, \varepsilon)$ -outcome indistinguishable.

Using the fact that the family $\mathcal{A}_{\mathcal{C},k}^{\text{MC}}$ is a subset of $\mathcal{A}_{\mathcal{C},\mathcal{G}}^{\text{MC}}$, one can show that degree- k multicalibration is *weaker* than the notion of multicalibration considered so far. With this in mind, it should not be surprising that running Algorithm 1 on $\mathcal{A}_{\mathcal{C},k}^{\text{MC}}$ instead of $\mathcal{A}_{\mathcal{C},\mathcal{G}}^{\text{MC}}$ immediately gives us the following tighter upper bound on the samples needed for degree- k multicalibration:

Theorem 12 *Running Algorithm 1 on $\mathcal{A} = \mathcal{A}_{\mathcal{C},k}^{\text{MC}}$ and appropriately chosen t, n yields a $(\mathcal{C}, \varepsilon)$ -degree- k multicalibrated predictor with probability at least 99% and samples at most*

$$tn \lesssim \left(\log(|\mathcal{C}|) + k \log \left(\frac{\ell}{k} \right) + \log \left(\frac{\ell}{\varepsilon} \right) \right) \cdot \frac{\log(\ell)}{\varepsilon^4}$$

i.i.d. individual-outcome pairs, where $\ell = |\mathcal{O}|$ and $\ell \geq k$.

In Appendix F.4, we prove the theorems of this section, and also give an upper bound on the sample complexity of *strict* multicalibration. The proofs will show, in particular, that the improvement in our Theorem 12 compared to Theorem 35 of Gopalan et al. (2022b) comes primarily from our deliberate use of *multiplicative* updates to \tilde{p} in Algorithm 1, as opposed to additive updates.

6. Graph Regularity as Structured Multigroup Fairness

Szemerédi's regularity lemma Szemerédi (1975) is a cornerstone result in extremal graph theory with a wide range of applications in combinatorics, number theory, computational complexity theory, and other areas of mathematics. Roughly speaking, it states that any large, dense graph can

be decomposed into parts that behave “pseudorandomly” in a certain precise sense. The *Frieze-Kannan weak regularity lemma* [Frieze and Kannan \(1996\)](#) is a related result in graph theory with a qualitatively weaker conclusion, but parameter dependencies much better suited for algorithmic applications.

The goal of this section is to show that regularity partitions of a graph correspond to predictors satisfying multi-group fairness and an additional structural condition on their level sets. In [Section 6.1](#), we state various definitions of graph regularity. In [Section 6.2](#), we state and prove the correspondence, which is the key result of this section.

6.1. Definitions of Graph Regularity

Let $G = (V, E)$ be a graph, by which we mean that V is a finite set and $E \subseteq V \times V$. For vertex subsets $S \subseteq V$ and $T \subseteq V$, let $e_G(S, T) = |(S \times T) \cap E|$ count the number of edges from S to T , and let $d_G(S, T) = e_G(S, T)/|S||T|$ denote the *density* of edges from S to T . When the graph G is clear from context, we will omit the subscript G from e_G and d_G .

To state Szemerédi’s regularity lemma, we must first recall the notion of an ε -regular pair:

Definition 13 *Let $X, Y \subseteq V$. We say that the pair (X, Y) is ε -regular if*

$$|d(S, T) - d(X, Y)| \leq \varepsilon$$

for all $S \subseteq X$ and $T \subseteq Y$ such that $|S| \geq \varepsilon|X|$ and $|T| \geq \varepsilon|Y|$.

Intuitively, a pair (X, Y) is ε -regular if edges from X to Y are distributed in a “pseudorandom” fashion. The Szemerédi regularity lemma finds a partition \mathcal{P} of the vertices of V such that most pairs of parts are ε -regular, in the following sense:

Definition 14 *A partition $\mathcal{P} = \{V_1, \dots, V_m\}$ of V satisfies Szemerédi ε -regularity if*

$$\sum_{\substack{j, k \in [m] \\ (V_j, V_k) \text{ not } \varepsilon\text{-regular}}} |V_j||V_k| \leq \varepsilon|V|^2.$$

In contrast to Szemerédi regularity, which gives fine-grained “local” regularity guarantees on the pairs of regular parts, the weaker regularity condition of [Frieze and Kannan \(1996\)](#) gives only a coarse “global” regularity guarantee:

Definition 15 *A partition $\mathcal{P} = \{V_1, \dots, V_m\}$ of V satisfies Frieze-Kannan ε -regularity if for all $S, T \subseteq V$,*

$$\left| e(S, T) - \sum_{j, k \in [m]} d(V_j, V_k) |S \cap V_j| |T \cap V_k| \right| \leq \varepsilon|V|^2.$$

Intermediate Regularity We will soon show that for a certain instantiation of the multi-group fairness framework, Szemerédi regularity corresponds to strict multicalibration, and Frieze-Kannan regularity corresponds to multiaccuracy. Inspired by this connection, we will also show that (ordinary) multicalibration corresponds to an *intermediate* notion of graph regularity that has, to our knowledge, not appeared in the prior literature:




Figure 2: A graph G , a regularity partition of G , and a multicalibrated edge predictor for G .

Definition 16 Let $X, Y, S, T \subseteq V$. We say that the pair (X, Y) is (S, T, ε) -regular if

$$|d(S \cap X, T \cap Y) - d(X, Y)| \leq \varepsilon.$$

Definition 17 A partition $\mathcal{P} = \{V_1, \dots, V_m\}$ of V satisfies intermediate ε -regularity if for all $S, T \subseteq V$,

$$\sum_{\substack{j, k \in [m] \\ (V_j, V_k) \text{ not } (S, T, \varepsilon)\text{-regular}}} |S \cap V_j| |T \cap V_k| \leq \varepsilon |V|^2.$$

We chose the name *intermediate regularity* to emphasize that it is a strictly stronger notion than Frieze-Kannan weak regularity, but still strictly weaker than Szemerédi regularity. In Appendix G, we prove these claimed relationships. In Appendix H, we present an algorithm for achieving intermediate regularity.

6.2. The Regularity-Multicalibration Theorem

The Szemerédi and Frieze-Kannan regularity lemmas state that any graph G has a partition satisfying ε -regularity (of the appropriate kind) whose number of parts is bounded by a function of ε , *independent of the size of the graph*. The partitions constructed in the proofs of these lemmas can be viewed as low-complexity approximations to the graph that fool a particular family of cryptographic distinguishers, as observed by Skorski (2017).

In this section, we show that these distinguishers fit neatly within the framework of multi-group fairness. To begin, we will recall that Frieze-Kannan weak regularity corresponds naturally to multiaccuracy, a result established in Trevisan et al. (2009). We then prove that Szemerédi regularity corresponds naturally to *strict* multicalibration with respect to the *same* collection of subpopulations. Taking this connection one step further, we show that (ordinary) multicalibration naturally gives rise to our new notion of *intermediate regularity*, which is stronger than Frieze-Kannan regularity but weaker than Szemerédi regularity.

Definition 18 In the edge prediction problem for a graph $G = (V, E)$, the population is $\mathcal{X} = V \times V$, each individual $i \in V \times V$ is a vertex pair drawn uniformly at random, and the true outcome of individual i is the single bit $o_i^* = \mathbf{1}[i \in E]$. The collection of protected subpopulations is $\mathcal{C} = \{\mathbf{1}_{S \times T} \mid S, T \subseteq V\}$. In this setting, we call $\tilde{p} : \mathcal{X} \rightarrow [0, 1]$ an edge predictor for G .

Theorem 19 (Regularity-Multicalibration) *Given a graph G , consider the following definitions of fairness for an edge predictor \tilde{p} for G and regularity for a vertex partition \mathcal{P} of G :*

- (i) *multiaccuracy and Frieze-Kannan regularity.*
- (ii) *multicalibration and intermediate regularity.*
- (iii) *strict multicalibration and Szemerédi regularity.*

For each such pair of definitions, there exists an absolute constant $0 < c < 1$ such that the following two implications hold for sufficiently small ε :

- (a) *If \mathcal{P} is ε -regular, the predictor \tilde{p} that outputs $d(V_j, V_k)$ on all of $V_j \times V_k \in \mathcal{P}^2$ is $(\mathcal{C}, \varepsilon^c)$ -fair.⁶*
- (b) *If \tilde{p} is $(\mathcal{C}, \varepsilon)$ -fair and the set of level sets of \tilde{p} is \mathcal{P}^2 for some partition \mathcal{P} , then \mathcal{P} is ε^c -regular.*

Proof Sketch For $V_j, V_k \in \mathcal{P}$, define $d_{jk} = d(V_j, V_k)$. In Appendix G, we will show through algebraic manipulations that the regularity criteria of the previous section are equivalent, up to a polynomial change in ε , to the following conditions:

$$\begin{aligned} \max_{S, T \subseteq V} \left| \sum_{j=1}^m \sum_{k=1}^m e(S \cap V_j, T \cap V_k) - d_{jk} |S \cap V_j| |T \cap V_k| \right| &\leq \varepsilon |V|^2 && \text{for Frieze-Kannan } \varepsilon\text{-regularity,} \\ \max_{S, T \subseteq V} \sum_{j=1}^m \sum_{k=1}^m |e(S \cap V_j, T \cap V_k) - d_{jk} |S \cap V_j| |T \cap V_k|| &\leq \varepsilon |V|^2 && \text{for intermediate } \varepsilon\text{-regularity,} \\ \sum_{j=1}^m \sum_{k=1}^m \max_{S, T \subseteq V} |e(S \cap V_j, T \cap V_k) - d_{jk} |S \cap V_j| |T \cap V_k|| &\leq \varepsilon |V|^2 && \text{for Szemerédi } \varepsilon\text{-regularity.} \end{aligned}$$

The relationships among our fairness criteria can be phrased similarly. Indeed, let $\Delta_{S,v}(\tilde{p}) = \mathbf{Pr}[i \in S, \tilde{o}_i = 1, \tilde{p}_i = v] - \mathbf{Pr}[i \in S, o_i^* = 1, \tilde{p}_i = v]$ for a predictor $\tilde{p} : \mathcal{X} \rightarrow [0, 1]$, a subpopulation $S \subseteq \mathcal{X}$ and a value $v \in [0, 1]$. Then, the requirements of Definitions 1, 2, and 3 reduce to:

$$\begin{aligned} \max_S \left| \sum_v \Delta_{S,v}(\tilde{p}) \right| &\leq \varepsilon && \text{for } (\mathcal{C}, \varepsilon)\text{-multiaccuracy,} \\ \max_S \sum_v |\Delta_{S,v}(\tilde{p})| &\leq \varepsilon && \text{for } (\mathcal{C}, \varepsilon)\text{-multicalibration,} \\ \sum_v \max_S |\Delta_{S,v}(\tilde{p})| &\leq \varepsilon && \text{for strict } (\mathcal{C}, \varepsilon)\text{-multicalibration,} \end{aligned}$$

where the maxima are taken over S such that $\mathbf{1}_S \in \mathcal{C}$ or $\mathbf{1}_{\mathcal{X} \setminus S} \in \mathcal{C}$ or $S = \mathcal{X}$ and the sums are taken over the range of \tilde{p} . Comparing the two displayed sets of inequalities will yield part (a) of the theorem. Part (b) will follow from a similar argument. For a full proof, see Appendix I.

6. For a partition $\mathcal{P} = \{V_1, \dots, V_m\}$ of the vertices V , the set $\mathcal{P}^2 = \{V_j \times V_k \mid j, k \in [m]\}$ denotes the partition of $V \times V$ obtained from all pairwise Cartesian products of parts of \mathcal{P} .

7. Hardcore Lemma for Real-Valued Functions

The leakage simulation lemma is connected with the hard-core lemma for deterministic Boolean functions as shown by Trevisan et al. (2009); Vadhan and Zheng (2013). We show that multicalibration enables stronger consequences, namely a hard-core lemma for real valued functions.

Informally, Impagliazzo’s hard-core lemma says that for any boolean function $f : \mathcal{X} \rightarrow \{0, 1\}$ that is hard on average against a class $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$, there is a large subset of \mathcal{X} , the *hard core*, whose size depends on the hardness of f , on which f is effectively pseudorandom. We prove an analogous statement for real valued functions (equivalently, random functions with boolean outcomes) $p^* : \mathcal{X} \rightarrow [0, 1]$, where hardness of being correct is replaced by hardness of approximation (in L_2 distance), while the hardcore/pseudorandom condition is replaced with a covariance condition. In other words, if it’s hard to approximate the function, then there’s some large set on which it’s hard to even have any non-negligible covariance.

In order to formally state the theorem, let $R_T(\mathcal{C})$ denote the class of functions of relative complexity T with respect to \mathcal{C} . Let a \mathcal{C} gate denote black-box computation of a function $c \in \mathcal{C}$. More precisely, every $\check{p} \in R_T(\mathcal{C})$ is a circuit containing T \mathcal{C} -gates and $\text{poly}(T)$ additional basic operations (e.g., floating-point arithmetic and Boolean logical operations). In particular, this captures the class of all functions which can be the outcome of the multicalibration algorithm after T rounds with an appropriate ε .

Theorem 20 (Hardcore lemma for probabilities) *Let $\alpha > 0, \gamma \in (0, \frac{1}{3}]$. Let $\mathcal{C} \subseteq [0, 1]^{\mathcal{X}}$ be a collection of real valued functions. Let $p^* : \mathcal{X} \rightarrow [0, 1]$ be a function which is hard to α -approximate by functions in $R_T(\mathcal{C})$ for $T = O(\alpha^{-6}\gamma^{-4})$, i.e., $\forall \check{p} \in R_T(\mathcal{C}), \mathbf{E}[(p_i^* - \check{p}_i)^2] > \alpha$.*

Then, there exists a hardcore set $S \subseteq \mathcal{X}$ where $\Pr[i \in S] > \frac{\alpha^2\gamma}{4}$ and $\forall c \in \mathcal{C}, \mathbf{Cov}(c_i, p_i^ \mid i \in S) < \gamma \mathbf{Var}(p_i^* \mid i \in S)$.*

The proof of this statement is based on the following intuition: First, the multicalibration algorithm gives us a relatively simple predictor \check{p} . Next, a multicalibrated predictor partitions \mathcal{X} into slices on which either (1) p^* has low variance (so \check{p} is highly accurate), or (2) \check{p} is not highly accurate (so p^* is high variance) but \mathcal{C} is not able to take advantage of this (nothing in \mathcal{C} is correlated with p^* on this slice). Finally, if \check{p} is far from p^* , then there must exist a set on which the latter condition is true. That set is the hardcore. For more detail, we refer the reader to Appendix J.

Acknowledgments

This work was supported in part by the Simons Foundation Grant 733782 and the Sloan Foundation Grant G-2020-13941. H. Lin was also supported by NSF grants CNS-1936825 (CAREER), CNS-2026774, and a Cisco Research Award.

References

- Noga Alon and Eyal Lubetzky. Codes and xor graph products. *Combinatorica*, 27:13–33, 2007.
- Noga Alon and Assaf Naor. Approximating the cut-norm via grothendieck’s inequality. In *ACM Symposium on Theory of Computing*, page 72–80, 2004.

- Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning*, 8(3-4):231–357, 2015.
- Sílvia Casacuberta, Cynthia Dwork, and Salil Vadhan. Complexity-theoretic implications of multicalibration: Hardcore sets, dense models, and pseudoentropy. *Unpublished manuscript*, 2023. See Casacuberta, Finding Simple Models of Complex Objects: From Regularity Lemmas to Algorithmic Fairness, undergraduate thesis, Harvard University, and <https://simons.berkeley.edu/workshops/multigroup-fairness-validity-statistical-judgment>.
- Yi-Hsiu Chen, Kai-Min Chung, and Jyun-Jie Liao. On the complexity of simulating auxiliary input. In *EUROCRYPT: Annual International Conference on the Theory and Applications of Cryptographic Techniques*, volume 3, pages 371–390, 2018.
- Kai-Min Chung, Edward Lui, and Rafael Pass. 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 12*, pages 66–92. Springer, 2015.
- Zhun Deng, Cynthia Dwork, and Linjun Zhang. HappyMap : A Generalized Multicalibration Method. In *Innovations in Theoretical Computer Science*, volume 251, pages 41:1–41:23, 2023.
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012.
- Cynthia Dwork, Michael P. Kim, Omer Reingold, Guy N. Rothblum, and Gal Yona. Outcome indistinguishability. In *ACM Symposium on Theory of Computing*, page 1095–1108, 2021.
- Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *IEEE Symposium on Foundations of Computer Science*, pages 293–302, 2008.
- Jacob Fox and László Miklós Lovász. A tight lower bound for szemerédi’s regularity lemma. *Combinatorica*, 37:911–951, 2017.
- Alan M. Frieze and Ravi Kannan. The regularity lemma and approximation schemes for dense problems. In *IEEE Symposium on Foundations of Computer Science*, pages 12–20, 1996.
- Parikshit Gopalan, Adam Tauman Kalai, Omer Reingold, Vatsal Sharan, and Udi Wieder. Omnipredictors. In *Innovations in Theoretical Computer Science*, pages 79:1–79:21, 2022a.
- Parikshit Gopalan, Michael P. Kim, Mihir Singhal, and Shengjia Zhao. Low-degree multicalibration. In *Conference on Learning Theory*, volume 178, pages 3193–3234, 2022b.
- Parikshit Gopalan, Lunjia Hu, Michael P. Kim, Omer Reingold, and Udi Wieder. Loss minimization through the lens of outcome indistinguishability. In *Innovations in Theoretical Computer Science*, 2023a.
- Parikshit Gopalan, Michael P. Kim, and Omer Reingold. Characterizing notions of omniprediction via multicalibration, 2023b. under submission.

- Varun Gupta, Christopher Jung, Georgy Noarov, Mallesh M. Pai, and Aaron Roth. Online multivalid learning: Means, moments, and prediction intervals. In *Innovations in Theoretical Computer Science*, volume 215, pages 82:1–82:24, 2022.
- Nika Haghtalab, Michael I. Jordan, and Eric Zhao. A unifying perspective on multi-calibration: Unleashing game dynamics for multi-objective learning, 2023.
- Úrsula Hébert-Johnson, Michael P. Kim, Omer Reingold, and Guy N. Rothblum. Multicalibration: Calibration for the (computationally-identifiable) masses. In *International Conference on Machine Learning*, pages 1939–1948, 2018.
- Russell Impagliazzo. Hard-core distributions for somewhat hard problems. In *IEEE Symposium on Foundations of Computer Science*, pages 538–545, 1995.
- Russell Impagliazzo. Algorithmic dense model theorems and weak regularity, 2009. Private communication.
- Dimitar Jetchev and Krzysztof Pietrzak. How to fake auxiliary input. In *Theory of Cryptography Conference*, pages 566–590, 2014.
- Christopher Jung, Changhwa Lee, Mallesh M. Pai, Aaron Roth, and Rakesh Vohra. Moment multicalibration for uncertainty estimation. In *Conference on Learning Theory*, volume 134, pages 2634–2678, 2021.
- Faisal Kamiran and Toon Calders. Classifying without discriminating. In *2009 2nd international conference on computer, control and communication*, pages 1–6. IEEE, 2009.
- Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International Conference on Machine Learning*, pages 2564–2572, 2018.
- Michael P. Kim, Amirata Ghorbani, and James Zou. Multiaccuracy: Black-box post-processing for fairness in classification. In *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*, page 247–254, 2019.
- Michael P. Kim, Christoph Kern, Shafi Goldwasser, Frauke Kreuter, and Omer Reingold. Universal adaptability: Target-independent inference that competes with propensity scoring. *Proceedings of the National Academy of Sciences*, 119(4):e2108097119, 2022.
- Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807*, 2016.
- Arkady Nemirovsky and David Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley, 1983.
- Omer Reingold, Luca Trevisan, Madhur Tulsiani, and Salil Vadhan. Dense subsets of pseudorandom sets. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 76–85. IEEE, 2008.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning – From Theory to Algorithms*. Cambridge University Press, 2014.

- Maciej Skorski. A cryptographic view of regularity lemmas: Simpler unified proofs and refined bounds. In *Theory and Applications of Models of Computation*, page 586–599, 2017.
- Endre Szemerédi. Regular partitions of graphs. Technical report, Stanford University, Stanford, CA, USA, 1975.
- Luca Trevisan, Madhur Tulsiani, and Salil P. Vadhan. Regularity, boosting, and efficiently simulating every high-entropy distribution. In *Computational Complexity Conference*, pages 126–136, 2009.
- Dino Pedreschi Salvatore Ruggieri Franco Turini. Discrimination-aware data mining. KDD, 2008.
- Salil Vadhan and Colin Jia Zheng. Characterizing pseudoentropy and simplifying pseudorandom generator constructions. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 817–836, 2012.
- Salil P. Vadhan and Colin Jia Zheng. A uniform min-max theorem with applications in cryptography. In *CRYPTO: Annual International Cryptology Conference*, volume 1, pages 93–110, 2013.
- Yufei Zhao. Graph theory and additive combinatorics, 2022.
- Jia Zheng. A uniform min-max theorem and characterizations of computational randomness, 2014. PhD thesis, Harvard University.

Appendix A. Relationships Among Definitions

First, we show that strict $(\mathcal{C}, \varepsilon)$ -multicalibration implies $(\mathcal{C}, \varepsilon)$ -multicalibration, which in turn implies $(\mathcal{C}, \varepsilon)$ -multiaccuracy:

Theorem 21 *If \tilde{p} is strictly $(\mathcal{C}, \varepsilon)$ -multicalibrated, then \tilde{p} is $(\mathcal{C}, \varepsilon)$ -multicalibrated.*

Proof If \tilde{p} is strictly $(\mathcal{C}, \varepsilon)$ -multicalibrated, then for all $c \in \mathcal{C}$,

$$\delta((c_i, \tilde{o}_i, \tilde{p}_i), (c_i, o_i^*, \tilde{p}_i)) = \mathbf{E} [\delta((c_i, \tilde{o}_i), (c_i, o_i^*) \mid \tilde{p}_i)] \leq \mathbf{E} \left[\max_{c' \in \mathcal{C}} \delta((c'_i, \tilde{o}_i), (c'_i, o_i^*) \mid \tilde{p}_i) \right] \leq \varepsilon,$$

so \tilde{p} is $(\mathcal{C}, \varepsilon)$ -multicalibrated, as well. ■

Theorem 22 *If \tilde{p} is $(\mathcal{C}, \varepsilon)$ -multicalibrated, then \tilde{p} is $(\mathcal{C}, \varepsilon)$ -multiaccurate.*

Proof If \tilde{p} is $(\mathcal{C}, \varepsilon)$ -multicalibrated, then for all $c \in \mathcal{C}$,

$$\delta((c_i, \tilde{o}_i), (c_i, o_i^*)) \leq \delta((c_i, \tilde{o}_i, \tilde{p}_i), (c_i, o_i^*, \tilde{p}_i)) \leq \varepsilon,$$

so \tilde{p} is $(\mathcal{C}, \varepsilon)$ -multiaccurate, as well. ■

Appendix B. Proof of Theorem 4

- (a) Let $\mathcal{X} = \{0, 1\}$ and $\mathcal{C} = \{\mathbf{1}_{\mathcal{X}}\}$. Consider an individual i drawn uniformly from \mathcal{X} whose outcome o_i^* is conditionally distributed as $\text{Ber}(1/2)$ given i . Then the predictor $\tilde{p}_j = j$ is $(\mathcal{C}, 0)$ -multiaccurate but not $(\mathcal{C}, 1/2 - \alpha)$ -multicalibrated for any $\alpha > 0$.
- (b) Let $\mathcal{X} = [m] \times [m]$ for some positive integer m , and let $\mathcal{C} = \{c_k : k \in [m]\}$ where

$$c_k(j) = \mathbf{1}[j_1 = k \text{ and } j_2 \leq k]$$

for each member $j = (j_1, j_2)$ of the population. Consider an individual $i = (i_1, i_2)$ drawn uniformly from \mathcal{X} whose outcome is $o_i^* = \mathbf{1}[i_1 \geq i_2]$, and let $\tilde{p}_j = j_1/m$. The range of \tilde{p} is $\{1/m, 2/m, \dots, 1\}$. A simple calculation shows that for each function $c_k \in \mathcal{C}$ and each value v in the range of \tilde{p} , we have that $\Pr[\tilde{p}_i = v] = 1/m$ and

$$\delta((c_{ki}, \tilde{o}_i), (c_{ki}, o_i^*) \mid \tilde{p}_i = v) = \begin{cases} 2v(1-v) & \text{if } v = k/m \\ 0 & \text{otherwise} \end{cases}$$

(recall from Section 2 that $c_{ki} = c_k(i)$). Therefore, as $m \rightarrow \infty$,

$$\max_{c_k \in \mathcal{C}} \delta((c_{ki}, \tilde{o}_i), (c_{ki}, o_i^*) \mid \tilde{p}_i) = \max_{k \in [m]} \frac{2(k/m)(1-k/m)}{m} \rightarrow 0$$

but

$$\mathbf{E} \left[\max_{c_k \in \mathcal{C}} \delta((c_{ki}, \tilde{o}_i), (c_{ki}, o_i^*) \mid \tilde{p}_i) \right] = \sum_{k=1}^m \frac{2(k/m)(1-k/m)}{m} \rightarrow \frac{1}{3},$$

so \tilde{p} is $(\mathcal{C}, \varepsilon)$ -multicalibrated but not strictly $(\mathcal{C}, 1/3 - \alpha)$ -multicalibrated for any $\alpha > 0$.

Appendix C. Rounding Predictions

In this section, we prove Theorem 5 and Corollary 6. In what follows, let \tilde{p} be a $(\mathcal{C}, \varepsilon)$ -multicalibrated predictor, and recall from Section 2 that \hat{p} denotes the discretization of \tilde{p} to a finite set $\mathcal{G} \subseteq \Delta\mathcal{O}$. If \mathcal{G} is an η -covering of $\Delta\mathcal{O}$, then the inequality $\delta(\hat{o}_i, \tilde{o}_i \mid i) \leq \eta$ holds almost surely, so

$$\mathbf{E} \left[\max_{c \in \mathcal{C}} \delta((c_i, \hat{o}_i), (c_i, o_i^*) \mid \hat{p}_i) \right] \leq \mathbf{E} \left[\max_{c \in \mathcal{C}} \delta((c_i, \tilde{o}_i), (c_i, o_i^*) \mid \hat{p}_i) \right] + \eta$$

and the expectation on the right hand side is precisely

$$\sum_{v \in \mathcal{G}} \max_{c \in \mathcal{C}} \delta((c_i, \tilde{o}_i), (c_i, o_i^*) \mid \hat{p}_i = v) \Pr[\hat{p}_i = v].$$

Each of $|\mathcal{G}|$ terms in the sum can be bounded as follows:

$$\begin{aligned} \delta((c_i, \tilde{o}_i), (c_i, o_i^*) \mid \hat{p}_i = v) \Pr[\hat{p}_i = v] &\leq \delta((c_i, \tilde{o}_i, \hat{p}_i), (c_i, o_i^*, \hat{p}_i)) \\ &\leq \delta((c_i, \tilde{o}_i, \tilde{p}_i), (c_i, o_i^*, \tilde{p}_i)) && \text{since } \hat{p}_i \text{ is a function of } \tilde{p}_i, \\ &\leq \varepsilon && \text{since } \tilde{p} \text{ is } (\mathcal{C}, \varepsilon)\text{-multicalibrated.} \end{aligned}$$

At this point, we have shown that \hat{p} is strictly $(\mathcal{C}, |\mathcal{G}|\varepsilon + \eta)$ -multicalibrated, which proves Theorem 5. To prove Corollary 6, we use the following lemma, which gives us a grid \mathcal{G} such that $|\mathcal{G}|\varepsilon^\ell + \eta < 4\varepsilon$ when $\eta = 3\varepsilon$ is sufficiently small.

Lemma 23 For all sufficiently small $\eta > 0$, the grid $\mathcal{G} = \Delta\mathcal{O} \cap \{0, \frac{1}{m}, \frac{2}{m}, \dots, \frac{m-1}{m}, 1\}^\ell$ with $m = \lceil (\ell - 1)/\eta \rceil$ is an η -covering of $\Delta\mathcal{O}$ of size $|\mathcal{G}| < (3/\eta)^{\ell-1}$.

Proof Given any $f \in \Delta\mathcal{O}$, we can find a grid point $g \in \mathcal{G}$ such that f and g differ by at most $1/m$ in all but one coordinate. Since $(\ell - 1)/m \leq \eta$, this means that \mathcal{G} is an η -covering. A counting argument shows that the size of \mathcal{G} is

$$|\mathcal{G}| = \binom{m + \ell - 1}{\ell - 1} \leq \left(\frac{e(m + \ell - 1)}{\ell - 1} \right)^{\ell-1} \leq \left(e \left(\frac{1}{\eta} + 2 \right) \right)^{\ell-1} < \left(\frac{3}{\eta} \right)^{\ell-1}$$

for all $\ell \geq 2$ and all sufficiently small $\eta > 0$. ■

Appendix D. Relationships to Prior Definitions

In this section, we explain the relationships of our new definitions to existing notions of multiaccuracy and multicalibration that appear in the algorithmic fairness literature. We emphasize that strict multicalibration has not been explicitly defined in prior works. Nevertheless, their algorithms actually achieve this stronger notion.

The original definitions of multiaccuracy and multicalibration from the algorithmic fairness literature roughly correspond to our notions of the same names when we restrict attention to binary outcomes and $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$. To facilitate the comparison, we state the following two definitions:

Definition 24 Assume $\mathcal{O} = \{0, 1\}$ and $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$. We say a predictor \tilde{p} is conditionally $(\mathcal{C}, \varepsilon)$ -multiaccurate if

$$\begin{aligned} & \text{for all } \mathbf{1}_S \in \mathcal{C} \text{ such that } \Pr[i \in S] \geq \varepsilon, \\ & \left| \Pr[o_i^* = 1 \mid i \in S] - \Pr[\tilde{p}_i = 1 \mid i \in S] \right| \leq \varepsilon. \end{aligned}$$

Definition 25 Assume $\mathcal{O} = \{0, 1\}$ and $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$. We say a predictor \tilde{p} is conditionally $(\mathcal{C}, \varepsilon)$ -multicalibrated if

$$\begin{aligned} & \text{for all } \mathbf{1}_S \in \mathcal{C} \text{ such that } \Pr[i \in S] \geq \varepsilon, \\ & \text{there exists } S' \subseteq S \text{ such that } \Pr[i \in S' \mid i \in S] \geq 1 - \varepsilon \text{ and for all } v \in \text{supp}(\tilde{p}_i \mid i \in S'), \\ & \left| \Pr[o_i^* = 1 \mid i \in S' \text{ and } \tilde{p}_i = v] - v \right| \leq \varepsilon. \end{aligned}$$

These conditional versions of multiaccuracy and multicalibration closely resemble their original definitions in Hébert-Johnson et al. (2018). They capture the intuition that the predictions of a multiaccurate (resp. multicalibrated) predictor are approximately accurate in expectation (resp. calibrated) on each subpopulation under consideration. It is also possible to give a conditional version of our definition of *strict* multicalibration:

Definition 26 Assume that $\mathcal{O} = \{0, 1\}$ and $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$. We say a predictor \tilde{p} is conditionally and strictly $(\mathcal{C}, \varepsilon)$ -multicalibrated if

$$\begin{aligned} & \text{there exists } V \subset [0, 1] \text{ such that } \Pr[\tilde{p}_i \in V] \geq 1 - \varepsilon \text{ and for all } v \in V, \\ & \text{for all } \mathbf{1}_S \in \mathcal{C} \text{ such that } \Pr[i \in S \mid \tilde{p}_i = v] \geq \varepsilon, \\ & \left| \Pr[o_i^* = 1 \mid i \in S \text{ and } \tilde{p}_i = v] - v \right| \leq \varepsilon. \end{aligned}$$




Figure 3: Relationships to Prior Definitions

Comparing Definitions 25 and 26 gives insight into the qualitative difference between multicalibration and strict multicalibration. Specifically, strict multicalibration *reverses the order of quantifiers* in the definition of multicalibration. If \tilde{p} is a strictly multicalibrated predictor, then most of its v -level sets satisfy the fairness guarantee uniformly across all protected subpopulations. In other words, \tilde{p} is \mathcal{C} -multiaccurate on its v -level set. In contrast, if \tilde{p} is multicalibrated but not strictly so, then each level set $\tilde{p}^{-1}(v)$ may fail the test of calibration on some subpopulation, and perhaps a different one for each v in the range of \tilde{p} . If one intends to use multicalibration as a certificate of *fairness* for a prediction algorithm, then such behavior is clearly undesirable.

The next theorem shows that when $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$ is closed under complement, the definitions in this section are equivalent to those of the previous section up to a polynomial change in ε . The proof is based on a simple application of Markov's inequality that previously appeared in [Gopalan et al. \(2022a\)](#).

Theorem 27 *Assume $\mathcal{O} = \{0, 1\}$ and $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$ is closed under complement. For each arrow from Definition A to Definition B in Figure 3, Definition A with parameters $(\mathcal{C}, \varepsilon)$ implies Definition B with parameters $(\mathcal{C}, \varepsilon^c)$ for sufficiently small $\varepsilon > 0$ and an absolute constant $c \in (0, 1)$.*

Proof We consider each of the six implications separately:

(1 \implies 24). If $c = \mathbf{1}_S \in \mathcal{C}$, then

$$|\Pr[o_i^* = 1 \mid i \in S] - \Pr[\tilde{o}_i = 1 \mid i \in S]| = \frac{|\Pr[o_i^* = 1, i \in S] - \Pr[\tilde{o}_i = 1, i \in S]|}{\Pr[i \in S]}.$$

The numerator of this fraction is at most $\delta((c_i, \tilde{o}_i), (c_i, o_i^*))$, which, by Definition 1, is at most ε . If the denominator satisfies $\Pr[i \in S] \geq \sqrt{\varepsilon}$, then the value of the fraction can be at most $\sqrt{\varepsilon}$. Thus, Definition 1 with parameter ε implies Definition 24 with parameter $\sqrt{\varepsilon}$.

(2 \implies 25). For $c = \mathbf{1}_S \in \mathcal{C}$ and $v \in \tilde{p}(\mathcal{X})$, consider the *multicalibration violation*

$$\nabla_{S,v} = \left| \Pr[o_i^* = 1 \mid i \in S \text{ and } \tilde{p}_i = v] - v \right| = \delta(\tilde{o}_i, o_i^* \mid \tilde{p}_i = v, i \in S),$$

and observe that

$$\mathbf{E}[\nabla_{S,\tilde{p}_i}] = \delta((\tilde{o}_i, \tilde{p}_i), (o_i^*, \tilde{p}_i) \mid i \in S) \leq \frac{\delta((c_i, \tilde{o}_i, \tilde{p}_i), (c_i, o_i^*, \tilde{p}_i))}{\Pr[i \in S]}.$$

By Definition 2, the numerator of this fraction is at most ε . If the denominator satisfies $\Pr[i \in S] \geq \varepsilon^{1/3}$, it follows that $\mathbf{E}[\nabla_{S,\tilde{p}_i}] \leq \varepsilon^{2/3}$. Let $S' = \{j \in S : \nabla_{S,\tilde{p}_j} \leq \varepsilon^{1/3}\}$. By Markov's inequality,

$$\Pr[i \notin S' \mid i \in S] \leq \frac{\mathbf{E}[\nabla_{S,\tilde{p}_i}]}{\varepsilon^{1/3}} \leq \varepsilon^{1/3}.$$

Thus, Definition 2 with parameter ε implies Definition 25 with parameter $\varepsilon^{1/3}$.

(3 \implies 26). Let $V = \{v \in \tilde{\rho}(\mathcal{X}) : \max_{c \in \mathcal{C}} \delta((c_i, \tilde{o}_i), (c_i, o_i^*) | \tilde{p}_i) \leq \varepsilon^{2/3}\}$. By Markov,

$$\Pr[\tilde{p}_i \notin V] \leq \frac{\mathbf{E} \left[\max_{c \in \mathcal{C}} \delta((c_i, \tilde{o}_i), (c_i, o_i^*) | \tilde{p}_i) \right]}{\varepsilon^{2/3}},$$

which is at most $\varepsilon^{1/3}$ by Definition 3. For $c = \mathbf{1}_S \in \mathcal{C}$ and $v \in V$, we have

$$\nabla_{S,v} = \delta(\tilde{o}_i, o_i^* | \tilde{p}_i = v, i \in S) \leq \frac{\delta((c_i, \tilde{o}_i), (c_i, o_i^*) | \tilde{p}_i = v)}{\Pr[i \in S | \tilde{p}_i = v]}.$$

The numerator of this fraction is at most $\varepsilon^{2/3}$ by construction of V . If the denominator satisfies $\Pr[i \in S | \tilde{p}_i = v] \geq \varepsilon^{1/3}$, then it follows that $\nabla_{S,v} \leq \varepsilon^{1/3}$. Thus, Definition 3 with parameter ε implies Definition 26 with parameter $\varepsilon^{1/3}$.

(24 \implies 1). For $c = \mathbf{1}_S \in \mathcal{C}$, either $\Pr[i \in S] \leq \varepsilon$ or $|\Pr[\tilde{o}_i | i \in S] - \Pr[o_i^* | i \in S]| \leq \varepsilon$ by Definition 24. Thus, their product satisfies $|\Pr[i \in S, o_i^* = 1] - \Pr[i \in S, \tilde{o}_i = 1]| \leq \varepsilon$. Since \mathcal{C} is closed under complement, the same inequality holds with $\mathcal{X} \setminus S$ in place of S . It follows that

$$\delta((c_i, \tilde{o}_i), (c_i, o_i^*)) = |\Pr[i \in S, o_i^* = 1] - \Pr[i \in S, \tilde{o}_i = 1]| + |\Pr[i \notin S, o_i^* = 1] - \Pr[i \notin S, \tilde{o}_i = 1]|$$

is at most 2ε . Thus, Definition 24 with parameter ε implies Definition 1 with parameter 2ε .

(25 \implies 2). For $c = \mathbf{1}_S \in \mathcal{C}$, we want to upper bound

$$\delta((c_i, \tilde{o}_i, \tilde{p}_i), (c_i, o_i^*, \tilde{p}_i)) = \delta((\tilde{o}_i, \tilde{p}_i), (o_i^*, \tilde{p}_i) | i \in S) \Pr[i \in S] + \delta((\tilde{o}_i, \tilde{p}_i), (o_i^*, \tilde{p}_i) | i \notin S) \Pr[i \notin S].$$

We will bound the two terms on the right side separately. By Definition 25, there exists $S' \subseteq S$ such that $\Pr[i \notin S' | i \in S] \leq \varepsilon$ and $\nabla_{S, \tilde{p}_j} \leq \varepsilon$ for all $j \in S'$. It follows that the first term satisfies

$$\delta((\tilde{o}_i, \tilde{p}_i), (o_i^*, \tilde{p}_i) | i \in S) \Pr[i \in S] \leq \Pr[i \in S \setminus S'] + \mathbf{E}[\nabla_{S, \tilde{p}_i} | i \in S'] \leq 2\varepsilon.$$

Since \mathcal{C} is closed under complement, we also have

$$\delta((\tilde{o}_i, \tilde{p}_i), (o_i^*, \tilde{p}_i) | i \notin S) \Pr[i \notin S] \leq 2\varepsilon.$$

Thus, Definition 25 with parameter ε implies Definition 2 with parameter 4ε .

(26 \implies 3). Take V as in Definition 26. Then

$$\mathbf{E} \left[\max_{c \in \mathcal{C}} \delta((c_i, \tilde{o}_i), (c_i, o_i^*) | \tilde{p}_i) \right] \leq \Pr[\tilde{p}_i \notin V] + \mathbf{E} \left[\max_{c \in \mathcal{C}} \delta((c_i, \tilde{o}_i), (c_i, o_i^*) | \tilde{p}_i) | \tilde{p}_i \in V \right].$$

By our choice of V , the first term satisfies $\Pr[\tilde{p}_i \notin V] \leq \varepsilon$. For $v \in V$, it remains to upper bound

$$\delta((c_i, \tilde{o}_i), (c_i, o_i^*) | \tilde{p}_i = v) = \Pr[i \in S | \tilde{p}_i = v] \nabla_{S,v} + \Pr[i \notin S | \tilde{p}_i = v] \nabla_{\mathcal{X} \setminus S, v}.$$

We will bound the two terms on the right side separately. By our choice of V , either $\Pr[i \in S | \tilde{p}_i = v] \leq \varepsilon$ or $\nabla_{S,v} \leq \varepsilon$. Thus, their product satisfies $\Pr[i \in S | \tilde{p}_i = v] \nabla_{S,v} \leq \varepsilon$. Similarly,

$\Pr[i \notin S \mid \tilde{p}_i = v] \nabla_{\mathcal{X} \setminus S, v} \leq \varepsilon$. Thus, Definition 26 with parameter ε implies Definition 3 with parameter 3ε . \blacksquare

Until this point, we have focused on the case $\mathcal{O} = \{0, 1\}$ and $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$. However, there are other definitions of multi-calibration in the algorithmic fairness literature that apply to more general sets \mathcal{O} and \mathcal{C} . One particularly noteworthy extension, introduced in Gopalan *et al.* Gopalan et al. (2022a), applies to the case that $\mathcal{O} = \{0, 1\}$ and $\mathcal{C} \subseteq [0, 1]^{\mathcal{X}}$. We include a rephrased statement here:

Definition 28 Assume $\mathcal{O} = \{0, 1\}$ and $\mathcal{C} \subseteq [0, 1]^{\mathcal{X}}$. We say \tilde{p} satisfies covariance-based $(\mathcal{C}, \varepsilon)$ -multi-calibration if

$$\mathbf{E} |\text{Cov}(c_i, o_i^* \mid \tilde{p}_i)| \leq \varepsilon$$

for all $c \in \mathcal{C}$.

Rather than measuring the statistical distance between (c_i, \tilde{o}_i) and (c_i, o_i^*) given \tilde{p}_i as we do, this definition measures the absolute value of the covariance of c_i and o_i^* given \tilde{p}_i . We conclude this section by showing that our version of multi-calibration is at least as strong as this covariance-based version whenever both are applicable.

Theorem 29 Assume $\mathcal{O} = \{0, 1\}$ and $\mathcal{C} \subseteq [0, 1]^{\mathcal{X}}$. If \tilde{p} is $(\mathcal{C}, \varepsilon)$ -multi-calibrated, then \tilde{p} also satisfies covariance-based $(\mathcal{C}, \varepsilon)$ -multi-calibration.

Proof Fix $c : \mathcal{X} \rightarrow [0, 1]$ and let $\mathcal{Y} \subseteq [0, 1]$ be the range of c . Since we assume \mathcal{X} is finite, so is \mathcal{Y} . Some straightforward algebra shows that

$$\text{Cov}(c_i, o_i^* \mid \tilde{p}_i) = \sum_{y \in \mathcal{Y}} \left(y - \frac{1}{2} \right) \begin{pmatrix} \Pr[c_i = y, o_i^* = 1 \mid \tilde{p}_i] & + \Pr[c_i = y \mid \tilde{p}_i] \Pr[\tilde{o}_i = 1 \mid \tilde{p}_i] \\ - \Pr[c_i = y, \tilde{o}_i = 1 \mid \tilde{p}_i] & - \Pr[c_i = y \mid \tilde{p}_i] \Pr[o_i^* = 1 \mid \tilde{p}_i] \end{pmatrix}.$$

We will split the above expression into a sum of two parts and bound the expected absolute value of each. First, because $|y - 1/2| \leq 1/2$ for each $y \in \mathcal{Y}$ and \tilde{p} is $(\mathcal{C}, \varepsilon)$ -multi-calibrated, we have

$$\mathbf{E} \left| \sum_{y \in \mathcal{Y}} \left(y - \frac{1}{2} \right) \left(\Pr[c_i = y, o_i^* = 1 \mid \tilde{p}_i] - \Pr[c_i = y, \tilde{o}_i = 1 \mid \tilde{p}_i] \right) \right| \leq \frac{1}{2} \mathbf{E} [\delta((c_i, o_i^*), (c_i, \tilde{o}_i) \mid \tilde{p}_i)] \leq \frac{\varepsilon}{2}.$$

Using the additional fact that $\sum_{y \in \mathcal{Y}} \Pr[c_i = y \mid \tilde{p}_i] = 1$, we see that

$$\mathbf{E} \left| \sum_{y \in \mathcal{Y}} \left(y - \frac{1}{2} \right) \Pr[c_i = y \mid \tilde{p}_i] \left(\Pr[\tilde{o}_i = 1 \mid \tilde{p}_i] - \Pr[o_i^* = 1 \mid \tilde{p}_i] \right) \right| \leq \frac{1}{2} \mathbf{E} [\delta(\tilde{o}_i, o_i^* \mid \tilde{p}_i)] \leq \frac{\varepsilon}{2}.$$

By the triangle inequality, we conclude that

$$\mathbf{E} |\text{Cov}(c_i, o_i^* \mid \tilde{p}_i)| \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2},$$

so \tilde{p} satisfies covariance-based $(\mathcal{C}, \varepsilon)$ -multi-calibration. \blacksquare

Appendix E. Omniprediction

In this section, we show how our new, statistical distance-based notions of multiaccuracy and multicalibration lend themselves naturally to applications by giving a remarkably simple extension of a state-of-the-art result in *omniprediction* to the multiclass setting. We will also extend the concept of omniprediction to consider loss functions that may depend on information of individuals. We remark that while our theorems hold whether the functions $c \in \mathcal{C}$ are discrete or continuous, they are only operationalizable for discrete-valued functions.

A concept introduced by [Gopalan et al. \(2022a\)](#), an *omnipredictor* is a single predictor capable of minimizing a wide range of *loss functions* $\ell \in \mathcal{L}$ while achieving competitive performance against a large class of *hypotheses* $c \in \mathcal{C}$. Informally speaking, the original omniprediction theorem [Gopalan et al. \(2022a\)](#) showed that any predictor \tilde{p} satisfying an appropriate multicalibration condition must also be an omnipredictor for all convex, Lipschitz, and bounded loss functions. However, a recent work [Gopalan et al. \(2023a\)](#) made significant strides by relaxing the assumptions of this theorem while strengthening its conclusion. The stronger version of the omniprediction theorem in [Gopalan et al. \(2023a\)](#) only assumes that \tilde{p} is multiaccurate and calibrated (not multicalibrated) and establishes omniprediction even for non-convex loss functions.

In what follows, let \mathcal{C} , as usual, be a collection of functions $c : \mathcal{X} \rightarrow \mathcal{Y}$, which we now call *hypothesis* functions. Also, let \mathcal{L} be a collection of *loss functions* $\ell : \mathcal{Y} \times \mathcal{O} \rightarrow [0, 1]$. Note that each loss function ℓ takes as input both an outcome $o \in \mathcal{O}$ and an *action* $y \in \mathcal{Y}$. It outputs a real number between 0 and 1 measuring the cost of choosing action y when the outcome is o . Also, consider the following notion of post-processing a prediction to minimize a loss function, which we have modified slightly from its form in [Gopalan et al. \(2022a, 2023a\)](#).

Definition 30 Say that $\text{post}_\ell : \Delta\mathcal{O} \rightarrow \mathcal{Y}$ is a post-processing function for the loss $\ell \in \mathcal{L}$ if

$$\text{post}_\ell(v) \in \operatorname{argmin}_{y \in \mathcal{Y}} \mathbf{E}_{o \sim v}[\ell(y, o)]$$

for each distribution $v \in \Delta\mathcal{O}$. For ease of notation, we also write $v^\ell = \text{post}_\ell(v)$.

We now recall the definition of an omnipredictor.

Definition 31 Say that $\tilde{p} : \mathcal{X} \rightarrow \Delta\mathcal{O}$ is a $(\mathcal{L}, \mathcal{C}, \varepsilon)$ -omnipredictor if

$$\mathbf{E}[\ell(\tilde{p}_i^\ell, o_i^*)] \leq \mathbf{E}[\ell(c_i, o_i^*)] + \varepsilon$$

for all $\ell \in \mathcal{L}$ and $c \in \mathcal{C}$.

In order to state the theorem of interest, we first emphasize an important special case of the definition of multicalibration in [Section 3](#):

Definition 32 A predictor \tilde{p} is ε -calibrated if $(\tilde{o}_i, \tilde{p}_i) \approx_\varepsilon (o_i^*, \tilde{p}_i)$.

Indeed, a predictor is ε -calibrated if and only if it is $(\{\mathbf{1}_{\mathcal{X}}\}, \varepsilon)$ -multicalibrated. With these definitions in hand, we are ready to present the main proof of this section. For clarity, we will first consider the special case that $\varepsilon = 0$.

The following two lemmas will be of use. The first says that calibrated predictions, even when post-processed, incur the same loss on real outcomes as on modeled outcomes.

Lemma 33 *If \tilde{p} is 0-calibrated, $f : \Delta\mathcal{O} \rightarrow \mathcal{Y}$ is any function, and $\ell \in \mathcal{L}$, then*

$$\mathbf{E}[\ell(f(\tilde{p}_i), o_i^*)] = \mathbf{E}[\ell(f(\tilde{p}_i), \tilde{o}_i)].$$

Proof 0-calibration means that $(\tilde{o}_i, \tilde{p}_i)$ and (o_i^*, \tilde{p}_i) have the same joint distribution. \blacksquare

The second lemma says that $c : \mathcal{X} \rightarrow \mathcal{Y}$ incurs the same loss on real outcomes as on modeled outcomes if the predictor is $(\mathcal{C}, 0)$ -multiaccurate.

Lemma 34 *If \tilde{p} is $(\mathcal{C}, 0)$ -multiaccurate, then*

$$\mathbf{E}[\ell(c_i, o_i^*)] = \mathbf{E}[\ell(c_i, \tilde{o}_i)]$$

for all $\ell \in \mathcal{L}$ and $c \in \mathcal{C}$.

Proof $(\mathcal{C}, 0)$ -multiaccuracy means that (c_i, o_i^*) and (c_i, \tilde{o}_i) have the same joint distribution. \blacksquare

We now state and prove a rephrased version of one of the main theorems of [Gopalan et al. \(2023a\)](#) in terms of our new language. In [Gopalan et al. \(2023a\)](#), the assumption of the theorem is that \tilde{p} is multiaccurate with respect to all $[0, 1]$ -bounded functions of the level sets of each $c \in \mathcal{C}$. In our statement of the theorem, these criteria are encapsulated naturally by our statistical distance-based formulation of multiaccuracy:

Theorem 35 *If \tilde{p} is ε_1 -calibrated and $(\mathcal{C}, \varepsilon_2)$ -multiaccurate, then \tilde{p} is an $(\mathcal{L}, \mathcal{C}, \varepsilon_1 + \varepsilon_2)$ -omnipredictor.*

Proof First consider the case that $\varepsilon_1 = \varepsilon_2 = 0$. Since \tilde{p}_i^ℓ is a function of \tilde{p}_i , we have

$$\begin{aligned} \mathbf{E}[\ell(\tilde{p}_i^\ell, o_i^*)] &= \mathbf{E}[\ell(\tilde{p}_i^\ell, \tilde{o}_i)] && \text{by Lemma 33,} \\ &\leq \mathbf{E}[\ell(c_i, \tilde{o}_i)] && \text{by Definition 30,} \\ &= \mathbf{E}[\ell(c_i, o_i^*)] && \text{by Lemma 34.} \end{aligned}$$

In the general case, standard properties of statistical distance ensure that the two expectation terms in Lemma 33 now differ by at most ε_1 , and the two expectation terms in Lemma 34 now differ by at most ε_2 . Here, we have used the assumption that the range of each $\ell \in \mathcal{L}$ is bounded between 0 and 1. Adding these two slack terms to the first and third lines, respectively, of the above calculation yields

$$\mathbf{E}[\ell(\tilde{p}_i^\ell, o_i^*)] \leq \mathbf{E}[\ell(c_i, o_i^*)] + \varepsilon_1 + \varepsilon_2,$$

so \tilde{p} is an $(\mathcal{L}, \mathcal{C}, \varepsilon_1 + \varepsilon_2)$ -omnipredictor. \blacksquare

For the sake of comparison, we also include a version of the original omniprediction proof of [Gopalan et al. \(2022a\)](#) in our current language. For simplicity, we only state the case of $\varepsilon_1 = \varepsilon_2 = 0$ but remark that additional Lipschitzness assumptions on \mathcal{L} would be required in the case of $\varepsilon_1, \varepsilon_2 > 0$.

Theorem 36 *Assume $\mathcal{O} = \{0, 1\}$ and $\mathcal{Y} \subseteq [0, 1]$. If \tilde{p} satisfies covariance-based $(\mathcal{C}, 0)$ -multi-calibrated and each $\ell \in \mathcal{L}$ is convex in its first input, then \tilde{p} is an $(\mathcal{L}, \mathcal{C}, 0)$ -omnipredictor.*

Proof Rephrasing the proof from [Gopalan et al. \(2022a\)](#) yields:

$$\begin{aligned}
 \mathbf{E}[\ell(\tilde{p}_i^\ell, o_i^*)] &= \mathbf{E}[\ell(\tilde{p}_i^\ell, \tilde{o}_i)] && \text{by Lemma 33} \\
 &\leq \mathbf{E}[\ell(\mathbf{E}[c_i \mid \tilde{p}_i], \tilde{o}_i)] && \text{by Definition 30} \\
 &= \mathbf{E}[\ell(\mathbf{E}[c_i \mid \tilde{p}_i], o_i^*)] && \text{by Lemma 33} \\
 &= \mathbf{E}[\ell(\mathbf{E}[c_i \mid \tilde{p}_i, o_i^*], o_i^*)] && \text{by Definition 28} \\
 &\leq \mathbf{E}[\ell(c_i, o_i^*)] && \text{by convexity of } \ell(-, 0) \text{ and } \ell(-, 1).
 \end{aligned}$$

In the second equality, we use the fact that $\mathbf{E}[c_i \mid \tilde{p}_i]$ is a function of \tilde{p}_i . In the third equality, we used the fact that $\mathbf{E}[c_i \mid \tilde{p}_i] = \mathbf{E}[c_i \mid \tilde{p}_i, o_i^*]$ (i.e., that c_i and o_i^* are conditionally uncorrelated given \tilde{p}_i). This is by definition of covariance-based $(\mathcal{C}, 0)$ -multi-calibration. \blacksquare

E.1. Loss Functions Dependent on Information of Individuals

So far the concept of omniprediction considers loss functions ℓ depending on an action $y \in \mathcal{Y}$ and outcome $o \in \mathcal{O}$. In this section, we extend the notion to consider loss functions that may depend on additional information about an individual, that is, the loss $\ell(y, o, z)$ also depend on information $z \in \mathcal{T}$ of an individual i w.r.t. whom the action y and outcome o are associated with. Here, we can consider different relevant information z about an individual captured by a class of functions $\mathcal{Z} = \{z : \mathcal{X} \rightarrow \mathcal{T}\}$. Now omniprediction w.r.t. loss functions \mathcal{L} , hypotheses \mathcal{C} , and auxiliary information functions \mathcal{Z} should guarantee that the predictor \tilde{p} has competitive performance against hypotheses in \mathcal{C} , in minimizing a large class of losses w.r.t. expressive information of individuals.

To formally state the result, we first slightly rephrase the setting of omniprediction. In prior works and above, we compare the loss incurred using an action, $y_i = \text{post}_\ell(\tilde{p}_i)$, derived from postprocessing the predicted outcome distribution \tilde{p}_i , with loss incurred using actions c_i prescribed by a hypothesis $c \in \mathcal{C}$. The postprocessing function post ([Definition 30](#)) ensures that y_i is Bayes optimal w.r.t. the composed function $\ell'(p, o) = \ell(\text{post}_\ell(p), o)$, in the sense that the expectation of ℓ' on input o drawn from a distribution p^* is minimized when the true probability distribution p^* is given as the input — namely, $p^* = \text{argmin}_{p \in \Delta\mathcal{O}} \mathbf{E}_{o \sim p^*}[\ell'(p, o)]$. We say that such loss functions are *Bayes optimal*. Then we can equivalently state previous results w.r.t. a class of Bayes optimal loss functions \mathcal{L}' and a hypothesis class \mathcal{C} mapping individuals in \mathcal{X} to outcome distributions $\Delta\mathcal{O}$ (instead of actions). Next, we describe our extension formally in this setting.

Definition 37 *Let \mathcal{L} be a collection of loss functions $\ell : \mathcal{Y} \times \mathcal{O} \times \mathcal{T} \rightarrow [0, 1]$. We say that \mathcal{L} is ε -Bayes optimal if every $\ell \in \mathcal{L}$ satisfies that*

$$\forall z \in \mathcal{T}, p^*, p \in \Delta\mathcal{O}, \quad \mathbf{E}_{o \sim p^*}[\ell(p^*, o, z)] \leq \mathbf{E}_{o \sim p^*}[\ell(p, o, z)] + \varepsilon$$

We now extend the definition of an omnipredictor. Let \mathcal{C} be a collection of hypothesis from \mathcal{X} to $\Delta\mathcal{O}$, and \mathcal{Z} be a collection of auxiliary information functions from \mathcal{X} to \mathcal{T} .

Definition 38 *Say that $\tilde{p} : \mathcal{X} \rightarrow \Delta\mathcal{O}$ is a $(\mathcal{L}, \mathcal{C}, \mathcal{Z}, \varepsilon)$ -omnipredictor if*

$$\mathbf{E}[\ell(\tilde{p}_i, o_i^*, z_i)] \leq \mathbf{E}[\ell(c_i, o_i^*, z_i)] + \varepsilon$$

for all $\ell \in \mathcal{L}$, $c \in \mathcal{C}$, and $z \in \mathcal{Z}$.

We show that our simple proof of omniprediction in the previous section can be easily adapted to accommodate the richer class of loss functions we consider. The main difference is that we now need the predictor \tilde{p} to be multicalibrated w.r.t. \mathcal{Z} and multiaccurate w.r.t. $\mathcal{C} \parallel \mathcal{Z} = \{c \parallel z : c \in \mathcal{C}, z \in \mathcal{Z}\}$.

Theorem 39 *Let \mathcal{L} be a collection of ε_3 -Bayes optimal loss functions. If \tilde{p} is $(\mathcal{Z}, \varepsilon_1)$ -multicalibrated and $(\mathcal{C} \parallel \mathcal{Z}, \varepsilon_2)$ -multiaccurate, then \tilde{p} is an $(\mathcal{L}, \mathcal{C}, \mathcal{Z}, \varepsilon_1 + \varepsilon_2 + \varepsilon_3)$ -omnipredictor.*

Proof First consider the case that $\varepsilon_1 = \varepsilon_2 = \varepsilon_3 = 0$. We have

$$\begin{aligned} \mathbf{E}[\ell(\tilde{p}_i, o_i^*, z_i)] &= \mathbf{E}[\ell(\tilde{p}_i, \tilde{o}_i, z_i)] && \text{by Definition 2 and } \tilde{p} \text{ is multicalibrated w.r.t. } \mathcal{Z}, \\ &\leq \mathbf{E}[\ell(c_i, \tilde{o}_i, z_i)] && \text{by Definition 37,} \\ &= \mathbf{E}[\ell(c_i, o_i^*, z_i)] && \text{by Lemma 34 and } \tilde{p} \text{ is multiaccurate w.r.t. } \mathcal{C} \parallel \mathcal{Z}. \end{aligned}$$

In the general case, standard properties of statistical distance ensure that the first and third equality of expectation terms differ by at most ε_1 and ε_2 (using the fact that the range of each $\ell \in \mathcal{L}$ is bounded between 0 and 1). The right hand side of the second inequality has an additional ε_3 term by definition of Bayes optimality of \mathcal{L} . Therefore, we obtain

$$\mathbf{E}[\ell(\tilde{p}_i, o_i^*, z_i)] \leq \mathbf{E}[\ell(c_i, o_i^*, z_i)] + \varepsilon_1 + \varepsilon_2 + \varepsilon_3,$$

so \tilde{p} is an $(\mathcal{L}, \mathcal{C}, \mathcal{Z}, \varepsilon_1 + \varepsilon_2 + \varepsilon_3)$ -omnipredictor. ■

Appendix F. Outcome Indistinguishability via No-Regret Learning

In this section, we present an algorithmic template that unifies two existing algorithms for achieving outcome indistinguishability (and hence multicalibration), through the lens of no-regret learning. Previously, in Section 5, we saw that an instantiation of the template yields algorithms with improved upper bounds on the sample complexity of multicalibration.

The two algorithms under consideration are similar in that both make iterative updates to an arbitrary initial predictor \tilde{p} . However, they differ in their implementations of the update rule. The first update rule selects $A \in \mathcal{A}$ that successfully distinguishes \tilde{p} from p^* with advantage ε , and makes an additive update to \tilde{p} resembling projected gradient descent Hébert-Johnson et al. (2018); Dwork et al. (2021). The second update rule also selects a distinguisher, but instead updates \tilde{p} in a multiplicative manner Kim et al. (2019).

To establish the claimed connection, we will first show in Appendix F.1 that the described algorithmic template can be instantiated with any update rule based on an algorithm with a *no-regret guarantee*. We will then discuss in Appendix F.2 how projected gradient descent and multiplicative weight updates can be viewed as instances of *mirror descent*, an algorithm with exactly the required no-regret guarantee. One benefit of this unified presentation via no-regret learning is that prior works require separate analyses for the two algorithms, but we only need a single, very simple, proof, relying only on the no-regret guarantee. We will also examine the relative merits of using projected gradient descent versus multiplicative weight updates for this role. (In brief, multiplicative weight updates work better in the multiclass setting, but projected gradient descent is more robust to a poor initialization.)

F.1. No-Regret Updates

We first recall the general framework of no-regret online learning. Consider T rounds of gameplay between two players, called the *decision-maker* and the *adversary*. In each round, the decision-maker chooses a *mixed strategy* (i.e., a probability distribution) over a finite set \mathcal{O} of available *pure strategies*. In response, the adversary adaptively chooses a *loss function*, which assigns an arbitrary numeric loss between 0 and 1 to each available pure strategy. At the end of the round, the decision-maker incurs a penalty equal to the expected loss of its chosen mixed strategy, while learning the entire description of the loss function (not just the penalty it incurred).

Intuitively, a decision-making algorithm satisfies a *no-regret guarantee* if the overall expected loss of a decision-maker employing the algorithm is no worse than the penalty the decision-maker would have incurred by playing any particular pure strategy against the same sequence of loss functions.

We now state a formal definition of no-regret learning. For simplicity, we focus our attention on decision-making algorithms whose strategy in round $t + 1$ is completely determined by what happened in round t , but this assumption is easy to relax.

Definition 40 A decision-making algorithm is specified by a distribution $D^{(1)} \in \Delta\mathcal{O}$ and a function update : $\Delta\mathcal{O} \times [0, 1]^{\mathcal{O}} \rightarrow \Delta\mathcal{O}$. We say it satisfies a (T, ε) -no-regret guarantee if

$$\frac{1}{T} \sum_{t=1}^T \mathbf{E}_{o \sim D^{(t)}} [L^{(t)}(o)] \leq \frac{1}{T} \sum_{t=1}^T L^{(t)}(o^*) + \varepsilon$$

for every pure strategy $o^* \in \mathcal{O}$, every sequence of loss functions $L^{(1)}, \dots, L^{(T)} \in [0, 1]^{\mathcal{O}}$, and the sequence of distributions $D^{(t)} \in \Delta\mathcal{O}$ given by $D^{(t+1)} = \text{update}(D^{(t)}, L^{(t)})$.

Example 1 The projected gradient descent and multiplicative weight update rules are

$$D^{(t+1)} = \text{proj}_{\Delta\mathcal{O}} \left(D^{(t)} - \eta L^{(t)} \right),$$

$$D^{(t+1)}(o) = \frac{D^{(t)}(o) e^{-\eta L^{(t)}(o)}}{\sum_{o' \in \mathcal{O}} D^{(t)}(o') e^{-\eta L^{(t)}(o')}},$$

respectively, for a parameter $\eta > 0$ called the step size.

We will state the standard no-regret guarantees for these update rules, along with appropriate initializations, in Appendix F.2.

Algorithm 2 shows how to achieve outcome indistinguishability via no-regret updates. Indeed, Algorithm 2 can be viewed as running $|\mathcal{X}|$ instances of a no-regret algorithm in parallel. Each instance corresponds to one member j of the population \mathcal{X} , and the distribution \tilde{p}_j corresponds to a mixed strategy over $o \in \mathcal{O}$. The predicted probabilities $\tilde{p}_j(o)$ are refined over multiple rounds.

The most important question is how the loss function is chosen in each round. Toward the goal of outcome indistinguishability against \mathcal{A} , given the current predictor \tilde{p} , the algorithm finds a distinguisher $A \in \mathcal{A}$ that has relatively high distinguishing advantage with respect to \tilde{p} . Such an adversary naturally defines a loss for each individual-outcome pair as follows:

$$L_j(o) = A(j, o, \tilde{p}).$$

We emphasize that though the no-regret algorithm is run separately for each member of \mathcal{X} , the choice of the distinguisher (and hence the loss functions) depends on the *entire* predictor \tilde{p} . Furthermore, it suffices to find a distinguisher with relatively high advantage as opposed to maximal advantage.

Finally, if \mathcal{X} is very large, it may be infeasible to run a separate instance of the no-regret algorithm for each $j \in \mathcal{X}$. This is not a problem because the collection of values $\tilde{p}_j(o)$ is implicitly defined by the distinguishers A found in each round, which in turn give an efficient representation of the predictor \tilde{p} as a whole.

Algorithm 2: Outcome Indistinguishability with No-Regret Updates

```

def Construct-via-No-Regret ( $\tilde{p}, \mathcal{A}, \varepsilon, \text{update}$ ):
    for  $A \in \mathcal{A}$  do
        // check for  $(\mathcal{A}, \varepsilon)$ -OI violation
        if  $\mathbf{E}[A(i, \tilde{o}_i, \tilde{p})] > \mathbf{E}[A(i, o_i^*, \tilde{p})] + \varepsilon$  then
            for  $j \in \mathcal{X}$  do
                for  $o \in \mathcal{O}$  do
                    // define loss function
                     $L_j(o) \leftarrow A(j, o, \tilde{p})$ 
                end
                // no-regret update
                 $\tilde{p}'_j \leftarrow \text{update}(\tilde{p}_j, L_j)$ 
            end
            // recurse
            return Construct-via-No-Regret ( $\tilde{p}', \mathcal{A}, \varepsilon, \text{update}$ )
        end
    end
    return  $\tilde{p}$ 
end

```

Theorem 41 *Suppose \mathcal{A} is closed under negation, meaning that $A \in \mathcal{A}$ iff $1 - A \in \mathcal{A}$, and that the function update satisfies a (T, ε) -no-regret guarantee when initialized at $p \in \Delta\mathcal{O}$. Set $\tilde{p}_j = p$ for each $j \in \mathcal{X}$. Then CONSTRUCT-VIA-NO-REGRET($\tilde{p}, \mathcal{A}, \varepsilon, \text{update}$) returns an $(\mathcal{A}, \varepsilon)$ -outcome-indistinguishable predictor in under T recursive calls.*

Proof Let $\tilde{p}^{(t)}$ denote the argument to the t^{th} recursive call (e.g., $\tilde{p}^{(1)} = \tilde{p}$), let $\tilde{o}_i^{(t)} \in \mathcal{O}$ be a random variable whose conditional distribution given i is specified by \tilde{p}_i , and suppose toward a contradiction that T or more recursive calls are made. The no-regret guarantee implies that

$$\frac{1}{T} \sum_{t=1}^T \left(\mathbf{E} \left[A(i, \tilde{o}_i^{(t)}, \tilde{p}^{(t)}) \right] - \mathbf{E} \left[A(i, o_i^*, \tilde{p}^{(t)}) \right] \right) \leq \varepsilon.$$

By design, if the algorithm does not terminate in T recursive calls, then each summand on the left is greater than ε , which leads to a contradiction. Thus, CONSTRUCT-VIA-NO-REGRET always

returns some predictor in under T recursive calls, which the stopping condition clearly ensures is $(\mathcal{A}, \varepsilon)$ -outcome-indistinguishable. \blacksquare

F.2. Mirror Descent Updates

In this section, we explain how projected gradient descent and multiplicative weight updates fit into the framework of no-regret learning and compare their advantages when used as the update rule in Algorithm 2. These two implementations of Algorithm 2 are typically analyzed by tracking a potential function measuring the “divergence” of a predictor \tilde{p} from the ground truth p^* as updates to \tilde{p} are made. The mirror descent perspective that we adopt in this section will clarify which “divergence” functions can be used in such an argument to derive a no-regret guarantee, while giving convergence rates for the corresponding update rules.

F.2.1. BACKGROUND

To begin, we state without proof some basic properties about the two algorithms under consideration. For more detail, we refer the reader to texts on convex optimization [Bubeck \(2015\)](#); [Nemirovsky and Yudin \(1983\)](#).

- **Projected Gradient Descent** Let $K \subset \mathbb{R}^n$ be a convex and compact constraint set. Then for any initialization $x_1 \in K$ and sequence $y_1, y_2, \dots \in \mathbb{R}^n$, the update rule

$$x_{t+1} = \text{proj}_K(x_t - \eta y_t)$$

with step size $\eta > 0$ satisfies the regret bound

$$\sum_{t=1}^T \langle y_t, x_t - x \rangle \leq \frac{\|x - x_1\|_2^2}{2\eta} + \frac{\eta}{2} \sum_{t=1}^T \|y_t\|_2^2$$

for all $T \in \mathbb{N}$ and $x \in K$. Here, $\|\cdot\|_p$ denotes the p -norm.

- **Multiplicative Weight Updates** Let $\Delta^{n-1} = \{w \in [0, 1]^n : \|w\|_1 = 1\}$. Then for any initialization $x_1 \in \Delta^{n-1} \cap \mathbb{R}_{>0}^n$ and sequence $y_1, y_2, \dots \in \mathbb{R}^n$, the update rule

$$x_{t+1,j} \propto x_{t,j} \exp(-\eta y_{t,j}), \quad x_{t+1} \in \Delta^{n-1}$$

with step size $\eta > 0$ satisfies the regret bound

$$\sum_{t=1}^T \langle y_t, x_t - x \rangle \leq \frac{D_{\text{KL}}(x \| x_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \|y_t\|_\infty^2$$

for all $T \in \mathbb{N}$ and $x \in \Delta^{n-1}$. Here, D_{KL} denotes the *Kullback-Leibler divergence* and $\|\cdot\|_\infty$ denotes the *maximum norm*.

- **Mirror Descent** Let $\|\cdot\|$ and $\|\cdot\|_*$ be dual norms on a finite-dimensional vector space V and its dual V^* , and let D be the *Bregman divergence associated with a mirror map that is 1-strongly convex with respect to $\|\cdot\|$ on $\Omega \subseteq V$* . Let $K \subset \bar{\Omega}$ be a convex and compact

constraint set. Then for any initialization $x_1 \in K \cap \Omega$ and sequence $y_1, y_2, \dots \in V^*$, the update rule

$$x_{t+1} = \operatorname{argmin}_{x \in K \cap \Omega} \langle y_t, x \rangle + \frac{D(x, x_t)}{\eta}$$

with step size $\eta > 0$ is well-defined and satisfies the regret bound

$$\sum_{t=1}^T \langle y_t, x_t - x \rangle \leq \frac{D(x, x_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \|y_t\|_*^2$$

for all $T \in \mathbb{N}$ and $x \in K$. Here, $\langle y, x \rangle$ denotes the real number $y(x)$.

The general update rule may be interpreted as selecting x_{t+1} that responds well to y_t without moving too far away from x_t , as enforced by the penalty term $D(x_{t+1}, x_t)/\eta$.

- **Relationships** One can show that projected gradient descent is a special case of mirror descent with $\|\cdot\| = \|\cdot\|_* = \|\cdot\|_2$ on \mathbb{R}^n and $D(p, q) = \frac{1}{2}\|p - q\|_2^2$ on $\Omega = \mathbb{R}^n$. Similarly, one can show that the multiplicative weights algorithm is a special case of mirror descent with $\|\cdot\| = \|\cdot\|_1$ and $\|\cdot\|_* = \|\cdot\|_\infty$ on \mathbb{R}^n and $K = \Delta^{n-1}$ and $D(p, q) = D_{\text{KL}}(p||q) = \sum_{j=1}^n p_j \log(p_j/q_j)$ on $\Omega = \mathbb{R}_{>0}^n$.

F.2.2. APPLICATION TO OI

Using the above notation, let $K = \Delta\mathcal{O}$ and $L = \max_{f:\mathcal{O} \rightarrow [0,1]} \|f\|_*$. Algorithm 3 specializes Algorithm 2 to the case of mirror descent updates, which includes projected gradient descent and multiplicative weight updates. In the pseudocode for Algorithm 3, the expression $\mathbf{E}_{o \sim p}[A(j, o, \tilde{p})]$ should be read as $\sum_{o \in \mathcal{O}} p(o)A(j, o, \tilde{p})$.

Algorithm 3: Outcome Indistinguishability with Mirror Descent Updates

```

def Construct-via-Mirror-Descent ( $\tilde{p}, \mathcal{A}, \varepsilon, D, \eta$ ):
    for  $A \in \mathcal{A}$  do
        // check for  $(\mathcal{A}, \varepsilon)$ -OI violation
        if  $\mathbf{E}[A(i, \tilde{o}_i, \tilde{p})] > \mathbf{E}[A(i, o_i^*, \tilde{p})] + \varepsilon$  then
            for  $j \in \mathcal{X}$  do
                // mirror descent update
                 $\tilde{p}'_j \leftarrow \operatorname{argmin}_{p \in \Delta\mathcal{O}} \mathbf{E}_{o \sim p}[A(j, o, \tilde{p})] + D(p, \tilde{p}_j)/\eta$ 
            end
            // recurse
            return Construct-via-Mirror-Descent ( $\tilde{p}', \mathcal{A}, \varepsilon, D, \eta$ )
        end
    end
    return  $\tilde{p}$ 
end
    
```

Theorem 42 $\text{CONSTRUCT-VIA-MIRROR-DESCENT}(\tilde{p}, \mathcal{A}, D, \varepsilon, \eta)$ with step size $\eta = \varepsilon/L^2$ returns an $(\mathcal{A}, \varepsilon)$ -outcome-indistinguishable predictor in under

$$2 \left(\frac{L}{\varepsilon} \right)^2 \mathbf{E}[D(p_i^*, \tilde{p}_i)]$$

recursive calls, where p_i^* denotes the conditional distribution of o_i^* given i .

Proof Let $\tilde{p}^{(t)}$ denote the argument to the t^{th} recursive call (e.g., $\tilde{p}^{(1)} = \tilde{p}$), and let $\tilde{o}_i^{(t)} \in \mathcal{O}$ be a random variable whose conditional distribution given i is specified by \tilde{p}_i . The mirror descent regret bound implies that

$$\sum_{t=1}^T \left(\mathbf{E} \left[A(i, \tilde{o}_i^{(t)}; \tilde{p}^{(t)}) \right] - \mathbf{E} \left[A(i, o_i^*; \tilde{p}^{(t)}) \right] \right) \leq \frac{\mathbf{E}[D(p_i^*, \tilde{p}_i)]}{\eta} + \frac{\eta T L^2}{2}.$$

By design, each summand on the left is greater than ε , which leads to a contradiction if the step size is $\eta = \varepsilon/L^2$ and there are $T \geq 2(L/\varepsilon)^2 \mathbf{E}[D(p_i^*, \tilde{p}_i)]$ calls. Thus, $\text{CONSTRUCT-VIA-MIRROR-DESCENT}$ always returns some predictor, which the stopping condition clearly ensures is $(\mathcal{A}, \varepsilon)$ -outcome-indistinguishable. \blacksquare

Example 2 Using projected gradient descent yields $L = \sqrt{|\mathcal{O}|}$ and a bound of

$$\frac{|\mathcal{O}| \mathbf{E} [\|p_i^* - \tilde{p}_i\|_2^2]}{\varepsilon^2} \leq \frac{2|\mathcal{O}|}{\varepsilon^2}$$

on the number of recursive calls. In the binary case, i.e., $\mathcal{O} = \{0, 1\}$, the update rule satisfies

$$\tilde{p}_j'(1) \leftarrow \text{proj}_{[0,1]} \left(\tilde{p}_j(1) - \frac{\eta}{2} (A(j, 1, \tilde{p}) - A(j, 0, \tilde{p})) \right),$$

which agrees with the original outcome indistinguishability algorithm [Dwork et al. \(2021\)](#).

Example 3 Using multiplicative weight updates yields $L = 1$ and a bound of $2\varepsilon^{-2} \mathbf{E}[D_{\text{KL}}(p_i^* \|\tilde{p}_i)]$ on the number of recursive calls. If we initialize \tilde{p}_j to the uniform distribution on \mathcal{O} for all $j \in \mathcal{X}$, then this bound reduces to $2\varepsilon^{-2} \log(|\mathcal{O}|)$, which has a better dependence on $|\mathcal{O}|$ than the bound for projected gradient descent but allows less flexibility in the initialization of \tilde{p} .

The original analyses of the generic multicalibration and outcome indistinguishability algorithms are not phrased in terms of no-regret bounds like our proofs of Theorems 41 and 42. Instead, they track changes to a *potential function* as \tilde{p} is iteratively updated. In our proof of Theorem 42, this potential function corresponds exactly to the quantity $\mathbf{E}[D(p_i^*, \tilde{p}_i)]$.

F.3. Weak Agnostic Learning

To prepare for the discussion of our new complexity upper bounds, we now present a variant of Algorithm 2 that abstracts the process of finding a distinguisher $A \in \mathcal{A}$ that distinguishes real from modeled outcomes with advantage ε . The abstraction we consider is based on that of a *weak agnostic learner*, which appeared in the original paper on multi-group fairness [Hébert-Johnson et al. \(2018\)](#).

Algorithm 4: Outcome Indistinguishability via Weak Agnostic Learning

```

def Construct-via-WAL ( $\varepsilon, \tilde{p}, \text{WAL}_{\mathcal{A}}, \text{update}$ ):
     $A \leftarrow \text{WAL}_{\mathcal{A}}(\varepsilon, \tilde{p})$  // check for  $(\mathcal{A}, \varepsilon)$ -OI violation
    if  $A = \perp$  then
        | return  $\tilde{p}$ 
    end
    for  $j \in \mathcal{X}$  do
        | for  $o \in \mathcal{O}$  do
            | |  $L_j(o) \leftarrow A(j, o, \tilde{p})$  // define loss function
            | end
            |  $\tilde{p}'_j \leftarrow \text{update}(\tilde{p}_j, L_j)$  // no-regret update
        | end
    return Construct-via-WAL ( $\varepsilon, \tilde{p}', \text{WAL}_{\mathcal{A}}, \text{update}$ ) // recurse
end

```

Definition 43 Let $\text{WAL}_{\mathcal{A}}$ be an algorithm that takes as input a parameter $\varepsilon > 0$ and a predictor \tilde{p} and outputs either a distinguisher $A \in \mathcal{A}$ or the symbol \perp . We assume that $\text{WAL}_{\mathcal{A}}$ has the ability to draw data samples that are i.i.d. copies of (i, o_i^*) . We say that $\text{WAL}_{\mathcal{A}}$ is a weak agnostic learner with failure probability β if the following two conditions hold:

- If there exists $A \in \mathcal{A}$ such that $\Delta_A(\tilde{p}) := \mathbf{E}[A(i, \tilde{o}_i, \tilde{p})] - \mathbf{E}[A(i, o_i^*, \tilde{p})] > \varepsilon$, then $\text{WAL}_{\mathcal{A}}(\varepsilon, \tilde{p})$ outputs $A' \in \mathcal{A}$ such that $\Delta_{A'}(\tilde{p}) > \varepsilon/2$ with probability at least $1 - \beta$.
- If every $A \in \mathcal{A}$ satisfies $\Delta_A(\tilde{p}) \leq \varepsilon$, then $\text{WAL}_{\mathcal{A}}(\varepsilon, \tilde{p})$ outputs either $A' \in \mathcal{A}$ such that $\Delta_{A'}(\tilde{p}) > \varepsilon/2$ or \perp with probability at least $1 - \beta$.

A simple application of Hoeffding’s inequality and a union bound yields the following lemma.

Lemma 44 For every family \mathcal{A} of distinguishers, there exists a weak agnostic learning algorithm $\text{WAL}_{\mathcal{A}}$ with failure probability β that draws at most $O(\log(|\mathcal{A}|/\beta)/\varepsilon^2)$ on input ε and \tilde{p} .

Algorithm 4 shows how to utilize $\text{WAL}_{\mathcal{A}}$ as a subroutine for achieving outcome indistinguishability.

Theorem 45 If \mathcal{A} is closed under complement, then there exists an algorithm $\text{WAL}_{\mathcal{A}}$, an algorithm update , and an initialization of \tilde{p} such that with probability at least $1 - \beta$, the algorithm $\text{CONSTRUCT}(\varepsilon, \tilde{p}, \text{WAL}_{\mathcal{A}}, \text{update})$ returns an $(\mathcal{A}, \varepsilon)$ -outcome-indistinguishable predictor in at most $\log(\ell)/\varepsilon^2$ recursive calls and using at most $\log(\ell) \log(|\mathcal{A}|/\beta)/\varepsilon^4$ samples, where $\ell = |\mathcal{O}|$.

Proof The proof of Theorem 42 and Example 3, along with the two properties in Definition 43, gives a $O(\log(\ell)/\varepsilon^2)$ upper bound on the number of updates. Lemma 44 upper bounds the number of samples needed per update. Choosing β appropriately and applying a union bound over the entire sequence of updates yields the claimed sample complexity upper bound. \blacksquare

F.4. Proofs for Section 5

Proof of Theorem 8 Throughout this proof, and for the remainder of this section, let

$$\Delta_A(\tilde{p}) = \left| \mathbf{E}[A(i, \tilde{o}_i, \tilde{p})] - \mathbf{E}[A(i, o_i^*, \tilde{p})] \right|$$

for a distinguisher $A \in \mathcal{A}$ and predictor $\tilde{p} : \mathcal{X} \rightarrow \Delta\mathcal{O}$. We say that Algorithm 1 *performs an update* in iteration s if it reaches the line “changed \leftarrow true” during iteration s of the outermost for-loop. We claim that the “exhaustive search” over $A \in \mathcal{A}$ in the algorithm correctly implements the weak agnostic learning step described in Appendix F.3. Indeed, a standard application of a Chernoff bound and union bound shows that for an appropriately chosen number $n \lesssim \log(|\mathcal{A}|t)/\varepsilon^2$ of samples per iteration, the following two properties hold with probability at least 99% across all iterations $s \in \{0, 1, \dots, t-1\}$:

- (a) If some $A \in \mathcal{A}$ satisfies $\Delta_A(\tilde{p}^{(s)}) > \varepsilon$, then the algorithm performs an update in iteration s .
- (b) If the algorithm performs an update in iteration s using $A^{(s)} \in \mathcal{A}$, then $\Delta_{A^{(s)}}(\tilde{p}^{(s)}) > \varepsilon/3$.

Since the algorithm only outputs $\tilde{p} \neq \perp$ after an iteration when no update was performed, property (a) immediately implies that such a predictor \tilde{p} must be $(\mathcal{A}, \varepsilon)$ -outcome-indistinguishable. It remains to show that Algorithm 1 will never output \perp for an appropriate value $t \lesssim \log(|\mathcal{O}|)/\varepsilon^2$. However, this follows immediately from property (b) and Example 3.

Proof of Theorem 10 Observe that the family $\mathcal{A}_{\mathcal{C}, \mathcal{G}}^{\text{MC}}$ has size $2^{|\mathcal{Y}||\mathcal{O}||\mathcal{G}|}|\mathcal{C}|$. By Lemma 23, we can choose an η -covering \mathcal{G} of $\Delta\mathcal{O}$ in such a way that $|\mathcal{G}| < (3/\eta)^{\ell-1}$. Thus, by Lemma 8, Algorithm 1 with input $\mathcal{A}_{\mathcal{C}, \mathcal{G}}^{\text{MC}}$ and parameters $\varepsilon/4$ and $\eta = 3\varepsilon/4$ outputs a $(\mathcal{A}_{\mathcal{C}, \mathcal{G}}^{\text{MC}}, \varepsilon/4)$ -outcome-indistinguishable predictor \tilde{p} with probability at least 99% using at most

$$tn \lesssim \left(\log(|\mathcal{C}|) + |\mathcal{Y}||\mathcal{O}| \left(\frac{4}{\varepsilon} \right)^{|\mathcal{O}|-1} \right) \cdot \frac{\log(|\mathcal{O}|)}{\varepsilon^4}$$

samples. Since \mathcal{G} is an η -covering of $\Delta\mathcal{O}$, the inequality $\delta(\hat{o}_i, \tilde{o}_i | i) \leq \eta$ holds almost surely, so

$$\delta((c_i, \hat{o}_i, \hat{p}_i), (c_i, o_i^*, \hat{p}_i)) \leq \delta((c_i, \tilde{o}_i, \hat{p}_i), (c_i, o_i^*, \hat{p}_i)) + \eta.$$

The first term is at most $\varepsilon/4$ by the definition of statistical distance and $(\mathcal{A}_{\mathcal{C}, \mathcal{G}}^{\text{MC}}, \varepsilon/4)$ -outcome-indistinguishability of \tilde{p} , so the right hand side is at most ε . We conclude that the discretized predictor \hat{p} is $(\mathcal{C}, \varepsilon)$ -multicalibrated.

Proof of Theorem 12 A standard counting argument shows that the family $\mathcal{A}_{\mathcal{C}, k}^{\text{MC}}$ has size $\binom{k+\ell-1}{\ell-1} \ell |\mathcal{C}|$. Thus, by Lemma 8, Algorithm 1 with input $\mathcal{A}_{\mathcal{C}, k}^{\text{MC}}$ and parameters $\varepsilon/4$ and $\eta = 3\varepsilon/4$ outputs a $(\mathcal{A}_{\mathcal{C}, k}^{\text{MC}}, \varepsilon)$ -outcome-indistinguishable predictor \tilde{p} with probability at least 99% using at most

$$tn \lesssim \left(\log(|\mathcal{C}|) + k \log \left(\frac{\ell}{k} \right) + \log \left(\frac{\ell}{\varepsilon} \right) \right) \cdot \frac{\log(\ell)}{\varepsilon^4}$$

samples. By definition, such a predictor \tilde{p} is $(\mathcal{C}, \varepsilon)$ -degree- k multicalibrated.

Strict Multicalibration To conclude this section, we also give an upper bound on the sample complexity of *strict* multicalibration.

Definition 46 Let $\mathcal{A}_{\mathcal{C},\mathcal{G}}^{\text{SMC}} = \{A_{\mathbf{c},E} \mid \mathbf{c} \in \mathcal{C}^{\mathcal{G}}, E \subseteq \mathcal{Y} \times \mathcal{O} \times \mathcal{G}\}$, where

$$A_{\mathbf{c},E}(j, o, \tilde{p}) = \mathbf{1}[(\mathbf{c}(\hat{p}_j)_j, o, \hat{p}_j) \in E]$$

for each member $j \in \mathcal{X}$, possible outcome $o \in \mathcal{O}$, and predictor $\tilde{p} : \mathcal{X} \rightarrow \Delta\mathcal{O}$.

Theorem 47 Running Algorithm 1 on $\mathcal{A} = \mathcal{A}_{\mathcal{C},\mathcal{G}}^{\text{SMC}}$ and appropriately chosen \mathcal{G}, t, n yields a predictor \tilde{p} such that \hat{p} is strictly $(\mathcal{C}, \varepsilon)$ -multicalibrated with probability at least 99%. The algorithm samples at most

$$tn \lesssim (\log(|\mathcal{C}|) + |\mathcal{Y}||\mathcal{O}|) \cdot \left(\frac{4}{\varepsilon}\right)^{|\mathcal{O}|-1} \cdot \frac{\log(|\mathcal{O}|)}{\varepsilon^4}$$

i.i.d. individual-outcome pairs.

Proof Observe that the family $\mathcal{A}_{\mathcal{C},\mathcal{G}}^{\text{SMC}}$ has size $2^{|\mathcal{Y}||\mathcal{O}||\mathcal{G}||\mathcal{C}|^{|\mathcal{G}|}}$. By Lemma 23, we can choose an η -covering \mathcal{G} of $\Delta\mathcal{O}$ in such a way that $|\mathcal{G}| < (3/\eta)^{\ell-1}$. Thus, by Lemma 8, Algorithm 1 with input $\mathcal{A}_{\mathcal{C},\mathcal{G}}^{\text{SMC}}$ and parameters $\varepsilon/4$ and $\eta = 3\varepsilon/4$ outputs a $(\mathcal{A}_{\mathcal{C},\mathcal{G}}^{\text{SMC}}, \varepsilon/4)$ -outcome-indistinguishable predictor \tilde{p} with probability at least 99% using at most

$$tn \lesssim (\log(|\mathcal{C}|) + |\mathcal{Y}||\mathcal{O}|) \cdot \left(\frac{4}{\varepsilon}\right)^{|\mathcal{O}|-1} \cdot \frac{\log(|\mathcal{O}|)}{\varepsilon^4}$$

samples. Since \mathcal{G} is an η -covering of $\Delta\mathcal{O}$, the inequality $\delta(\hat{o}_i, \tilde{o}_i \mid i) \leq \eta$ holds almost surely, so

$$\mathbf{E} \left[\max_{c \in \mathcal{C}} \delta((c_i, \hat{o}_i), (c_i, o_i^*) \mid \hat{p}_i) \right] \leq \mathbf{E} \left[\max_{c \in \mathcal{C}} \delta((c_i, \tilde{o}_i), (c_i, o_i^*) \mid \hat{p}_i) \right] + \eta.$$

The expectation on the right hand side is at most $\varepsilon/4$ by the definition of statistical distance and $(\mathcal{A}_{\mathcal{C},\mathcal{G}}^{\text{SMC}}, \varepsilon/4)$ -outcome-indistinguishability of \tilde{p} , so the right hand side is at most ε . We conclude that the discretized predictor \hat{p} is strictly $(\mathcal{C}, \varepsilon)$ -multicalibrated. \blacksquare

F.5. Improvements in Special Cases

In the case that $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$, we can refine Theorem 10 so that its bound depends on the *Vapnik–Chervonenkis dimension* $\text{VC}(\mathcal{C})$ of \mathcal{C} instead of the logarithm of the cardinality of \mathcal{C} . To do so, we first state some definitions and lemmas that follow directly from basic properties of VC dimension that can be found in standard texts [Shalev-Shwartz and Ben-David \(2014\)](#).

Definition 48 For finite sets \mathcal{S} and $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{S}}$, the VC dimension $\text{VC}(\mathcal{H})$ of \mathcal{H} is the size of the largest subset $\mathcal{T} \subseteq \mathcal{S}$ such that every possible function $\mathcal{T} \rightarrow \{0, 1\}$ is the restriction of some function in \mathcal{H} . Such a set \mathcal{T} is said to be shattered by \mathcal{H} . Given a family \mathcal{A} of distinguishers and a predictor \tilde{p} , we write $\text{VC}(\mathcal{A}_{\tilde{p}})$ to denote the VC dimension of the collection $\mathcal{A}_{\tilde{p}} \subseteq \{0, 1\}^{\mathcal{X} \times \mathcal{O}}$ of functions $A_{\tilde{p}} : \mathcal{X} \times \mathcal{O} \rightarrow \{0, 1\}$ given by $A_{\tilde{p}}(j, o) = A(j, o, \tilde{p})$ for $A \in \mathcal{A}$.

Lemma 49 *For every family \mathcal{A} of distinguishers and every $\varepsilon > 0$, there exists a weak agnostic learning algorithm $\text{WAL}_{\mathcal{A},\varepsilon}$ with failure probability β that draws at most $O((\text{VC}(\mathcal{A}_{\tilde{p}}) + \log(1/\beta))/\varepsilon^2)$ samples.*

With these tools in hand, we now state and prove the main result of this section.

Theorem 50 *Fix $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$ and let $\ell = |\mathcal{O}|$. There is an algorithm that takes as input*

$$O\left(\text{VC}(\mathcal{C}) + \ell \left(\frac{4}{\varepsilon}\right)^{\ell-1}\right) \cdot \frac{\log(\ell)}{\varepsilon^4}$$

i.i.d. individual-outcome pairs and outputs a strictly $(\mathcal{C}, \varepsilon)$ -multicalibrated predictor w.p. 99%.

Proof Given \mathcal{C} , \mathcal{G} , and $E \subseteq \{0, 1\} \times \mathcal{O} \times \mathcal{G}$, let $\mathcal{A}_{\mathcal{C},\mathcal{G},E}^{\text{MC}} = \{A_{c,E} \mid c \in \mathcal{C}\}$ where

$$A_{c,E}(j, o, \tilde{p}) = \mathbf{1}[(c_j, o, \hat{p}_j) \in E]$$

for each member $j \in \mathcal{X}$, possible outcome $o \in \mathcal{O}$, and predictor $\tilde{p} : \mathcal{X} \rightarrow \Delta\mathcal{O}$. Note that

$$\mathcal{A}_{\mathcal{C},\mathcal{G},\tilde{p}}^{\text{MC}} = \bigcup_E \mathcal{A}_{\mathcal{C},\mathcal{G},E,\tilde{p}}^{\text{MC}},$$

where the union is taken over all $2^{2^{|\mathcal{O}||\mathcal{G}|}}$ possible choices of E . We claim that

$$\text{VC}(\mathcal{A}_{\mathcal{C},\mathcal{G},E,\tilde{p}}^{\text{MC}}) \leq \text{VC}(\mathcal{C})$$

for any E and any predictor \tilde{p} . Once this is shown, it will follow from Lemma 49 with failure probability $\beta/2^{2^{|\mathcal{O}||\mathcal{G}|}}$ and a union bound over the choice of E that there exists a weak agnostic learning algorithm $\text{WAL}_{\mathcal{A}_{\mathcal{C},\mathcal{G},\tilde{p}}^{\text{MC}},\varepsilon}$ with failure probability β that uses only $O((\text{VC}(\mathcal{C}) + \log(2^{2^{|\mathcal{O}||\mathcal{G}|}}/\beta))/\varepsilon^2)$ samples. Proceeding as in the proof of Lemma 45 and Theorem 10 yields the desired sample complexity bound.

It remains to prove the claim. To this end, fix E and \tilde{p} and let $\mathcal{T} = \{(j_1, o_1), \dots, (j_d, o_d)\} \subseteq \mathcal{X} \times \mathcal{O}$ be a maximum size set shattered by $\mathcal{A}_{\mathcal{C},\mathcal{G},E,\tilde{p}}^{\text{MC}}$. Consider the element $(j_d, o_d) \in \mathcal{T}$. In order for there to exist two distinguishers $A_{c,E}$ and $A_{c',E}$ such that $A_{c,E}(j_d, o_d, \tilde{p}) = 0$ and $A_{c',E}(j_d, o_d, \tilde{p}) = 1$, it must be the case that $c(j_d) \neq c'(j_d)$, or else we would have $A_{c,E}(j_d, o_d, \tilde{p}) = A_{c',E}(j_d, o_d, \tilde{p})$. More generally, if the restrictions of $A_{c,E}$ and $A_{c',E}$ to \mathcal{T} are distinct, then the restrictions of c and c' to $\{j_1, \dots, j_d\}$ must also be distinct. Since the distinguishers shatter \mathcal{T} , it follows that $\{j_1, \dots, j_d\} \subseteq \mathcal{X}$ is shattered by \mathcal{C} and hence that $d \leq \text{VC}(\mathcal{C})$. This completes the proof of the claim. \blacksquare

F.6. A Randomized Approach

In this section, we give a variant of Theorem 10 using an alternate algorithm for achieving multicalibration in the case that $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$ that we believe may be of interest. We will present the algorithm in this section with a slightly simpler notion of weak agnostic learning than in the preceding sections. We caution the reader that the sample complexity upper bound we derive in Theorem 53 will not be as tight as that of Theorem 10.

Definition 51 Suppose $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$. Let $\text{WAL}_{\mathcal{C}}$ be an algorithm that takes as input a parameter $\varepsilon > 0$ and a sequence of labeled data $(x_1, y_1), (x_2, y_2), \dots \in \mathcal{X} \times [-1, 1]$, and outputs either a function $c \in \mathcal{C}$ or the symbol \perp . Consider a fixed joint distribution over $\mathcal{X} \times [-1, 1]$, and a pair (x, y) drawn from this distribution. We say that $\text{WAL}_{\mathcal{C}}$ is a weak agnostic learner with failure probability β with respect to this distribution if the following two conditions hold when $\text{WAL}_{\mathcal{C}, \varepsilon}$ is run on input ε and i.i.d. copies of (x, y) :

- If there exists $c \in \mathcal{C}$ such that $\mathbf{E}[c_x y] > \varepsilon$, then $\text{WAL}_{\mathcal{C}}$ outputs $c' \in \mathcal{C}$ such that $\mathbf{E}[c'_x y] > \varepsilon/2$ with probability at least $1 - \beta$.
- If every $c \in \mathcal{C}$ satisfies $\mathbf{E}[c_x y] \leq \varepsilon$, then $\text{WAL}_{\mathcal{C}}$ outputs $c' \in \mathcal{C}$ such that $\mathbf{E}[c'_x y] > \varepsilon/2$ or \perp with probability at least $1 - \beta$.

Lemma 52 Let $\mathcal{A} = \mathcal{A}_{\mathcal{C}, \mathcal{G}}^{\text{MC}}$ and $\varepsilon' = \varepsilon/8\sqrt{|\mathcal{O}||\mathcal{G}|}$. The procedure $\text{SELECT-DISTINGUISHER}(\varepsilon', \tilde{p}, \text{WAL}_{\mathcal{C}})$ in Algorithm 5 samples at most $O((\text{VC}(\mathcal{C}) + \log(1/\beta)) \log(1/\beta)/\varepsilon'^2)$ copies of (i, o_i^*) and its output satisfies the following two properties:

- If there exists $A \in \mathcal{A}$ with $\Delta_A > \varepsilon$, then the procedure returns $A' \in \mathcal{A}$ with $\Delta_{A'} > \varepsilon'/2$ with probability at least $1 - \beta$.
- If every $A \in \mathcal{A}$ satisfies $\Delta_A \leq \varepsilon$, then the procedure returns either $A' \in \mathcal{A}$ with $\Delta_{A'} > \varepsilon'/2$ or \perp with probability at least $1 - \beta$.

Proof Suppose first that there exists a distinguisher $A_{c,E} \in \mathcal{A}$ with advantage $\Delta_{A_{c,E}} > \varepsilon$, where $c \in \mathcal{C}$ and $E \subseteq \{0, 1\} \times \mathcal{O} \times \mathcal{G}$. This means that

$$\Pr[(c_i, \tilde{o}_i, \hat{p}_i) \in E] - \Pr[(c_i, o_i^*, \hat{p}_i) \in E] > \varepsilon,$$

which can also be written as

$$\mathbf{E} [\mathbf{1}[(c_i, \tilde{o}_i, \hat{p}_i) \in E] - \mathbf{1}[(c_i, o_i^*, \hat{p}_i) \in E]] > \varepsilon.$$

It follows that at least one of the two inequalities

$$\mathbf{E} \left[c_i \left(\mathbf{1}[(1, \tilde{o}_i, \hat{p}_i) \in E] - \mathbf{1}[(1, o_i^*, \hat{p}_i) \in E] \right) \right] > \frac{\varepsilon}{2}$$

or

$$\mathbf{E} \left[(1 - c_i) \left(\mathbf{1}[(0, \tilde{o}_i, \hat{p}_i) \in E] - \mathbf{1}[(0, o_i^*, \hat{p}_i) \in E] \right) \right] > \frac{\varepsilon}{2}$$

must hold. If \mathcal{C} is closed under complement, we may assume without loss of generality that the first inequality holds. Using the expression LABEL defined in the algorithm, we may rewrite this inequality as

$$\mathbf{E} [c_i \cdot \text{LABEL}(i, o_i^*, \tilde{p}, E)] > \frac{\varepsilon}{2}.$$

Consider a fixed iteration $s \in [S]$ in the main loop of procedure $\text{SELECT-DISTINGUISHER}$. Using standard anti-concentration inequalities, we will show that

$$\mathbf{E} \left[c_i \cdot \text{LABEL}(i, o_i^*, \tilde{p}, E_s) \mid E_s \right] > \varepsilon'$$

Algorithm 5: Outcome Indistinguishability via Randomized Distinguisher Selection

```

def Label ( $j, o, \tilde{p}, E$ ) :
    |  $o' \sim \tilde{p}_j$  // draw  $o' \in \mathcal{O}$  from the distribution  $\tilde{p}_j$ 
    | return  $\mathbf{1}[(1, o', \hat{p}_j) \in E] - \mathbf{1}[(1, o, \hat{p}_j) \in E]$ 
end

def Select-Distinguisher ( $\varepsilon, \tilde{p}, \text{WAL}_{\mathcal{C}}$ ) :
    |  $S \leftarrow O(\log(1/\beta))$  // set parameters
    |  $T \leftarrow O((\text{VC}(\mathcal{C}) + \log(S/\beta))/\varepsilon^2)$ 
    | for  $s = 1, \dots, S$  do // one-sided error amplification
    | |  $E_s \sim 2^{\{1\} \times \mathcal{O} \times \mathcal{G}}$  // draw  $E_s \subseteq \{1\} \times \mathcal{O} \times \mathcal{G}$  uniformly at random
    | | for  $t = 1, \dots, T$  do
    | | |  $(i_{st}, o_{i_{st}}^*) \sim (i, o_i^*)$  // draw independent copy of  $(i, o_i^*)$ 
    | | |  $(x_{st}, y_{st}) \leftarrow ((i_{st}, o_{i_{st}}^*), \text{Label}(i_{st}, o_{i_{st}}^*, \tilde{p}, E_s))$ 
    | | | end
    | | |  $c_s \leftarrow \text{WAL}_{\mathcal{C}}(\varepsilon, (x_{s1}, y_{s1}), \dots, (x_{sT}, y_{sT}))$  // check for
    | | |  $(\mathcal{C}, \varepsilon)$ -multicalibration violation
    | | | if  $c_s \neq \perp$  then
    | | | | return  $A_{c_s, E_s}$ 
    | | | end
    | | end
    | | return  $\perp$ 
end

def Construct ( $\varepsilon, \tilde{p}, \text{WAL}_{\mathcal{C}}, \text{update}$ ) :
    |  $\varepsilon' \leftarrow \varepsilon/8\sqrt{|\mathcal{O}||\mathcal{G}|}$ 
    |  $A \leftarrow \text{Select-Distinguisher}(\varepsilon', \tilde{p}, \text{WAL}_{\mathcal{C}})$ 
    | if  $A = \perp$  then
    | | return  $\tilde{p}$ 
    | | end
    | | for  $j \in \mathcal{X}$  do
    | | | for  $o \in \mathcal{O}$  do
    | | | |  $L_j(o) \leftarrow A(j, o, \tilde{p})$  // define loss function
    | | | | end
    | | | |  $\tilde{p}'_j \leftarrow \text{update}(\tilde{p}_j, L_j)$  // no-regret update
    | | | | end
    | | | return Construct ( $\varepsilon, \tilde{p}', \text{WAL}_{\mathcal{C}}, \text{update}$ ) // recurse
    | end
end

```

with probability at least $\Omega(1)$ (over the randomness in the choice of $E_s \subseteq \{1\} \times \mathcal{O} \times \mathcal{G}$).

Once this is shown, it will follow that for an appropriate choice of the number of iterations $S = O(\log(1/\beta))$, with probability at least $1 - \beta/2$ over the draws of E_1, \dots, E_S , there will exist some iteration $s \in [S]$ such that the set E_s has the above property.

Additionally, $\text{WAL}_{\mathcal{C}}$ succeeds in all S iterations with probability at least $1 - \beta/2$ since it is fed $T = O((\text{VC}(\mathcal{C}) + \log(S/\beta)/(\varepsilon')^2))$ fresh labeled samples in each iteration.

Consequently, $\text{SELECT-DISTINGUISHER}$ returns a distinguisher A_{c_s, E_s} such that

$$\Delta_{A_{c_s, E_s}} = \mathbf{E} \left[c_{si} \cdot \text{LABEL}(i, o_i^*, \tilde{p}, E_s) \mid c_s, E_s \right] > \frac{\varepsilon}{2}$$

with probability at least $1 - \beta$ over all of the internal randomness of the algorithm (i.e., the draws of E_1, \dots, E_S and the samples $(i_{st}, o_{i_{st}}^*)$ fed to $\text{WAL}_{\mathcal{C}}$).

In the case that every distinguisher $A_{c, E} \in \mathcal{A}$ satisfies $\Delta_{A_{c, E}} \leq 3\sqrt{|\mathcal{O}||\mathcal{G}|}\varepsilon$, the error guarantee of $\text{WAL}_{\mathcal{C}}$ similarly ensures that the procedure returns either $A' \in \mathcal{A}$ with $\Delta_{A'} > \varepsilon/2$ or \perp with probability at least $1 - \beta$.

It remains to prove the claim that the existence of a set $E \subseteq \{0, 1\} \times \mathcal{O} \times \mathcal{G}$ satisfying

$$\mathbf{E}[c_i \cdot \text{LABEL}(i, o_i^*, \tilde{p}, E)] > \frac{\varepsilon}{2}$$

implies that

$$\mathbf{E}[c_i \cdot \text{LABEL}(i, o_i^*, \tilde{p}, E_s) \mid E_s] > \varepsilon'$$

with probability at least $\Omega(1)$ over the draw of a uniformly random $E_s \subseteq \{1\} \times \mathcal{O} \times \mathcal{G}$. To begin, define the sign $\sigma_{E_s}(o, v)$ to be 1 if $(1, o, v) \in E_s$ and -1 if $(1, o, v) \notin E_s$. It is clear that the collection of random variables $\{\sigma_{E_s}(o, v) \mid o \in \mathcal{O}, v \in \mathcal{G}\}$ are independent Rademacher random variables.

Some algebra shows that

$$\mathbf{E}[c_i \text{LABEL}(i, o_i^*, \tilde{p}, E_s) \mid E_s] = \frac{1}{2} \sum_{\substack{o \in \mathcal{O} \\ v \in \mathcal{G}}} \sigma_{E_s}(o, v) \left(\Pr[c_i = 1, \tilde{o}_i = o, \hat{p}_i = v] - \Pr[c_i = 1, o_i^* = o, \hat{p}_i = v] \right),$$

which clearly has mean 0 and standard deviation

$$\frac{1}{2} \left(\sum_{\substack{o \in \mathcal{O} \\ v \in \mathcal{G}}} \left(\Pr[c_i = 1, \tilde{o}_i = o, \hat{p}_i = v] - \Pr[c_i = 1, o_i^* = o, \hat{p}_i = v] \right)^2 \right)^{\frac{1}{2}}.$$

By the Cauchy-Schwarz inequality, this standard deviation is at least

$$\frac{1}{2\sqrt{|\mathcal{O}||\mathcal{G}|}} \sum_{\substack{o \in \mathcal{O} \\ v \in \mathcal{G}}} \left| \Pr[c_i = 1, \tilde{o}_i = o, \hat{p}_i = v] - \Pr[c_i = 1, o_i^* = o, \hat{p}_i = v] \right|,$$

which is at least

$$\frac{1}{2\sqrt{|\mathcal{O}||\mathcal{G}|}} \mathbf{E}[c_i \cdot \text{LABEL}(i, o_i^*, \tilde{p}, E)] > \frac{\varepsilon}{4\sqrt{|\mathcal{O}||\mathcal{G}|}} = 2\varepsilon'.$$

To conclude the proof, we observe any symmetric random variable with mean 0 and standard deviation $2\varepsilon'$ must exceed ε' with probability $\Omega(1)$ (e.g., by the Paley-Zygmund inequality). \blacksquare

Theorem 53 Fix $\mathcal{C} \subseteq \{0, 1\}^{\mathcal{X}}$ and let $\ell = |\mathcal{C}|$. There is a constant $c > 0$ and an algorithm that takes as input $(c/\varepsilon)^{2\ell+c} \cdot \text{VC}(\mathcal{C})$ i.i.d. individual-outcome pairs and outputs a $(\mathcal{C}, \varepsilon)$ -multicalibrated predictor w.p. 99% while only accessing \mathcal{C} through calls to $\text{WAL}_{\mathcal{C}}$.

Proof The proof of Theorem 42 with a step size proportional to ε' , along with the two properties in Lemma 52, gives a $O(\log(\ell)/\varepsilon'^2)$ upper bound on the number of updates. Lemma 52 also upper bounds the number of samples needed per update by $O(\text{VC}(\mathcal{C})/\varepsilon'^2)$. Choosing β appropriately and applying a union bound over the entire sequence of updates and substituting $\varepsilon' = \Theta(\varepsilon/\sqrt{|\mathcal{C}||\mathcal{G}|})$ yields the claimed sample complexity upper bound. ■

Appendix G. Properties of Intermediate Regularity

This section is devoted to the following two results, which establish the strict separation of our notion of intermediate regularity from Szemerédi regularity and from Frieze-Kannan regularity.

Theorem 54 There is an absolute constant $c \in (0, 1)$ such that for all sufficiently small $\varepsilon > 0$:

- For any graph G , if the vertex partition \mathcal{P} satisfies intermediate ε -regularity, then \mathcal{P} satisfies Frieze-Kannan ε^c -regularity.
- There exists a graph G and a vertex partition \mathcal{P} satisfying Frieze-Kannan ε -regularity but not intermediate c -regularity.

Theorem 55 There is an absolute constant $c \in (0, 1)$ such that for all sufficiently small $\varepsilon > 0$:

- For any graph G , if the vertex partition \mathcal{P} satisfies Szemerédi ε -regularity, then \mathcal{P} satisfies intermediate ε^c -regularity.
- There exists a graph G such that any vertex partition \mathcal{P} satisfying intermediate ε -regularity does not satisfy Szemerédi $\frac{c}{\sqrt{\log^*(1/\varepsilon)}}$ -regularity.

In order to prove Theorems 54 and 55, it will be useful to introduce an alternative characterization of Szemerédi regularity, based on the notion of *irregularity*:

Definition 56 Let $X, Y \subseteq V$. The irregularity of the pair (X, Y) is

$$\text{irreg}(X, Y) = \max_{\substack{S \subseteq X \\ T \subseteq Y}} |e(S, T) - d(X, Y)|S||T||.$$

Specifically, it is known that Szemerédi ε -regularity (Definition 14) is equivalent to having irregularity at most $\varepsilon|V|^2$, up to a polynomial change in ε , where the irregularity of a partition is defined as follows:

Definition 57 The irregularity of a partition $\mathcal{P} = \{V_1, \dots, V_m\}$ of V is

$$\text{irreg}(\mathcal{P}) = \sum_{j, k \in [m]} \text{irreg}(V_j, V_k).$$

For more on this equivalence, we refer the reader to [Skorski \(2017\)](#). One can state an alternate version of the definition of intermediate regularity ([Definition 17](#)) that is equivalent up to a polynomial change in the ε parameter. Specifically, one can check that intermediate ε -regularity is equivalent to having (S, T) -irregularity at most $\varepsilon|V|^2$ for all $S, T \subseteq V$, up to a polynomial change in ε , where the (S, T) -irregularity of a partition is defined as follows:

Definition 58 *Let $X, Y, S, T \subseteq V$. The (S, T) -irregularity of the pair (X, Y) is*

$$\text{irreg}_{S,T}(X, Y) = \left| e(S \cap X, T \cap Y) - d(X, Y) |S \cap X| |T \cap Y| \right|.$$

Definition 59 *The (S, T) -irregularity of a partition $\mathcal{P} = \{V_1, \dots, V_m\}$ of V is*

$$\text{irreg}_{S,T}(\mathcal{P}) = \sum_{j,k \in [m]} \text{irreg}_{S,T}(V_j, V_k).$$

We now formally state and prove the claimed equivalence between [Definitions 17](#) and [59](#):

Theorem 60 *If \mathcal{P} satisfies intermediate ε -regularity, then \mathcal{P} has (S, T) -irregularity at most $2\varepsilon|V|^2$ for all $S, T \subseteq V$. Conversely, if \mathcal{P} has (S, T) -irregularity at most $\varepsilon|V|^2$ for all $S, T \subseteq V$, then \mathcal{P} satisfies intermediate $\sqrt{\varepsilon}$ -regularity.*

Proof To prove the forward direction, suppose that the partition $\mathcal{P} = \{V_1, \dots, V_m\}$ satisfies intermediate ε -regularity. If (V_j, V_k) is an (S, T, ε) -regular pair, then

$$\text{irreg}_{S,T}(V_j, V_k) \leq \varepsilon |S \cap V_j| |T \cap V_k|.$$

If (V_j, V_k) is not an (S, T, ε) -regular pair, we only have the bound

$$\text{irreg}_{S,T}(V_j, V_k) \leq |S \cap V_j| |T \cap V_k|.$$

Consequently,

$$\text{irreg}_{S,T}(\mathcal{P}) \leq \sum_{\substack{j,k \in [m] \\ (V_j, V_k) \text{ is } (S,T,\varepsilon)\text{-regular}}} \varepsilon |S \cap V_j| |T \cap V_k| + \sum_{\substack{j,k \in [m] \\ (V_j, V_k) \text{ not } (S,T,\varepsilon)\text{-regular}}} |S \cap V_j| |T \cap V_k|.$$

The first sum is clearly at most $\varepsilon|V|^2$, and the second sum is at most $\varepsilon|V|^2$ by [Definition 17](#).

To prove the converse direction, suppose that \mathcal{P} has (S, T) -irregularity at most $\varepsilon|V|^2$ for all $S, T \subseteq V$. If a pair (V_j, V_k) is not $(S, T, \sqrt{\varepsilon})$ -regular, then $\text{irreg}_{S,T}(V_j, V_k) \geq \sqrt{\varepsilon} |S \cap V_j| |T \cap V_k|$. It follows that the partition \mathcal{P} under consideration satisfies

$$\varepsilon|V|^2 \geq \text{irreg}_{S,T}(\mathcal{P}) \geq \sum_{\substack{j,k \in [m] \\ (V_j, V_k) \text{ not } (S,T,\sqrt{\varepsilon})\text{-regular}}} \sqrt{\varepsilon} |S \cap V_j| |T \cap V_k|,$$

and dividing both sides by $\sqrt{\varepsilon}$ allows us to conclude that \mathcal{P} satisfies intermediate $\sqrt{\varepsilon}$ -regularity. ■

With these alternative characterizations of Szemerédi regularity and intermediate regularity in hand, we are ready to prove [Theorems 54](#) and [55](#), establishing the strict separation of intermediate regularity from prior notions.

Proof of Theorem 54 For the first part, it suffices to show that \mathcal{P} satisfies Frieze-Kannan ε -regularity if it has (S, T) -irregularity at most $\varepsilon|V|^2$ for all $S, T \subseteq V$. To this end, fix $S, T \subseteq V$ and observe that

$$\left| e(S, T) - \sum_{j,k \in [m]} d_{jk} |S \cap V_j| |T \cap V_k| \right| = \left| \sum_{j,k \in [m]} (e(S \cap V_j, T \cap V_k) - d_{jk} |S \cap V_j| |T \cap V_k|) \right|,$$

where $d_{jk} = d(V_j, V_k)$. By the triangle inequality, the right hand side is at most

$$\sum_{j,k \in [m]} \left| e(S \cap V_j, T \cap V_k) - d(V_j, V_k) |S \cap V_j| |T \cap V_k| \right| = \text{irreg}_{S,T}(\mathcal{P}) \leq \varepsilon|V|^2.$$

For the second part, let $G = (V, E)$ be any graph with the property that

$$\left| e_G(S, T) - \frac{1}{2}|S||T| \right| < \varepsilon|V|^2$$

for all $S, T \subseteq V$. The existence of such a *quasirandom* graph of density $1/2$ follows from standard probabilistic arguments, but explicit constructions are also known [Zhao \(2022\)](#). We will now modify the graph G into a graph G' with a partition \mathcal{P} satisfying Frieze-Kannan ε -regularity but not intermediate $1/9$ -regularity. To do so, let $G' = (V', E')$ be the graph with $V' = V \times [2]$ and

$$E' = \{((v_1, b_1), (v_2, b_2)) \mid (v_1, v_2) \in E \text{ or } b_1 \neq b_2 \text{ but not both}\}.$$

This graph G' can be realized as the *Xor product* [Alon and Lubetzky \(2007\)](#) of G with a graph consisting of a single edge. We claim that the partition $\mathcal{P} = \{\{(v, 1), (v, 2)\} \mid v \in V\}$ of V' has the desired properties. To check Frieze-Kannan ε -regularity, observe that the density of edges from any part of \mathcal{P} to another is precisely $1/2$, so it suffices to show that

$$\left| e_{G'}(S', T') - \frac{1}{2}|S'||T'| \right| \leq \varepsilon|V'|^2$$

for any $S', T' \subseteq V'$. To this end, for $S' \subseteq V'$ and $b \in [2]$, let $S'_b = \{v \in V \mid (v, b) \in S'\}$. Then

$$\begin{aligned} \left| e_{G'}(S', T') - \frac{1}{2}|S'||T'| \right| &= \left| |E \cap (S'_1 \times T'_1)| - \frac{1}{2}|S'_1||T'_1| \right| \\ &\quad + \left| |(V^2 \setminus E) \cap (S'_1 \times T'_2)| - \frac{1}{2}|S'_1||T'_2| \right| \\ &\quad + \left| |(V^2 \setminus E) \cap (S'_2 \times T'_1)| - \frac{1}{2}|S'_2||T'_1| \right| \\ &\quad + \left| |E \cap (S'_2 \times T'_2)| - \frac{1}{2}|S'_2||T'_2| \right|, \end{aligned}$$

and, by our initial choice of G , each of the four terms on the right hand is at most $\varepsilon|V|^2 = \varepsilon|V'|^2/4$. To check that intermediate $1/9$ -regularity fails, let $S' = T' = \{(v, b) \in V' \mid b = 1\}$. Then

$$\text{irreg}_{S',T'}(\mathcal{P}) = \sum_{v_1, v_2 \in V} \left| \mathbf{1}[(v_1, v_2) \in E] - \frac{1}{2} \right| = \frac{1}{2}|V|^2 = \frac{1}{8}|V'|^2.$$

Proof of Theorem 55 For the first part, it suffices by Theorem 60 to show that the (S, T) -irregularity of \mathcal{P} is bounded above by its irregularity. To this end, observe that for any $V_j, V_k \in \mathcal{P}$, we have that

$$\text{irreg}(V_j, V_k) = \max_{S, T \subseteq V} \text{irreg}_{S, T}(V_j, V_k).$$

Therefore, for any particular $S, T \subseteq V$, we have that

$$\text{irreg}_{S, T}(\mathcal{P}) = \sum_{j, k \in [m]} \text{irreg}_{S, T}(V_j, V_k) \leq \sum_{j, k \in [m]} \text{irreg}(V_j, V_k) = \text{irreg}(\mathcal{P}).$$

The second part follows readily from the following two facts. The first fact is a lower bound on the number of parts required to achieve Szemerédi ε -regularity. Specifically, Fox and Lovász (2017) showed that there exists a graph for which every vertex partition \mathcal{P} with

$$\text{irreg}(\mathcal{P}) \leq \varepsilon |V|^2$$

requires the number of parts $|\mathcal{P}|$ to be at least a tower of twos of height $\Omega(1/\varepsilon^2)$. The second fact is an upper bound on the number of parts required to achieve intermediate ε -regularity. Specifically, we will argue in Appendix H that every graph has a vertex partition \mathcal{P} with

$$\max_{S, T \subseteq V} \text{irreg}_{S, T}(\mathcal{P}) \leq \varepsilon |V|^2$$

and $|\mathcal{P}| \leq 4^{1/\varepsilon^2}$. Comparing these upper and lower bounds yields the claimed separation between intermediate and Szemerédi regularity.

Appendix H. Algorithm for Intermediate Regularity

Theorem 19 suggests that intermediate regularity might be achievable via a modified multicalibration algorithm. This is indeed the case, and standard analyses show that Algorithm 6, initialized with the trivial partition $\mathcal{P} = \{V\}$, computes such a partition with complexity summarized by the second row of Table 1. The algorithm can also be viewed as a modification of a standard algorithm for Frieze-Kannan regularity. The SELECT subroutine of Algorithm 6 implements the algorithm from Alon and Naor (2004) that takes as input a function $f : V \times V \rightarrow [-1, 1]$ and outputs $S, T \subseteq V$ such that $\left| \sum_{(u, v) \in S \times T} f(u, v) \right| > \frac{1}{2} \max_{S', T' \subseteq V} \left| \sum_{(u, v) \in S' \times T'} f(u, v) \right|$ in $\text{poly}(|V|)$ time.

Regularity Notion	Number of Parts	Time Complexity ⁷
Frieze-Kannan	$\exp(\text{poly}(1/\varepsilon))$	$\text{poly}(n/\varepsilon)$
Intermediate	$\exp(\text{poly}(1/\varepsilon))$	$\text{poly}(n) \exp(\exp(\text{poly}(1/\varepsilon)))$
Szemerédi	$\text{tower}(\text{poly}(1/\varepsilon))$	$\text{poly}(n) \text{tower}(\text{poly}(1/\varepsilon))$

Table 1: Some Upper Bounds on Achieving Graph Regularity ($n = |V|$, $\text{tower}(h) = \underbrace{2^{2^{\cdot^{\cdot^2}}}}_{h \text{ times}}$)

At this point, it is instructive to compare Algorithm 6, which achieves intermediate regularity, to the two algorithms of Skorski (2017), which target Frieze-Kannan and Szemerédi regularity.

⁷ The time complexity bound for Frieze-Kannan regularity is for computing an implicit representation of the partition.

Algorithm 6: Construction of Intermediate ε -Regularity Partition

```

def REFINE ( $\varepsilon, G, \mathcal{P}$ ) :
    for  $\sigma \in \{-1, +1\}^{m \times m}$  do
         $S, T \leftarrow$  SELECT ( $\sum_{j=1}^m \sum_{k=1}^m \sigma_{jk} (\mathbf{1}_E - d(V_j, V_k)) \mathbf{1}_{V_j \times V_k}$ )
        if  $\text{irreg}_{S,T}(\mathcal{P}) > \frac{1}{2} \varepsilon |V|^2$  then
             $\mathcal{P}' \leftarrow$  common refinement of  $S, T$ , and  $\mathcal{P}$ 
            return REFINE ( $\varepsilon, G, \mathcal{P}'$ )
        end
    end
    return  $\mathcal{P}$ 
end
    
```

Indeed, using the informal notation of Section 1.3 of Skorski (2017), the three algorithms can be viewed as computing low-complexity approximations $\mathcal{P} = \{V_1, \dots, V_k\}$ to the graph with respect to the following families of test functions:

$$\begin{aligned} \mathcal{F} &= \{f : f = \pm \mathbf{1}_{T \times S}\} && \text{(for Frieze-Kannan Regularity)} \\ \mathcal{F} &= \left\{ f : f = \sum_{j,k} \pm \mathbf{1}_{(T \cap V_j) \times (S \cap V_k)} \right\} && \text{(for Intermediate Regularity)} \\ \mathcal{F} &= \left\{ f : f = \sum_{j,k} \pm \mathbf{1}_{T_{j,k} \times S_{j,k}} \right\} && \text{(for Szemerédi Regularity)} \end{aligned}$$

Above, the sums are taken over $j, k \in [m]$, and we consider all $T \subseteq V$, all $S \subseteq V$, all $T_{j,k} \subseteq V_j$, all $S_{j,k} \subseteq V_k$, and the \pm symbol is understood to be a factor of $+1$ or -1 , possibly different for each term in the sum.

Intuitively, these three families of functions are testing for multiaccuracy, multicalibration, and strict multicalibration with respect to the *same* collection of subpopulations, namely $\mathcal{C} = \{T \times S : T \subseteq V, S \subseteq V\}$, and this can be seen by viewing the products $V_j \times V_k$ of parts of the partition \mathcal{P} as playing the role of “level sets of \tilde{p} ” in the language of multi-group fairness.

Appendix I. Proof of Theorem 19

Throughout this proof, we will use the notation d_{jk} and $\Delta_{S,v}$ from the proof sketch in Appendix 6.2, as well as the alternate characterizations of Szemerédi and intermediate regularity from Appendix G.

- (a) As usual, let V_1, \dots, V_m denote the parts of \mathcal{P} , and consider any fixed sets $S, T \subseteq V$ and value $v \in [0, 1]$. The construction of \tilde{p} ensures that $\tilde{p}_{(a,b)} = d(V_j, V_k)$ for any vertex pair $(a, b) \in V_j \times V_k$, so some algebra yields

$$\Delta_{S \times T, v}(\tilde{p}) = \sum_{\substack{(j,k) \in [m]^2 \text{ s.t.} \\ \tilde{p}(V_j \times V_k) = v}} e(S \cap V_j, T \cap V_k) - d(V_j, V_k) |S \cap V_j| |T \cap V_k|,$$

along with the useful fact that $\Delta_{\mathcal{X} \setminus (S \times T), v}(\tilde{p}) = -\Delta_{S \times T, v}(\tilde{p})$. By taking absolute values, summing over v in the range of \tilde{p} , and taking the max over $S, T \subseteq V$ in various orders on both sides of the above equation, we deduce the following three inequalities:

$$\begin{aligned} \max_{S, T \subseteq V} \left| \sum_{v \in \tilde{p}(\mathcal{X})} \Delta_{S \times T, v}(\tilde{p}) \right| &\leq \max_{S, T \subseteq V} \left| \sum_{j=1}^m \sum_{k=1}^m e(S \cap V_j, T \cap V_k) - d_{jk} |S \cap V_j| |T \cap V_k| \right|, \\ \max_{S, T \subseteq V} \sum_{v \in \tilde{p}(\mathcal{X})} |\Delta_{S \times T, v}(\tilde{p})| &\leq \max_{S, T \subseteq V} \sum_{j=1}^m \sum_{k=1}^m |e(S \cap V_j, T \cap V_k) - d_{jk} |S \cap V_j| |T \cap V_k||, \\ \sum_{v \in \tilde{p}(\mathcal{X})} \max_{S, T \subseteq V} |\Delta_{S \times T, v}(\tilde{p})| &\leq \sum_{j=1}^m \sum_{k=1}^m \max_{S, T \subseteq V} |e(S \cap V_j, T \cap V_k) - d_{jk} |S \cap V_j| |T \cap V_k||. \end{aligned}$$

The above three inequalities, show, respectively, that for an appropriate absolute constant $c \in (0, 1)$ and sufficiently small $\varepsilon > 0$, Frieze-Kannan ε -regularity of \mathcal{P} implies ε^c -multiaccuracy of \tilde{p} , that intermediate ε -regularity of \mathcal{P} implies ε^c -multicalibration of \tilde{p} , and that Szemerédi ε -regularity of \mathcal{P} implies strict ε^c -multicalibration of \tilde{p} .

- (b) For a fixed set $S \subseteq V$ of vertices, let $\tilde{S} \subseteq V$ denote a random set of vertices sampled as follows: for each $j \in [m]$, independently include all vertices of V_j in \tilde{S} with probability $|S \cap V_j|/|V_j|$. With this notation, a simple algebraic calculation shows that for any fixed $j, k \in [m]$ such that $V_j \times V_k$ exactly coincides with the v -level set of \tilde{p} , and for any fixed sets $S, T \subseteq V$, we have

$$e(S \cap V_j, T \cap V_k) - d(V_j, V_k) |S \cap V_j| |T \cap V_k| = \Delta_{S \times T, v}(\tilde{p}) - \mathbf{E}_{\tilde{S}, \tilde{T}} \left[\Delta_{\tilde{S} \times \tilde{T}, v}(\tilde{p}) \right].$$

We will manipulate this key equation in three different ways to derive the three versions of this part of the theorem. First, by summing over $j, k \in [m]$, taking absolute values, and applying the triangle inequality, we see that

$$\left| e(S, T) - \sum_{j, k \in [m]} d_{jk} |S \cap V_j| |T \cap V_k| \right| \leq |\Delta_{S \times T}(\tilde{p})| + \mathbf{E}_{\tilde{S}, \tilde{T}} \left| \sum_{v \in \tilde{p}(\mathcal{X})} \Delta_{\tilde{S} \times \tilde{T}, v}(\tilde{p}) \right|.$$

For small enough $\varepsilon, c > 0$, this shows that ε -multiaccuracy of \tilde{p} implies Frieze-Kannan ε^c -regularity of \mathcal{P} . If we were to instead take absolute values of both sides before summing over $j, k \in [m]$ and applying the triangle inequality, we would see that

$$\sum_{j, k \in [m]} \left| e(S \cap V_j, T \cap V_k) - d_{jk} |S \cap V_j| |T \cap V_k| \right| \leq \sum_{v \in \tilde{p}(\mathcal{X})} |\Delta_{S \times T, v}(\tilde{p})| + \sum_{v \in \tilde{p}(\mathcal{X})} |\Delta_{V \times V, v}(\tilde{p})|,$$

which shows that ε -multicalibration of \tilde{p} implies intermediate ε^c -regularity of \mathcal{P} . Finally, if we had chosen to take the maximum over S, T before summing over $j, k \in [m]$ and applying the triangle inequality, we would have seen that

$$\sum_{j, k \in [m]} \text{irreg}(V_j, V_k) \leq \sum_{v \in \tilde{p}(\mathcal{X})} \max_{S, T \subseteq V} |\Delta_{S \times T, v}(\tilde{p})| + \sum_{v \in \tilde{p}(\mathcal{X})} |\Delta_{V \times V, v}(\tilde{p})|,$$

which shows that strict ε -multicalibration of \tilde{p} implies Szemerédi ε^c -regularity of \mathcal{P} .

Appendix J. Proof of Theorem 20

We first prove Lemma 62, which states that we can obtain a relatively simple predictor satisfying both (a variant of) statistical-distance multicalibration and the additional guarantee that the predictor is perfectly accurate in expectation on all its slices. This is obtained by a careful analysis of a simple post-processing of the predictor obtained by our multicalibration-through-outcome-indistinguishability Algorithm 1 when instantiated with a suitable class of distinguishers.

The change is that instead of requiring indistinguishability w.r.t. the real valued c 's (which would require close to perfect accuracy), we ask for indistinguishability with respect to *outcomes* when treating c_i as a probability. That is, instead of bounding $\delta((c_i, \check{o}_i, \hat{p}_i), (c_i, o_i^*, \hat{p}_i))$, we bound $\delta((o_i^c, \check{o}_i, \hat{p}_i), (o_i^c, o_i^*, \hat{p}_i)) < \varepsilon$, where $o_i^c \sim \text{Ber}(c_i)$.

We'll also need a new family of distinguishers.

Definition 61 Let $\mathcal{B}_{\mathcal{C}, \mathcal{G}}^{\text{MC}} = \{B_{c,E} \mid c \in \mathcal{C}, E \subseteq \mathcal{Y} \times \mathcal{O} \times \mathcal{G}\}$, for $\mathcal{C} \in \Delta \mathcal{Y}^{\mathcal{X}}$, where

$$B_{c,E}(j, o, \tilde{p}) = \mathbf{1}[(o_j^c \leftarrow c_j, o, \hat{p}_j) \in E]$$

for each member $j \in \mathcal{X}$, possible outcome $o \in \mathcal{O}$, and predictor $\tilde{p} : \mathcal{X} \rightarrow \Delta \mathcal{O}$.

Note that Definition 9 is a special case of Definition 61 for trivial distributions.

Lemma 62 For every $\varepsilon > 0, \eta > 0$, there is a function $\check{p} \in R_{T(\varepsilon)}(\mathcal{C})$, $\check{p} : \mathcal{X} \rightarrow [0, 1]$ and a partitioning function $\hat{p} : \mathcal{X} \rightarrow \mathcal{G}$, with \mathcal{G} an η -cover of $[0, 1]$ s.t.

1. for all $c \in \mathcal{C}$, $\delta((o_i^c, \check{o}_i, \hat{p}_i), (o_i^c, o_i^*, \hat{p}_i)) < \varepsilon$.
2. \check{p} is perfectly accurate in expectation on the slices $\hat{p} = v$.
3. $T(\varepsilon) = O\left(\frac{1}{\varepsilon^2}\right)$

Proof Run Algorithm 1 using the collection of distinguishers $\mathcal{B}_{\mathcal{C}, \mathcal{G}}^{\text{MC}}$, where \mathcal{C} is as provided in the statement of the theorem and \mathcal{G} is the standard η covering, to obtain \tilde{p} that is $\left(\mathcal{B}_{\mathcal{C}, \mathcal{G}}^{\text{MC}}, \frac{\varepsilon}{2}\right)$ -outcome-indistinguishable from p^* and has relative complexity $T(\varepsilon) = O(\varepsilon^{-2})$ to \mathcal{C} . This gives us a predictor satisfying

$$\delta((o_i^c, \check{o}_i, \hat{p}_i), (o_i^c, o_i^*, \hat{p}_i)) < \frac{\varepsilon}{2}$$

However, it does not necessarily satisfy Claim 2; we will modify \tilde{p} to obtain a new predictor \check{p} that satisfies both Claims 1 and 2. This is achieved by shifting \tilde{p} on the level sets of \hat{p} , incurring an additive term of $|\mathcal{G}|$ on the complexity of \hat{p} . Speaking intuitively, Claim 1 still holds for this new \hat{p} since we are only improving the accuracy of \hat{p} ; we prove this intuition to be (nearly) correct. More formally, define

$$\check{p}_i := \tilde{p}_i + \tau(\hat{p}_i),$$

where

$$\tau_v := \mathbf{E}[p_i^* - \tilde{p}_i \mid \hat{p}_i = v].$$

\check{p}_i satisfies Claim 1 by construction. Finally, it remains to show that Claim 1 still holds for this new \check{p} . First, we show that the average magnitude of these shifts must be small:

$$\begin{aligned} \mathbf{E}_v[|\tau_v|] &= \mathbf{E}_v[|\mathbf{E}[p_i^* - \tilde{p}_i \mid \hat{p}_i = v]|] \\ &\leq \delta((o_i^c, \check{o}_i, \hat{p}_i), (o_i^c, o_i^*, \hat{p}_i)) \end{aligned}$$

We show that this implies that shifting at most doubles the statistical difference.

$$\begin{aligned}
 \delta((o_i^c, \check{o}_i, \hat{p}_i), (o_i^c, o_i^*, \hat{p}_i)) &= \mathbf{E}_y[\delta((\check{o}_i, \hat{p}_i), (o_i^*, \hat{p}_i) | o_i^c = y)] \\
 &= \mathbf{E}_y \left[\mathbf{E}_v \left[\left| \mathbf{E}_i[p_i^* - \check{p}_i | \hat{p}_i = v] \right| | o_i^c = y \right] \right] \\
 &= \mathbf{E}_y \left[\mathbf{E}_v \left[\left| \mathbf{E}_i[p_i^* - \tilde{p}_i - \tau_v | \hat{p}_i = v] \right| | o_i^c = y \right] \right] \\
 &\leq \mathbf{E}_y \left[\mathbf{E}_v \left[\left(\left| \mathbf{E}_i[p_i^* - \tilde{p}_i | \hat{p}_i = v] \right| + |\tau_v| \right) | o_i^c = y \right] \right]
 \end{aligned}$$

Applying the above fact,

$$\begin{aligned}
 &\leq \mathbf{E}_y \left[\mathbf{E}_v \left[\left| \mathbf{E}_i[p_i^* - \tilde{p}_i | \hat{p}_i = v] \right| | o_i^c = y \right] \right] + \delta((o_i^c, \tilde{o}_i, \hat{p}_i), (o_i^c, o_i^*, \hat{p}_i)) \\
 &= 2\delta((o_i^c, \tilde{o}_i, \hat{p}_i), (o_i^c, o_i^*, \hat{p}_i)) \\
 &\leq \varepsilon.
 \end{aligned}$$

■

We will continue to use \mathcal{G} to denote the set of rounded to values/slices for a predictor. In this section, we will always use \mathcal{G} of the form $(0, \eta, \dots, \lceil \frac{1}{\eta} \rceil \cdot \eta, 1)$, so that $m = |\mathcal{G}| = \lceil \frac{1}{\eta} \rceil + 1$. For clarity, we will assume from now on that $\eta = 1/m$ to avoid having to clutter the presentation with rounding.

The next lemma lets us find large slices on which the target function has large variance, under the assumption that \tilde{p} fails to α -approximate p^* .

Lemma 63 *Let $\check{p} : [0, 1]^{\mathcal{X}}$ and $\hat{p} : \mathcal{G}^{\mathcal{X}}$. If $\mathbf{E}[(\check{p}_i - p_i^*)^2] > \alpha$, then there exists $v \in \mathcal{G}$ s.t.*

1. $\Pr[\hat{p}_i = v] \geq 2\alpha\eta$
2. and $\mathbf{Var}(p_i^* | \hat{p}_i = v) \geq \frac{\alpha}{2} - 3\eta$

Proof First, we show that we can approximately break down $\mathbf{E}[(\check{p}_i - p_i^*)^2] > \alpha$ as the average of the variance of p^* on the slices defined by \hat{p} .

$$\begin{aligned}
 \alpha &< \mathbf{E}[(\check{p}_i - p_i^*)^2] \\
 &= \sum_{v \in \mathcal{G}} \Pr[\hat{p}_i = v] \mathbf{E}[(\check{p}_i - p_i^*)^2 | \hat{p}_i = v]
 \end{aligned}$$

Letting μ_v denote the expected value of p_i^* on slice v which, by construction, is equivalent to that of \check{p}_i ,

$$\begin{aligned}
 &= \sum_{v \in \mathcal{G}} \Pr[\hat{p}_i = v] \mathbf{E}[(\check{p}_i - \mu_v + \mu_v - p_i^*)^2 | \hat{p}_i = v] \\
 &= \sum_{v \in \mathcal{G}} \Pr[\hat{p}_i = v] \mathbf{E}[(p_i^* - \mu_v)^2 - 2(p_i^* - \mu_v)(\check{p}_i - \mu_v) + (\check{p}_i - \mu_v)^2] \\
 &= \sum_{v \in \mathcal{G}} \Pr[\hat{p}_i = v] (\mathbf{Var}(p_i^* | \hat{p}_i = v) - 2 \mathbf{Cov}(p_i^*, \check{p}_i | \hat{p}_i = v) + \mathbf{Var}(\check{p}_i | \hat{p}_i = v))
 \end{aligned}$$

Since \check{p}_i is bounded within an interval of size η ,

$$\leq \sum_{v \in \mathcal{G}} \Pr[\hat{p}_i = v] (\mathbf{Var}(p_i^* | \hat{p}_i = v) + 3\eta).$$

Next, we apply a Markov type argument to show that the total probability of all the slices with large p^* variance must be at least $\alpha/2$. Let \mathcal{B} denote the slices on which $\mathbf{Var}(p_i^* | \hat{p}_i = v) \geq \frac{\alpha}{2} - 3\eta$.

$$\begin{aligned}
 &\sum_{v \in \mathcal{G}} \Pr[\hat{p}_i = v] \mathbf{Var}(p_i^* | \hat{p}_i = v) \\
 &= \sum_{v \in \mathcal{B}} \Pr[\hat{p}_i = v] \mathbf{Var}(p_i^* | \hat{p}_i = v) + \sum_{v \in \mathcal{G} - \mathcal{B}} \Pr[\hat{p}_i = v] \mathbf{Var}(p_i^* | \hat{p}_i = v)
 \end{aligned}$$

Since p^* is bounded in $[0, 1]$, it has variance at most $1/4$,

$$\begin{aligned}
 &\leq \sum_{v \in \mathcal{B}} \Pr[\hat{p}_i = v] \cdot \frac{1}{4} + \sum_{v \in \mathcal{G} - \mathcal{B}} \Pr[\hat{p}_i = v] \left(\frac{\alpha}{2} - 3\eta \right) \\
 &= \frac{1}{4} \Pr[\mathcal{B}] + \frac{\alpha}{2} - 3\eta.
 \end{aligned}$$

Combining with the first inequality, we conclude that

$$\Pr[\mathcal{B}] \geq 2\alpha$$

And because \mathcal{G} is the uniform η covering, we conclude that there exists a $v \in \mathcal{B}$ satisfying $\Pr[\hat{p}_i = v] \geq 2\alpha\eta$. \blacksquare

The final two lemmas let us show that if there is no hardcore, then statistical distance must be large.

Lemma 64 $\delta((o_i^c, \check{o}_i), (o_i^c, o_i^*) | \hat{p}_i) \geq \mathbf{Cov}(c_i, p_i^* | \hat{p}_i) - \eta$.

Proof We first show that $\mathbf{Cov}(c_i, p_i^* | \hat{p}_i) = \mathbf{E}[o_i^c(o_i^* - \hat{o}_i) | \hat{p}_i]$.

$$\begin{aligned}
 \mathbf{Cov}(c_i, p_i^* | \hat{p}_i) &= \mathbf{E}[c_i(p_i^* - \hat{p}_i) | \hat{p}_i] \\
 &= \mathbf{E}[o_i^c(o_i^* - \hat{o}_i) | \hat{p}_i]
 \end{aligned}$$

Next, we show that $\mathbf{E}[o_i^c(o_i^* - \hat{o}_i)|\hat{p}_i] \leq \delta((o_i^c, \hat{o}_i, \hat{p}_i), (o_i^c, o_i^*, \hat{p}_i) | \hat{p}_i)$.

$$\begin{aligned} \mathbf{E}[o_i^c(o_i^* - \hat{o}_i)|\hat{p}_i] &= \mathbf{E}\left[\left(o_i^c - \frac{1}{2}\right)(o_i^* - \hat{o}_i)|\hat{p}_i\right] \\ &= \mathbf{E}\left[\frac{1}{2}(\mathbf{Pr}[o_i^c = 1](\mathbf{Pr}[o_i^* = 1] - \mathbf{Pr}[\tilde{o}_i = 1]) + \mathbf{Pr}[o_i^c = 0](\mathbf{Pr}[o_i^* = 0] - \mathbf{Pr}[\tilde{o}_i = 0]))\right] \\ &= \frac{1}{2}(\mathbf{Pr}[o_i^c = o_i^*] - \mathbf{Pr}[o_i^c - \tilde{o}_i]) \\ &\leq \delta((o_i^c, \hat{o}_i), (o_i^c, o_i^*) | \hat{p}_i) \end{aligned}$$

Putting the two together, we get

$$\delta((o_i^c, \hat{o}_i), (o_i^c, o_i^*) | \hat{p}_i) \geq \mathbf{Cov}(c_i, p_i^* | \hat{p}_i)$$

To conclude the proof, observe that $\delta((c_i, o_i^*), (c_i, \hat{o}_i) | \hat{p}_i) \leq \delta((c_i, o_i^*), (c_i, \check{o}_i) | \hat{p}_i) + \eta$. \blacksquare

Lemma 65 *If there exists $v^* \in \mathcal{G}$ satisfying:*

1. $\mathbf{Pr}[\hat{p}_i = v^*] > \beta$
2. $\mathbf{Var}(p_i^* | \hat{p}_i = v^*) > \lambda$
3. $\mathbf{Cov}(c, p^* | \hat{p}_i = v^*) > \gamma \mathbf{Var}(p_i^* | \hat{p}_i = v^*)$

Then

$$\delta((c_i, \check{o}_i, \hat{p}_i), (c_i, o_i^*, \hat{p}_i)) > \beta\gamma\lambda - \beta\eta$$

Proof This lemma is a simple consequence of the Lemma 64 and the definition of conditional statistical distance.

$$\begin{aligned} \delta((c_i, \check{o}_i, \hat{p}_i), (c_i, o_i^*, \hat{p}_i)) &= \sum_v \mathbf{Pr}[\hat{p}_i = v] \delta((c_i, \check{o}_i, \hat{p}_i), (c_i, o_i^*, \hat{p}_i) | \hat{p}_i = v) \\ &\geq \sum_v \mathbf{Pr}[\hat{p}_i = v] (\mathbf{Cov}(c, p^* | \hat{p}_i = v) - \eta) \\ &\geq \mathbf{Pr}[\hat{p}_i = v^*] (\mathbf{Cov}(c, p^* | \hat{p}_i = v^*) - \eta) \\ &> \beta(\gamma \mathbf{Var}(p_i^* | \hat{p}_i = v^*) - \eta) \\ &> \beta\gamma\lambda - \beta\eta. \end{aligned}$$

\blacksquare

Theorem 20 now follows by setting parameters and gluing these lemmas together.

Proof [Proof of Theorem 20] First, we apply Lemma 62 with parameters $\varepsilon = \frac{\alpha^3\gamma^2}{16}, \eta = \frac{\alpha\gamma}{8}$ to attain a $\check{p} \in R_{T(\varepsilon)}(\mathcal{C})$. $T(\varepsilon) = O\left(\frac{1}{\varepsilon^2}\right) = O\left(\frac{1}{\alpha^6\gamma^4}\right)$. Also, by the α -hardness assumption, we can apply Lemma 63 to attain a v^* satisfying $\mathbf{Pr}[\hat{p}_i = v^*] \geq 2\alpha\eta$ and $\mathbf{Var}(p_i^* | \hat{p}_i = v^*) \geq \frac{\alpha}{2} - 3\eta$. Finally, by Lemma 65, we conclude that *if there is no hard-core set* (i.e. the conditions of the lemma are satisfied with $\beta = 2\alpha\eta$, $\lambda = \frac{\alpha}{2} - 3\eta$, and γ as provided), then

$$\delta((c_i, \check{o}_i, \hat{p}_i), (c_i, o_i^*, \hat{p}_i)) > \beta\gamma\lambda - \beta\eta > \frac{\alpha^3\gamma^2}{16} = \varepsilon.$$

This violates the conclusion from our application of Lemma 62, that

$$\delta((c_i, \check{o}_i, \hat{p}_i), (c_i, o_i^*, \hat{p}_i)) < \varepsilon.$$

We therefore conclude that there must exist a hardcore set as specified in the theorem statement. ■