# Likelihood-free hypothesis testing

Patrik Róbert Gerber, Yury Polyanskiy, Senior Member, IEEE,

#### Abstract

Consider the problem of binary hypothesis testing. Given Z coming from either  $\mathbb{P}^{\otimes m}$  or  $\mathbb{Q}^{\otimes m}$ , to decide between the two with small probability of error it is sufficient and in most cases necessary to have  $m \simeq 1/\epsilon^2$ , where  $\epsilon$ measures the separation between  $\mathbb{P}$  and  $\mathbb{Q}$  in total variation (TV). This, however, requires complete knowledge of the distributions and can be done, for example, using the Neyman-Pearson test. In this paper we consider a variation of the problem, which we call likelihood-free (or simulation-based) hypothesis testing, where access to  $\mathbb{P}$  and  $\mathbb{Q}$  is given through n iid observations from each. In the case when  $\mathbb{P}, \mathbb{Q}$  are assumed to belong to a non-parametric family  $\mathcal{P}$ , we demonstrate the existence of a fundamental trade-off between n and m given by  $nm \simeq n_{GoF}^2(\epsilon, \mathcal{P})$ , where  $n_{GoF}$  is the minimax sample complexity of testing between the hypotheses  $H_0: \mathbb{P} = \mathbb{Q}$  vs  $H_1: \mathsf{TV}(\mathbb{P}, \mathbb{Q}) \ge \epsilon$ . We show this for three families of distributions, and also study the larger family of all discrete distributions for which we obtain a more complicated trade-off that exhibits a phase-transition. The test that we propose, based on the  $L^2$ -distance statistic of Ingster, simultaneously achieves all points on the trade-off curve for the regular classes. Our results demonstrate the possibility of testing without fully estimating the distributions, provided  $m \gg 1/\epsilon^2$ .

## **CONTENTS**

Introduction						
I-A	Informal statement of the main result	4				
I-B	Related work	5				
I-C	Contributions	5				
I-D	Structure	6				
I-E	Notation	6				
Statistical rates, non-parametric classes and tests						
II-A	Five fundamental problems in Statistics	6				
II-B	Four classes of distributions	7				
II-C	Tests for LFHT	8				
	II-C1 Scheffé's test	8				
	II-C2 Likelihood-free Neyman-Pearson test	8				
	II-C3 Huber's and Birgé's robust tests	9				
	II-C4 Ingster's $L^2$ -distance test	9				
Results		10				
III-A	General reductions	10				
III-B	Sample complexity of likelihood-free hypothesis testing					
III-C	$L^2$ -robust likelihood-free hypothesis testing					
III-D	Beyond total variation	13				
Sketch	proof of main results	14				
IV-A	Upper bounds for Theorems 1 to 4					
	IV-A1 Bounded discrete distributions	15				
	IV-A2 Smooth densities	15				
	IV-A3 Gaussian sequence model	16				
	Introdu I-A I-B I-C I-D I-E Statisti II-A II-B II-C III-A III-B III-C III-D Sketch IV-A	Introduction         I-A       Informal statement of the main result         I-B       Related work         I-C       Contributions         I-D       Structure         I-E       Notation         Statistical rates, non-parametric classes and tests         II-A       Five fundamental problems in Statistics         I-B       Four classes of distributions         II-C       Tests for LFHT         II-C1       Scheffé's test         II-C2       Likelihood-free Neyman-Pearson test         II-C3       Huber's and Birgé's robust tests         II-C4       Ingster's L <sup>2</sup> -distance test         III-B       Sample complexity of likelihood-free hypothesis testing         III-C       L <sup>2</sup> -robust likelihood-free hypothesis testing         III-D       Beyond total variation         Sketch proof of main results         IV-A       Upper bounds for Theorems 1 to 4.         IV-A1       Bounded discrete distributions         IV-A2       Smooth densities         IV-A3       Gaussian sequence model				

Patrik Róbert Gerber and Yury Polyanskiy are with MIT.

		IV-A4 Discrete distributions 10
	IV-B	Lower bounds for Theorems 1 to 4
		IV-B1 Smooth densities
		IV-B2 Bounded discrete distributions
		IV-B3 Gaussian sequence model
		IV-B4 Discrete distributions 1
V	Open p	problems 12
Apper	ndix A:	Upper bounds of Theorem 1 and 2
	A-A	The class $\mathcal{P}_{Db}$
	A-B	The class $\mathcal{P}_{H}$
	A-C	The class $\mathcal{P}_{G}$
	A-D	The class $\mathcal{P}_{D}$ 2.
Apper	ndix B:	Lower bounds of Theorem 1 and 2
	B-A	The class $\mathcal{P}_{H}$
	B-B	The class $\mathcal{P}_{G}$
	B-C	The classes $\mathcal{P}_{Db}$ and $\mathcal{P}_{D}$
		B-C1 Valiant's wishful thinking theorem
Apper	ndix C:	Proof of Theorem 4 34
	C-A	Upper bound
	C-B	Lower bound
Apper	ndix D:	Auxiliary technical results 30
	D-A	Proof of Lemma 1
	D-B	Proof of Proposition 4
	D-C	Lemma 8
	D-D	Proof of Proposition 3
Refere	ences	42

## I. INTRODUCTION

A setting that we call *likelihood-free inference (LFI)*, also known as simulation based inference (SBI), has independently emerged in many areas of science over the past decades. Given an expensive to collect ("experimental") dataset and the ability to simulate from a high fidelity, often mechanistic, stochastic model, whose output distribution (and likelihood) is intractable and inapproximable, how does one perform model selection, parameter estimation or construct confidence sets? The list of disciplines where such highly complex black-box simulators are used is long, and include particle physics, astrophysics, climate science, epidemiology, neuroscience and ecology to just name a few. For some of the above fields, such as climate modeling, the bottleneck resource is in fact the simulated data as opposed to the experimental data. In either case, understanding the trade-off between the number of simulations and experiments necessary to do valid inference is crucial. Our aim in this paper is to introduce a theoretical framework under which LFI can be studied using the tools of non-parametric statistics and information theory.

To illustrate we draw an example from high energy physics, where LFI methods are used and developed extensively. The discovery of the Higgs boson in 2012 [1], [2] is regarded as the crowning achievement of the Large Hadron collider (LHC) - the most expensive instrument ever built. Using a composition of complex simulators [3]–[7] modeling the standard model and the detection process, physicists are able to simulate the results of LHC experiments. Given actual data  $Z_1, \ldots, Z_m$  from the collider, to verify existence of the Higgs boson one tests whether the null hypothesis (physics without the Higgs boson, or  $Z_i \stackrel{iid}{\sim} \mathbb{P}_0$ ) or the alternative hypothesis (physics with the

Higgs boson, or  $Z_i \stackrel{iid}{\sim} \mathbb{P}_1$ ) describes the experimental data more accurately. Note that standard Neyman-Pearson likelihood ratio test is not implementable since  $\mathbb{P}_0$  and  $\mathbb{P}_1$  are only available via simulators.

How was this statistical test actually performed? In essence, a probabilistic classifier C was trained on simulated data to distinguish the two hypotheses (specifically, it was a boosted decision tree classifier). Then, the proportion of real data points falling in the set  $S = \{x \in \mathbb{R}^d : C(x) \le t\}$  is computed, where t is chosen to maximize an asymptotic approximation of the power. Finally, p-values are reported based on the asymptotic distribution under a Poisson sampling model [8], [9]. Summarizing, the "Higgs boson" test was performing the simple comparison

$$\frac{1}{m}\sum_{i=1}^{m}\mathbbm{1}\{Z_i\in S\}\lessgtr\gamma, \tag{Scheffé}$$

where  $Z_1, \ldots, Z_m$  are the real data and  $\gamma$  is some threshold. Such count-based tests (named after Scheffé in folklore [10, Section 6]) are natural.

Notice that Scheffé's test converts each observation  $Z_i$  into a binary 0/1 value. This extreme quantization certainly helps robustness, but may (and should) raise the suspicion of loss of power. Indeed, when the distributions under both hypotheses are completely known, an optimal (Neyman-Pearson) test thresholds the sum of *real-valued* logarithms of the likelihood-ratio. Thus, it is natural to expect that a good test should aggregate non-binary values (We survey some other natural tests in Section II-C). This work is based on one such test that we describe next.

Given some estimates  $\hat{p}_0$ ,  $\hat{p}_1$  of the density of the null and alternative distributions based on simulated samples, our test proceeds via the comparison

$$\frac{2}{m}\sum_{i=1}^{m}(\hat{p}_{0}(Z_{i}) - \hat{p}_{1}(Z_{i})) \leq \gamma$$
(I.1)

where  $Z_1, \ldots, Z_m$  are the real data. Tests of this kind originate from the famous goodness-of-fit work of Ingster [11], which corresponds to taking  $\hat{p}_0 = p_0$ , as the null-density is known exactly.<sup>1</sup> The surprising observation of Ingster was that such a test is able to reject the null hypothesis that  $Z_i \stackrel{iid}{\sim} p_0$  even when the true distribution of Z is much closer to  $p_0$  than described by the optimal density-estimation rate; in other words goodness-of-fit testing is significantly easier than estimation. In fact we will use  $\gamma = \|\hat{p}_0\|_2^2 - \|\hat{p}_1\|_2^2$  in which case (I.1) boils down to the comparison of two squared  $L^2$ -distances.

Our overall goal is to understand the trade-off between the number n of simulated observations and the size of the actual data set m. The characterization of this tradeoff is reminiscent of the rate-regions in multi-user information theory, but there is a certain important difference that we wanted to emphasize for the reader. In information theory, the problem is most often stated in the form "given a distribution  $P_{X,Y,Z}$  (or a channel  $P_{Y,Z|X}$ ) find the rate region", with the distribution being completely specified ahead of time. In minimax statistics, however, distributions are apriori only known to belong to a certain class. In *estimation problems* the fundamental limits are thus defined by minimizing the estimation error over this class, and the theoretical goal is to characterize the worst-case rate at which this error converges to zero as the sample size grows to infinity. The definition of the fundamental limit in *testing problems*, however, is more subtle. If separation is fixed and the number of samples is taken to infinity then the rate of convergence trivializes (becomes exponential in n). By now a standard definition of fundamental limit, as suggested by Ingster (following ideas of Pittman efficiency), is to vary  $\varepsilon$  with n and to find the fastest possible decrease of  $\varepsilon$  so as to still have an acceptable probability of error. This is the approach taken in the literature on goodness-of-fit and two-sample testing, and also the one we adopt here.

Specifically, we assume that it is known apriori that the two distributions  $\mathbb{P}_0$ ,  $\mathbb{P}_1$  belong to a known class  $\mathcal{P}$  and are  $\varepsilon$ -separated under total variation. Given a large number n of samples simulated from  $\mathbb{P}_0$  and  $\mathbb{P}_1$  and m samples  $Z_1, \ldots, Z_m$  from the experiment our goal is to test which of the  $\mathbb{P}_i$  generated the data. If n is sufficiently large to estimate  $\mathbb{P}_i$  (in total variation) to precision  $\varepsilon/10$  then one can perform the hypothesis test with an information-theoretically optimal (even under oracle knowledge of  $\mathbb{P}_i$ 's)  $m \approx 1/\varepsilon^2$  experimental samples. However, looking at

<sup>&</sup>lt;sup>1</sup>In the case of discrete distributions on a finite (but large) alphabet, the idea was rediscovered by the computer science community for goodnessof-fit and two-sample testing (first in [12]). Moreover, the difference of  $L^2$ -norms statistic was first studied (to the best of our knowledge) in [13]. See Section I-B for more on the latter.

the test (Scheffé) one may wonder if the full estimation of the distributions  $\mathbb{P}_i$  is needed, or whether perhaps a suitable decision boundary could be found with a lot fewer simulated samples n. Unfortunately, our *first main result* disproves this intuition: any test using the minimal  $m \simeq 1/\varepsilon^2$  dataset size will require n so large as to be enough to estimate the distributions of  $\mathbb{P}_0$  and  $\mathbb{P}_1$  to within accuracy  $\simeq \varepsilon$ , which is the distance separating the two hypotheses. In particular, any method minimizing m performs no different in the worst case, than pairing off-the-shelf density estimators  $\hat{p}_0, \hat{p}_1$  and applying (Scheffé) with  $S = \{\hat{p}_1 \ge \hat{p}_0\}$ .

This result appears rather pessimistic and seems to invalidate the whole attraction of LFI, which after all attempts to circumvent the exorbitant number of simulation samples required for fully learning high-dimensional distributions. Fortunately, *our second result* offers a resolution: if more data samples  $m \gg 1/\varepsilon^2$  are collected, then testing is possible with *n* much smaller than required for density estimation. More precisely, when neither  $p_0$  nor  $p_1$  are known except through *n* i.i.d. samples from each, the test (I.1) is able to detect which of the two distributions generated the Z-sample, *even when the number of samples n* is insufficient for any estimate  $\hat{p}_i$  to be within a distance  $\approx \varepsilon = \mathsf{TV}(p_0, p_1)$  from the true values. In other words, the test is able to reliably detect the true hypotheses even though the estimates  $\hat{p}_i$  themselves have accuracy that is orders of magnitude larger than the separation  $\varepsilon$  between the hypotheses.

In summary, this paper shows that likelihood-free hypothesis testing (LFHT) is possible without learning the densities when  $m \gg 1/\varepsilon^2$ , but not otherwise. It turns out that the simple test (I.1) has minimax optimal sample complexity up to constants in both n and m in all "regular" (cf. Remark 6) settings.

## A. Informal statement of the main result

Let us formalize the problem using the notation used throughout the rest of the paper. Suppose that we observe true data  $Z \sim \mathbb{P}_{Z}^{\otimes m}$  and that we have two candidate parameter settings for our simulator, from which we generate two artificial datasets  $X \sim \mathbb{P}_{X}^{\otimes n}$  and  $Y \sim \mathbb{P}_{Y}^{\otimes n}$ . If we are convinced that one of the settings accurately reflects reality, we are faced with the problem of testing the hypothesis

$$H_0: \mathbb{P}_{\mathsf{X}} = \mathbb{P}_{\mathsf{Z}}$$
 versus  $H_1: \mathbb{P}_{\mathsf{Y}} = \mathbb{P}_{\mathsf{Z}}.$  (I.2)

*Remark* 1. We emphasize that  $\mathbb{P}_X$  and  $\mathbb{P}_Y$  are known only through the *n* simulated samples. Thus, (I.2) can be interpreted as binary hypothesis testing with approximately specified hypotheses. Alternatively, using the language of machine learning, we may think of this problem as having *n* labeled samples from both classes, and *m* unlabeled samples. The twist is that the unlabeled samples are guaranteed to have the same common label (i.e., be purely coming all from a single class). One can think of many examples of this setting occurring in genetic, medical and other studies.

To put (I.2) in a minimax framework, suppose that the output distribution of the simulator is constrained to lie in a known set  $\mathcal{P}$  and that  $\mathbb{P}_X$ ,  $\mathbb{P}_Y$  are  $\varepsilon$ -separated with respect to the total variation distance TV. Clearly (I.2) becomes 'easier' if we have a lot of data (large sample sizes n and m) or if the hypotheses are well-separated (large  $\varepsilon$ ). We are interested in characterizing the pairs of values (n, m) as functions of  $\varepsilon$  and  $\mathcal{P}$ , for which the hypothesis test (I.2) can be performed with constant type-I and type-II error. Letting  $n_{GoF}(\varepsilon, \mathcal{P})$  denote the minimax sample complexity of goodness-of-fit testing (Definition 2), we show for *several different classes* of  $\mathcal{P}$ , that (I.2) is possible with total error, say, 5% if and only if

$$m\gtrsim 1/arepsilon^2$$
 and  $n\gtrsim n_{\mathsf{GoF}}$  and  $mn\gtrsim n_{\mathsf{GoF}}^2$ .

We also observe, to our knowledge for the first time in the literature, that  $n_{GoF}^2 \varepsilon^2 \simeq n_{Est}$  for these classes<sup>2</sup>, where  $n_{Est}(\varepsilon, \mathcal{P})$  denotes the minimax complexity of density estimation to  $\varepsilon$ -accuracy (Definition 4) with respect to total variation. This observation provides additional meaning to the mysterious formula of Ingster [11] for the goodness-of-fit testing rate for the class of  $\beta$ -smooth densities over  $[0, 1]^d$ , see Table I below. More importantly, however, it allows us to interpret (I.2) as an "interpolation" between different fundamental statistical procedures, namely

<sup>&</sup>lt;sup>2</sup>A possible reason for this observation having been missed previously is that fundamental limits in statistics are usually presented in the form of *rates* of loss decrease with n, e.g.  $r_{\mathsf{Est}}(n) \triangleq n_{\mathsf{Est}}^{-1}(n) = 1/n^{\beta/(2\beta+d)}$  and  $r_{\mathsf{GoF}}(n) \triangleq n_{\mathsf{GoF}}^{-1}(n) = 1/n^{\beta/(2\beta+d/2)}$  for  $\beta$ -smooth densities. Unlike  $n_{\mathsf{Est}} \asymp n_{\mathsf{GoF}}^2 \epsilon^2$  there seems to be no simple relation between  $r_{\mathsf{Est}}$  and  $r_{\mathsf{GoF}}$ .

- $A \leftrightarrow$  Binary hypothesis testing,
- $B \leftrightarrow$  Estimation followed by robust binary hypothesis testing,
- $C \leftrightarrow$  Two-sample testing,
- $D \leftrightarrow$  Goodness-of-fit testing,

corresponding to the extreme points A, B, C, D on Figure 1.

#### B. Related work

The problem (I.2) has initially appeared (to the best of our knowledge) in the work of Ziv and Gutman [14], [15] where in fact an *M*-ary version of the problem was studied. Ziv coined the name *classification with empirically observed statistics* for emphasizing the fact that hypotheses are defined in terms of samples. In terms of results, given two arbitrary, unknown  $\mathbb{P}_X$ ,  $\mathbb{P}_Y$  over a finite alphabet, they show that *n* has to grow at least linearly with *m* to get exponentially decaying error in the latter. Gutman proposes a test whose error exponent is second order optimal, as shown later by [16]. Recent work [17]–[20] extends this problem to settings such as distributed or sequential testing. (Note that we study the opposite regime of fixed-error and focus on dependence of the sample complexity on the separation  $\varepsilon$  of the two alternatives.)

A line of inquiry closer to ours began in [13], [21] where the authors study (I.2) with n = m over the class of discrete distributions p with  $\min_i p_i \simeq \max_i p_i \simeq 1/n^{\alpha}$ , which they call  $\alpha$ -large sources. Disregarding the dependence on the TV-separation  $\varepsilon$  (effectively setting  $\varepsilon$  to a constant), they find that achieving non-trivial minimax error is possible if and only if  $\alpha \le 2$ , using in fact the same difference of squared  $L^2$ -distances test (I.1) that we study in this paper. Follow-up work [22] studies the case  $m \ne n$  and the class of distributions on alphabet [k] with  $\max_i p_i \le 1/k$  showing that non-trivial minimax error is possible if and only if  $k \le \min(n^2, nm)$ . Although the authors omit it from their theorem statements, by the discussion after Proposition VI.1. of [22] one can deduce the minimax optimal dependence on  $\varepsilon$  in the regime  $m, n \le k$ , recovering some of the results of this paper. In summary, we see that previous literature have not addressed the m vs n trade-off (as a function of  $\varepsilon$ ) in all regimes, and only focused on the regular discrete case.

A similar trade-off to ours appears in two-sample testing with unequal sample sizes [23], [24]. This is no coincidence, in Section III-A we show via reductions that in the case  $m \ge n$  our problem (I.2) is equivalent to two-sample testing. Note that for this problem nontrivial behaviour arises only in classes for which  $n_{\text{TS}} \ne n_{\text{GoF}}$  (cf. Remark 6).

(Scheffé) has been considered previously [25]–[31] and is known as a 'classification accuracy' test (CAT) by some. In follow-up work to the present paper [32] it was shown that CATs are in fact (near-)minimax optimal in all settings studied here.

#### C. Contributions

Though the likelihood-free hypothesis testing problem (I.2) has previously appeared under various disguises and was studied in different regimes for the class of bounded discrete distributions, it omitted the key question of understanding dependence of the sample complexity on the separation  $\varepsilon$ . Our work fully characterizes the dependence on the separation  $\varepsilon$  (Theorems 1 and 2). We discover existence of a rather non-trivial trade-off between the *m* and *n* showing that in likelihood-free setting statistical performance (*n*) can be traded for computational resources (*m*). Our results are shown for not just one but multiple distribution classes. In addition, we also demonstrate that LFHT naturally interpolates between its special cases corresponding to goodness-of-fit testing, two-sample testing and density-estimation. As a by-product we observe the relation  $n_{GoF}^2 \varepsilon^2 \approx n_{Est}$  that holds over several classes of distributions and measures of separation (indicating its universality). On the technical side we provide a unified upper bound analysis for all regular classes (cf Remark 6) considered and prove matching lower bounds using techniques of Tsybakov, Ingster and Valiant. Our upper bound analysis is inspired by Ingster [11], [33] whose  $L^2$ -squared distance testing approach (originally designed for goodness-of-fit in smooth-density classes) seems to have been reinvented in the discrete-alphabet world later [12], [13], [21]. Compared to classical works, the new ingredient needed in the discrete case is a "flattening" reduction [24], [34], which we also utilize. Several minor results are also shown along the way, namely, robustness with respect to  $L^2$ -misspecification (Theorem 3) and characterization of  $n_{GoF}$  for the class of  $\beta$ -smooth densities with  $\beta \leq 1$  under Hellinger separation (Theorem 4).

#### D. Structure

Section II defines the statistical problems and the classes of distributions that are studied in this paper; moreover various natural tests for likelihood-free hypothesis testing. Section III contains our main results and the discussion linking to goodness-of-fit and two-sample testing, estimation and robustness. In Section IV we provide a sketch of our proofs for these results. Finally, in Section V we discuss possible future directions of research and the Appendix contains the detailed proofs of Theorems 1 to 4 and auxilary results.

#### E. Notation

For  $k \in \mathbb{N}$  we write  $[k] \triangleq \{1, 2, \dots, k\}$ . For  $x, y \in \mathbb{R}$  we write  $x \wedge y \triangleq \min(x, y), x \vee y \triangleq \max(x, y)$ . We use the Bachmann–Landau notation  $\Omega, \Theta, \mathcal{O}$  as usual and write  $f \leq g$  for  $f = \mathcal{O}(g)$  and  $f \asymp g$  for  $f = \Theta(g)$ . For  $c \in \mathbb{R}$  and  $A \subseteq \mathbb{R}^2$  we write  $cA \triangleq \{(ca_1, ca_2) \in \mathbb{R}^2 : (a_1, a_2) \in A\}$ . For two sets  $A, B \subseteq \mathbb{R}^2$  we write  $A \asymp B$  if there exists a constant c > 0 with  $\frac{1}{c}A \subseteq B \subseteq cA$ . For two probability measures  $\mu, \nu$  dominated by  $\eta$  with densities p, q we define the following divergences:  $\mathsf{TV}(\mu, \nu) \triangleq \frac{1}{2} \int |p - q| d\eta, \mathsf{H}(\mu, \nu) \triangleq (\int (\sqrt{p} - \sqrt{q})^2 d\eta)^{1/2}$ ,  $\mathsf{KL}(\mu \| \nu) \triangleq \int p \log(p/q) d\eta$ ,  $\chi^2(\mu \| \nu) \triangleq \int \frac{(p-q)^2}{q} d\eta$ . Abusing notation, we sometimes write (p, q) as arguments instead of  $(\mu, \nu)$ . Given a divergence D and joint measures  $P_{XY}, Q_{XY}$  we write  $D(P_{Y|X} \| Q_{Y|X} | P_X) \triangleq \mathbb{E}_{X \sim P_X} D(P_{Y|X} \| Q_{Y|X})$ . We write  $\| \cdot \|_p$  for the  $L^p$  and  $\ell^p$  norms, where the base measure shall be clear from the context.

## II. STATISTICAL RATES, NON-PARAMETRIC CLASSES AND TESTS

## A. Five fundamental problems in Statistics

Formally, we define a hypothesis as a set of probability measures. Given two hypotheses  $H_0$  and  $H_1$  on some space  $\mathcal{X}$ , we say that a function  $\psi : \mathcal{X} \to \{0, 1\}$  successfully tests the two hypotheses against each other if

$$\max_{i=0,1} \max_{P \in H_i} \mathbb{P}_{S \sim P}(\psi(S) \neq i) \le 1/3.$$
(II.1)

*Remark* 2. For our purposes, the constant 1/3 above is unimportant and could be replaced by any number less than 1/2. Indeed, throughout the paper we are interested in the asymptotic order of the sample complexity, and sample splitting followed by a majority vote arbitrarily decreases the overall error probability of any successful tester at the cost of a constant factor in the sample complexity.

Throughout this section let  $\mathcal{P}$  be a class of probability distributions on  $\mathcal{X}$ . Suppose we observe independent samples  $X \sim \mathbb{P}_{X}^{\otimes n}$ ,  $Y \sim \mathbb{P}_{Y}^{\otimes n}$  and  $Z \sim \mathbb{P}_{Z}^{\otimes m}$  whose distributions  $\mathbb{P}_{X}, \mathbb{P}_{Y}, \mathbb{P}_{Z} \in \mathcal{P}$  are *unknown* to us. Finally,  $\mathbb{P}_{0}, \mathbb{P}_{1} \in \mathcal{P}$  refer to distributions that are *known* to us. We now define five fundamental problems in statistics that we refer to throughout this paper.

Definition 1. Binary hypothesis testing is the problem of testing

$$H_0: \mathbb{P}_{\mathsf{X}} = \mathbb{P}_0 \qquad \text{against} \qquad H_1: \mathbb{P}_{\mathsf{X}} = \mathbb{P}_1$$
 (HT)

based on the sample X. We denote by  $n_{\mathsf{HT}}(\varepsilon, \mathcal{P})$  the smallest number such that for all  $n \ge n_{\mathsf{HT}}$  and all  $\mathbb{P}_0, \mathbb{P}_1$ with  $\mathsf{TV}(\mathbb{P}_0, \mathbb{P}_1) \ge \varepsilon$  there exists a function  $\psi : \mathcal{X}^n \to \{0, 1\}$  which given X as input successfully tests (in the sense of (II.1))  $H_0$  against  $H_1$ .

It is well known that the complexity of binary hypothesis testing is controlled by the Hellinger divergence.

**Lemma 1.** For all  $\varepsilon$  and  $\mathcal{P}$  with  $|\mathcal{P}| \ge 2$ ,  $n_{\mathsf{HT}}(\varepsilon, \mathcal{P}) = \Theta(\sup_{\mathbb{P}_0, \mathbb{P}_1 \in \mathcal{P}: \mathsf{TV}(\mathbb{P}_0, \mathbb{P}_1) \ge \varepsilon} \mathsf{H}^{-2}(\mathbb{P}_0, \mathbb{P}_1))$  where the implied constant is universal.

*Proof.* We include the proof in Section D-A for completeness.

In fact, for all  $\mathcal{P}$  mentioned in this paper  $n_{\text{HT}} = \Theta(1/\varepsilon^2)$  holds. Therefore, going forward we refrain from the general notation  $n_{\text{HT}}$  and simply write  $1/\varepsilon^2$ .

Definition 2. Goodness-of-fit testing is the problem of testing

$$H_0: \mathbb{P}_{\mathsf{X}} = \mathbb{P}_0 \quad \text{against} \quad H_1: \mathsf{TV}(\mathbb{P}_{\mathsf{X}}, \mathbb{P}_0) \ge \varepsilon$$
 (GoF)

based on the sample X. Write  $n_{\mathsf{GoF}}(\varepsilon, \mathcal{P})$  for the smallest value such that for all  $n \ge n_{\mathsf{GoF}}$  and  $\mathbb{P}_0 \in \mathcal{P}$  there exists a function  $\psi : \mathcal{X}^n \to \{0, 1\}$  which given X as input successfully tests (in the sense of (II.1))  $H_0$  against  $H_1$ .

Definition 3. Two-sample testing is the problem of testing

$$H_0: \mathbb{P}_{\mathsf{X}} = \mathbb{P}_{\mathsf{Z}}$$
 against  $H_1: \mathsf{TV}(\mathbb{P}_{\mathsf{X}}, \mathbb{P}_{\mathsf{Z}}) \ge \varepsilon$  (TS)

based on the samples X, Z. Write  $\mathcal{R}_{\mathsf{TS}}(\varepsilon, \mathcal{P}) \subseteq \mathbb{R}^2$  for the maximal set such that for all  $(n,m) \in \mathbb{N}^2$  with  $n \geq x, m \geq y$  for some  $(x, y) \in \mathcal{R}_{\mathsf{TS}}$ , there exists a function  $\psi : \mathcal{X}^n \times \mathcal{X}^m \to \{0, 1\}$  which given X, Z as input successfully tests (in the sense of (II.1)) between  $H_0$  and  $H_1$ . We will use the abbreviation  $n_{\mathsf{TS}}(\varepsilon, \mathcal{P}) = \min\{\ell \in \mathbb{N} : (\ell, \ell) \in \mathcal{R}_{\mathsf{TS}}(\varepsilon, \mathcal{P})\}$ .

**Definition 4.** The sample complexity of **estimation** is the smallest value  $n_{\mathsf{Est}}(\varepsilon, \mathcal{P})$  such that for all  $n \ge n_{\mathsf{Est}}$  there exists an estimator  $\hat{\mathbb{P}}_X$  which given X as input satisfies

$$\mathbb{E}\mathsf{T}\mathsf{V}(\hat{\mathbb{P}}_{\mathsf{X}}, \mathbb{P}_{\mathsf{X}}) \le \varepsilon. \tag{Est}$$

Definition 5. Likelihood-free hypothesis testing is the problem of testing

$$H_0: \mathbb{P}_{\mathsf{Z}} = \mathbb{P}_{\mathsf{X}} \quad \text{against} \quad \mathbb{P}_{\mathsf{Z}} = \mathbb{P}_{\mathsf{Y}}$$
 (LF)

based on the samples X, Y, Z. Write  $\mathcal{R}_{\mathsf{LF}}(\varepsilon, \mathcal{P}) \subseteq \mathbb{R}^2$  for the maximal set such that for all  $(n, m) \in \mathbb{N}^2$  with  $n \geq x, m \geq y$  for some  $(x, y) \in \mathcal{R}_{\mathsf{LF}}$ , there exists a function  $\psi : \mathcal{X}^n \times \mathcal{X}^m \times \mathcal{X}^m \to \{0, 1\}$  which given X, Y, Z as input, successfully tests (in the sense of (II.1))  $H_0$  against  $H_1$  provided  $\mathsf{TV}(\mathbb{P}_X, \mathbb{P}_Y) \geq \varepsilon$ .

*Remark* 3. Requiring  $\mathcal{R}_{LF}$  to be maximal is well defined because  $\mathcal{R}_{LF} \ni (n_0, m_0) \le (n, m)$  coordinate-wise implies we can take  $(n, m) \in \mathcal{R}_{LF}$ , since  $\psi$  can simply disregard extra samples.

*Remark* 4. (GoF) can be thought of as a version of (HT) where only the null is known and the alternative is specified up to an i.i.d. sample. This leads naturally to the generalization (LF) where both hypotheses are known only up to i.i.d. samples.

*Remark* 5. All five definitions above can be modified to measure separation with respect to an arbitrary function d instead of TV. We will write  $n_{GoF}(\varepsilon, d, P)$  etc. for the corresponding values.

## B. Four classes of distributions

To state our results, we need to introduce the nonparametric classes of distributions that we consider in this paper.

(i) Smooth density. Let  $C(\beta, d, C)$  denote the set of functions  $f : [0, 1]^d \to \mathbb{R}$  that are  $\lceil \beta - 1 \rceil$ -times differentiable and satisfy

$$\|f\|_{\mathcal{C}_{\beta}} \triangleq \max\left(\max_{0 \le |\alpha| \le \lceil \beta - 1\rceil} \|f^{(\alpha)}\|_{\infty}, \sup_{x \ne y \in [0,1]^d, |\alpha| = \lceil \beta - 1\rceil} \frac{|f^{(\alpha)}(x) - f^{(\alpha)}(y)|}{\|x - y\|_2^{\beta - \lceil \beta - 1\rceil}}\right) \le C,$$

where  $|\alpha| = \sum_{i=1}^{d} \alpha_i$  for the multiindex  $\alpha \in \mathbb{N}^d$ . We write  $\mathcal{P}_{\mathsf{H}}(\beta, d, C_{\mathsf{H}})$  for the class of distributions with Lebesgue-densities in  $\mathcal{C}(\beta, d, C_{\mathsf{H}})$ .

(ii) Gaussian sequence model on the Sobolev ellipsoid. Define the Sobolev ellipsoid  $\mathcal{E}(s, C)$  of smoothness s > 0and size C > 0 as  $\{\theta \in \mathbb{R}^{\mathbb{N}} : \sum_{j=1}^{\infty} j^{2s} \theta_j^2 \leq C\}$ . For  $\theta \in \mathbb{R}^{\infty}$  let  $\mu_{\theta} = \bigotimes_{i=1}^{\infty} \mathcal{N}(\theta_i, 1)$ . We define our second class as

$$\mathcal{P}_{\mathsf{G}}(s, C_{\mathsf{G}}) \triangleq \{\mu_{\theta} : \theta \in \mathcal{E}(s, C_{\mathsf{G}})\}.$$

#### (iii)-(iv) **Distributions on a finite alphabet.** For $k \in \mathbb{N}$ , let

$$\mathcal{P}_{\mathsf{D}}(k) \triangleq \{ \text{all distributions on the finite alphabet } [k] \},\$$

$$\mathcal{P}_{\mathsf{Db}}(k, C_{\mathsf{Db}}) \triangleq \{ p \in \mathcal{P}_{\mathsf{D}}(k) : \|p\|_{\infty} \le C_{\mathsf{Db}}/k \},\$$

where  $C_{Db} > 1$  is a constant. In other words,  $\mathcal{P}_{Db}$  are those discrete distributions that are bounded by a constant multiple of the uniform distribution.

*Remark* 6. We call  $\mathcal{P}_{Db}$  the "regular discrete" class. We'll see that it behaves similarly to  $\mathcal{P}_{H}$  and  $\mathcal{P}_{G}$  but different from  $\mathcal{P}_{D}$ . More generally we call the classes  $\mathcal{P}_{H}, \mathcal{P}_{G}, \mathcal{P}_{Db}$  "regular", characterized by the fact that  $n_{GoF} \approx n_{TS}$  and consequently  $\mathcal{R}_{TS} \approx \{(n, m) : n \land m \ge n_{TS}\}$ .

## C. Tests for LFHT

In this section we discuss various types of tests that can be considered for (LF).

- (i) Scheffé's test
- (ii) Likelihood-free Neyman-Pearson test
- (iii) Huber's and Birgé's robust tests
- (iv) Ingster's  $L^2$ -distance test

Tests (i-ii) are based on the idea of learning (from the simulated samples) a set or a function separating  $\mathbb{P}_X$  from  $\mathbb{P}_Y$ . Tests (iii-iv) use the simulated samples to obtain density estimates of  $\mathbb{P}_X, \mathbb{P}_Y$  directly. All of them, however, are of the form

$$\sum_{i=1}^{m} s(Z_i) \leq 0 \tag{II.2}$$

with only the function s varying.

1) Scheffé's test: Variants of Scheffé's test using machine-learning enabled classifiers are the subject of current research in two-sample testing [27]–[31] and are used in practice for LFI specifically in high energy physics, cf. Section I. Thus, understanding the performance of Scheffé's test in the context of (LF) is of great theoretical and practical importance. Suppose that using the simulated samples we train a probabilistic classifier  $C : \mathcal{X} \to [0, 1]$  on the labeled data  $\bigcup_{i=1}^{n} \{(X_i, 0), (Y_i, 1)\}$ . The specific form of the classifier here is arbitrary and can be anything from logistic regression to a deep neural network. Given thresholds  $t, \gamma \in [0, 1]$  chosen to satisfy our risk appetite for type-1 vs type-2 errors, Scheffé's test proceeds via the comparison

$$\frac{1}{m}\sum_{i=1}^{m}\mathbb{1}\{C(Z_i) \ge t\} \le \gamma.$$
(II.3)

We see that (II.3) is of the form (II.2) with  $s(z) = (\mathbb{1}\{C(z) \ge t\} - \gamma)/m$ . The follow-up work [32] studies the performance of Scheffé's test in great detail, finding that it is (near)-minimax optimal in all cases considered in this paper.

2) Likelihood-free Neyman-Pearson test: If the distributions  $\mathbb{P}_X$ ,  $\mathbb{P}_Y$  are known then the (minimax optimal) Neyman-Pearson test corresponds to

$$\sum_{i=1}^{m} s_{\mathsf{NP}}(Z_i) \leq \gamma \qquad s_{\mathsf{NP}}(z) = \log\left(\frac{\mathrm{d}\mathbb{P}_{\mathsf{X}}}{\mathrm{d}\mathbb{P}_{\mathsf{Y}}}(z)\right),\tag{II.4}$$

where  $\gamma$  is again chosen to satisfy our type-1 vs type-2 error trade-off preferences. However, in our setting  $\mathbb{P}_X$ ,  $\mathbb{P}_Y$  are known only up to i.i.d. samples. Notice that  $s_{NP}$  minimizes the population *cross-entropy (or logistic) loss*, that is

$$s_{\mathsf{NP}} = \operatorname*{arg\,min}_{s} \mathbb{E}_{z \sim \mathbb{P}_{\mathsf{X}}}[\ell(s(z), 1)] + \mathbb{E}_{z \sim \mathbb{P}_{\mathsf{Y}}}[\ell(s(z), 0)],$$

where  $\ell(s, y) = \log(1 + e^s) - ys$ . In practice, the majority of today's classifiers are obtained by running some form of gradient descent on the problem

$$\hat{s} = \mathop{\arg\min}_{s \in \mathcal{G}} \mathbb{E}_{z \sim \hat{\mathbb{P}}_{\mathsf{X}}}[\ell(s(z), 1)] + \mathbb{E}_{z \sim \hat{\mathbb{P}}_{\mathsf{Y}}}[\ell(s(z), 0)],$$

where  $\mathcal{G}$  is (for example) a parametric class of neural networks and  $\hat{\mathbb{P}}_X$ ,  $\hat{\mathbb{P}}_Y$  are empirical distributions. Given such an estimate  $\hat{s}$ , we can replace the unknown  $s_{NP}$  in (II.4) by  $\hat{s}$  to obtain the *likelihood-free* Neyman-Pearson test. For recent work on this approach in LFI see e.g. [35]. Studying properties of this test is outside the scope of this paper.

3) Huber's and Birgé's robust tests: The next approach is based on the idea of robust testing, first proposed by Huber [36], [37]. Huber's seminal result implies that if one has approximately correct distributions  $\hat{\mathbb{P}}_X$ ,  $\hat{\mathbb{P}}_Y$  satisfying

$$\mathsf{TV}(\hat{\mathbb{P}}_{\mathsf{X}}, \mathbb{P}_{\mathsf{X}}) \lor \mathsf{TV}(\hat{\mathbb{P}}_{\mathsf{Y}}, \mathbb{P}_{\mathsf{Y}}) \leq \varepsilon/3 \qquad \text{and} \qquad \mathsf{TV}(\mathbb{P}_{\mathsf{X}}, \mathbb{P}_{\mathsf{Y}}) \geq \varepsilon,$$

then for some  $c_1 < c_2$  the test

$$\sum_{i=1}^{m} s_{\mathsf{H}}(Z_i) \leq 0 \qquad \text{where} \qquad s_{\mathsf{H}}(z) = \left\{ c_1 \vee \log \frac{\mathrm{d}\hat{\mathbb{P}}_{\mathsf{X}}}{\mathrm{d}\hat{\mathbb{P}}_{\mathsf{Y}}}(z) \right\} \wedge c_2$$

has type-1 and type-2 error bounded by  $\exp(-\Omega(m\varepsilon^2))$  (and is in fact minimax optimal for all sample sizes). From this we see that Scheffé's test can be interpreted as an approximation of the maximally robust Huber's test. Let  $\hat{\mathcal{L}}(z) = (d\hat{\mathbb{P}}_{\mathsf{Y}}/d\hat{\mathbb{P}}_{\mathsf{X}})(z)$  denote the likelihood-ratio of the estimates. The values of  $c_1, c_2$  are given as the solution to

$$\varepsilon/3 = \mathbb{E}_{z \sim \hat{\mathbb{P}}_{\mathsf{X}}} \left[ \mathbbm{1}\left\{ \hat{\mathcal{L}}(z) \le c_1 \right\} \frac{c_1 - \hat{\mathcal{L}}(z)}{1 + c_1} \right] = \mathbb{E}_{z \sim \hat{\mathbb{P}}_{\mathsf{Y}}} \left[ \mathbbm{1}\left\{ \hat{\mathcal{L}}(z) \ge c_2 \right\} \frac{\hat{\mathcal{L}}(z) - c_2}{1 + c_2} \right],$$

which can be easily approximated to high accuracy given samples from  $\hat{\mathbb{P}}_X$ ,  $\hat{\mathbb{P}}_Y$ . This suggests both a theoretical construction (since  $\hat{\mathbb{P}}_X$ ,  $\hat{\mathbb{P}}_Y$  can be obtained with high probability from simulation samples via the general estimator of Yatracos [38]) and a practical rule: instead of the possibly brittle likelihood-free Neyman-Pearson test (*ii*), one should try clamping the estimated log-likelihood ratio from above and below.

Similar results hold due to Birgé [39], [40] in the case when distance is measured by Hellinger divergence:

$$\mathsf{H}(\mathbb{P}_{\mathsf{X}},\mathbb{P}_{\mathsf{X}}) \lor \mathsf{H}(\mathbb{P}_{\mathsf{Y}},\mathbb{P}_{\mathsf{Y}}) \leq \varepsilon/3 \qquad \text{and} \qquad \mathsf{H}(\mathbb{P}_{\mathsf{X}},\mathbb{P}_{\mathsf{Y}}) \geq \varepsilon.$$

For ease of notation, let  $\hat{p}_X, \hat{p}_Y$  denote the densities of  $\hat{\mathbb{P}}_X, \hat{\mathbb{P}}_Y$  with respect to some base measure  $\mu$ . Regarding  $\sqrt{\hat{p}_X}$  and  $\sqrt{\hat{p}_Y}$  as unit vectors of the Hilbert space  $L^2(\mu)$ , let  $\gamma : [0,1] \to L^2(\mu)$  be the constant speed geodesic on the unit sphere of  $L^2(\mu)$  with  $\gamma(0) = \sqrt{\hat{p}_X}$  and  $\gamma(1) = \sqrt{\hat{p}_Y}$ . It is easily checked that each  $\gamma_t$  is positive (i.e. square-root-densities form a geodesically convex subset of the unit sphere of  $L^2(\mu)$ ) and Birgé showed that the test

$$\sum_{i=1}^{m} \log \left( \frac{\gamma_{1/3}^2}{\gamma_{2/3}^2}(Z_i) \right) \leq 0$$

has both type-I and type-II errors bounded by  $\exp(-\Omega(m\varepsilon^2))$ .

4) Ingster's  $L^2$ -distance test: Finally, we re-introduce the statistic (I.1), based on Ingster's goodness-of-fit testing idea. For simplicity we focus on the case of discrete distributions. This case is more general than may first appear: for example in the case of smooth densities on  $[0, 1]^d$  one can simply take a regular grid (whose resolution is determined by the smoothness of the densities) and count the number of datapoints falling in each cell. Let  $\hat{p}_X$ ,  $\hat{p}_Y$ ,  $\hat{p}_Z$  denote the empirical probability mass functions of the finitely supported distributions  $\hat{\mathbb{P}}_X$ ,  $\hat{\mathbb{P}}_Y$ ,  $\hat{\mathbb{P}}_Z$ . The test proceeds via the comparison

$$\|\hat{p}_{\mathsf{X}} - \hat{p}_{\mathsf{Z}}\|_{2} \leq \|\hat{p}_{\mathsf{Y}} - \hat{p}_{\mathsf{Z}}\|_{2}.$$
 (II.5)

Squaring both sides and rearranging, we arrive at the form

$$\frac{1}{m}\sum_{i=1}^{m}(\hat{p}_{\mathsf{Y}}(Z_i) - \hat{p}_{\mathsf{X}}(Z_i)) \leq \gamma,$$

where  $\gamma = (\|\hat{p}_{\mathsf{Y}}\|^2 - \|\hat{p}_{\mathsf{X}}\|^2)/2$ . As mentioned in the introduction, variants of this  $\ell^2$ -distance based test have been invented and re-invented multiple times for goodness-of-fit [11], [12] and two-sample testing [41], [42]. The exact statistic (II.5) with application to  $\mathcal{P}_{\mathsf{Db}}$  has appeared in [13], [21], and Huang and Meyn [22] proposed an ingenious improvement restricting attention exclusively to bins whose counts are one of (2,0), (1,1), (0,2) for the samples (X, Z) or (Y, Z). We attribute (II.5) to Ingster because his work on goodness-of-fit testing for smooth densities is the first occurrence of the idea of comparing empirical  $\ell^2$  norms, but we note that [13] and [12] arrive at this influential idea apparently independently.

We emphasize the following subtlety. Let us rewrite (II.5) as

$$\|\hat{p}_{\mathsf{X}} - \hat{p}_{\mathsf{Z}}\|_{2}^{2} - \|\hat{p}_{\mathsf{Y}} - \hat{p}_{\mathsf{Z}}\|_{2}^{2} \leq 0.$$
 (II.6)

As we argue below, this difference results in an optimal test regardless of n and m for  $\mathcal{P}_{Db}$ . However, it does not mean that each term by itself is a meaningful estimate of the corresponding distance: rejecting the null by thresholding  $\|\hat{p}_{\mathsf{X}} - \hat{p}_{\mathsf{Z}}\|_2^2$  would not work. Indeed, the variance of  $\|p_{\mathsf{X}} - \hat{p}_{\mathsf{Z}}\|_2^2$  is so large that it requires taking  $m \gg 1/\varepsilon^2$  (in fact, requires m to be at least  $n_{\mathsf{GoF}}$ ). The "magic" of the  $L_2$ -difference test is that the two terms in (II.6) separately have high variance (and are not good estimators of their means), but their difference cancels the high-variance terms.

*Remark* 7. While testing (LF), practitioners are usually interested in obtaining a *p*-value rather than purely a decision whether to reject the null hypothesis. For this we propose the following scheme. Let  $\sigma_1, \ldots, \sigma_P$  be i.i.d. uniformly random permutations on n + m elements. Let  $\hat{T} = \|\hat{p}_X - \hat{p}_Z\|_2^2 - \|\hat{p}_Y - \hat{p}_Z\|_2^2$  be our statistic, and write  $\hat{T}_i$  for the statistic  $\hat{T}$  evaluated on the permuted dataset where  $\{X_1, \ldots, X_n, Z_1, \ldots, Z_m\}$  are shuffled according to  $\sigma_i$ . Under the null the random variables  $\hat{T}, \hat{T}_1, \ldots, \hat{T}_P$  are exchangeable, thus reporting the empirical upper quantile of  $\hat{T}$  in this sample yields an unbiased estimator of the *p*-value. Studying the power of this procedure is beyond the scope of this work.

#### **III. RESULTS**

Before presenting our results for the specific classes introduced in Section II-B, we give general reductions valid for any nonparametric class  $\mathcal{P}$  and separation measure d (cf Remark 5).

#### A. General reductions

**Proposition 1.** Let  $\mathcal{P}$  be a generic family of distributions and  $d : \mathcal{P}^2 \to \mathbb{R}$  be any function used to measure separation. There exists a universal constant c > 0 such that for  $n, m \in \mathbb{N}$  the following implications hold.

$$(n,m) \in \mathcal{R}_{\mathsf{LF}} \implies m \ge n_{\mathsf{HT}},$$
 (III.1)

$$(n,m) \in \mathcal{R}_{\mathsf{TS}} \implies n \wedge m \ge n_{\mathsf{GoF}}$$
 (III.2)

$$(n,m) \in \mathcal{R}_{\mathsf{LF}} \implies cn \ge n_{\mathsf{GoF}},$$
 (III.3)

$$(n,m) \in \mathcal{R}_{\mathsf{TS}} \implies (n,m) \in \mathcal{R}_{\mathsf{LF}},$$
 (III.4)

$$m \ge n \text{ and } (n,m) \in \mathcal{R}_{\mathsf{LF}} \implies (cn,cm) \in \mathcal{R}_{\mathsf{TS}},$$
 (III.5)

where we omit the argument  $(\varepsilon, d, P)$  for simplicity. In particular,

$$\mathbb{N}_{n\leq m}^2 \cap \mathcal{R}_{\mathsf{LF}} \asymp \mathbb{N}_{n\leq m}^2 \cap \mathcal{R}_{\mathsf{TS}},\tag{III.6}$$

where  $\mathbb{N}_{n\leq m}^2 = \{(n,m)\in\mathbb{N}^2:n\leq m\}.$ 

*Proof.* In what follows, let  $\Psi_{LF}$ ,  $\Psi_{TS}$  be minimax optimal tests for (LF) and (TS) respectively. Throughout the proof we omit the arguments ( $\varepsilon$ , d,  $\mathcal{P}$ ) for notational simplicity.

**Reducing hypothesis testing to** (LF) Suppose  $(n, m) \in \mathcal{R}_{LF}$ . Let  $\mathbb{P}_0, \mathbb{P}_1 \in \mathcal{P}$  be given with  $d(\mathbb{P}_0, \mathbb{P}_1) \geq \varepsilon$  and suppose Z is an i.i.d. sample with m observations. We wish to test the hypothesis  $H_0 : Z_i \sim \mathbb{P}_0$  against  $H_1 : Z_i \sim \mathbb{P}_1$ . To this end generate n i.i.d. observations X, Y from  $\mathbb{P}_0, \mathbb{P}_1$  respectively, and simply output  $\Psi_{LF}(X, Y, Z)$ . This shows that if  $(n, m) \in \mathcal{R}_{LF}$  then  $m \geq n_{HT}$  and concludes the proof of (III.1).

Reducing goodness-of-fit testing to two-sample testing Suppose  $(n, m) \in \mathcal{R}_{TS}$ . Then obviously  $(n \wedge m, \infty) \in \mathcal{R}_{TS}$ . However, two-sample testing with sample sizes  $n \wedge m, \infty$  is equivalent to goodness-of-fit testing with a sample size of  $n \wedge m$ . Therefore,  $n \wedge m \ge n_{GoF}$  must hold, concluding the proof of (III.2).

**Reducing goodness-of-fit testing to** (LF) Suppose  $(n, m) \in \mathcal{R}_{LF}$  with  $m \leq n$ . Let a distribution  $\mathbb{P}_0 \in \mathcal{P}$  be given as well as an i.i.d. sample X of size cn with unknown distribution  $\mathbb{P}_X$ , where  $c \in \mathbb{N}$  is a large integer. We want to test  $H_0 : \mathbb{P}_X = \mathbb{P}_0$  against  $H_1 : \mathbb{P}_X \in \mathcal{P}, d(\mathbb{P}_X, \mathbb{P}_0) \geq \varepsilon$ . Generate  $c \times 2$  i.i.d. samples  $Y^{(i)}, Z^{(i)}$  for  $i = 1, \ldots, c$  of size n, m respectively, all from  $\mathbb{P}_0$ . Split the sample X into c batches  $X^{(i)}, i = 1, \ldots, c$  of size n each and form the variables

$$A_i = \Psi_{\mathsf{LF}}(X^{(i)}, Y^{(i)}, Z^{(i)}) - \Psi_{\mathsf{LF}}(X^{(i)}, Y^{(i)}, X^{(i+1)}_{1:m})$$

for  $i = 1, 3, ..., 2\lfloor c/2 \rfloor - 1$ , where  $X_{1:m}^{(i)}$  denotes the first *m* observations in the batch  $X^{(i)}$ . Note that the  $A_i$  are i.i.d. and bounded random variables. Under the null hypothesis we have  $\mathbb{E}A_i = 0$ , while under the alternative they have mean  $\mathbb{E}A_i \ge 1/3$  (since  $\Psi_{\mathsf{LF}}$  is a successful tester in the sense of (II.1)). Therefore, a constant number c/2 observations suffice to decide whether  $\mathbb{P}_{\mathsf{X}} = \mathbb{P}_0$  or not. In particular,  $cn \ge n_{\mathsf{GoF}}$  which concludes the proof of (III.3) for the case  $m \le n$ . The case  $n \le m$  follows from (III.5) and (III.2).

**Reducing** (LF) to two-sample testing Suppose  $(n, m) \in \mathcal{R}_{TS}$ . Let three samples X, Y, Z be given, of sizes a, a, b from the unknown distributions  $\mathbb{P}_X, \mathbb{P}_Y, \mathbb{P}_Z$  respectively, where  $\{a, b\} = \{n, m\}$ . We want to test the hypothesis  $H_0 : \mathbb{P}_X = \mathbb{P}_Z$  against  $H_1 : \mathbb{P}_Y = \mathbb{P}_Z$ , where  $d(\mathbb{P}_X, \mathbb{P}_Y) \ge \varepsilon$  under both. Then, the test

$$\Psi_{\mathsf{LF}}(X,Y,Z) \triangleq \Psi_{\mathsf{TS}}(X,Z)$$

shows that  $(n, m), (m, n) \in \mathcal{R}_{LF}$  and concludes the proof of (III.4).

**Reducing two-sample testing to** (LF) Suppose  $(n, m) \in \mathcal{R}_{LF}$  where  $m \ge n$ . Let two samples X, Y be given, from the unknown distributions  $\mathbb{P}_X, \mathbb{P}_Y \in \mathcal{P}$  and of sample size cn, cm respectively, where  $c \in \mathbb{N}$  is a large integer. We wish to test the hypothesis  $H_0: \mathbb{P}_X = \mathbb{P}_Y$  against  $H_1: d(\mathbb{P}_X, \mathbb{P}_Y) \ge \varepsilon$ . Split the samples X, Y into  $2 \times c$  batches  $X^{(i)}, Y^{(i)}, i = 1, \ldots, c$  of sizes n, m respectively, and form the variables

$$A_{i} = \Psi_{\mathsf{LF}}(X^{(i)}, Y^{(i)}_{1:n}, Y^{(i+1)}) - \Psi_{\mathsf{LF}}(Y^{(i)}_{1:n}, X^{(i)}, Y^{(i+1)})$$

for  $i = 1, 3, ..., 2\lfloor c/2 \rfloor - 1$ , where  $Y_{1:n}^{(i)}$  denotes the first *n* observations in the batch  $Y^{(i)}$ . The variables  $A_i$  are i.i.d. and bounded. Under the null hypothesis we have  $\mathbb{E}A_i = 0$  while under the alternative  $\mathbb{E}A_i \ge 1/3$  holds. Therefore a constant number c/2 observations suffice to decide whether  $\mathbb{P}_X = \mathbb{P}_Y$  or not. In particular,  $(cn, cm) \in \mathcal{R}_{\mathsf{TS}}$  which concludes the proof of (III.5).

Equivalence between two-sample testing and (LF) Equation (III.6) follows immediately from (III.5) and (III.4).  $\Box$ 

Equation (III.6) tells us that the problems of likelihood-free hypothesis testing and two-sample testing are equivalent, *but only for*  $m \ge n$ , i.e. when we have more real data than simulated data. We will see in the next section (and on Figure 1 visually) that this distinction is necessary.

## B. Sample complexity of likelihood-free hypothesis testing

In this section we present our results on the sample complexity of (LF) for specific classes  $\mathcal{P}$  with separation measured by d = TV. In all results below the parameters  $\beta$ , d,  $C_H$ , s,  $C_G$ ,  $C_{Db}$  are regarded as constants, we only care about the dependence on the separation  $\varepsilon$  and the alphabet size k (in the case of  $\mathcal{P}_D$ ,  $\mathcal{P}_{Db}$ ). Where convenient we omit the arguments of  $n_{GoF}$ ,  $n_{TS}$ ,  $\mathcal{R}_{TS}$ ,  $n_{Est}$ ,  $\mathcal{R}_{LF}$  to ease notation, whose value should be clear from the context.

**Theorem 1.** Under TV-separation, for each choice  $\mathcal{P} \in {\mathcal{P}_{H}, \mathcal{P}_{G}, \mathcal{P}_{Db}}$  we have

$$\mathcal{R}_{\mathsf{LF}}(\varepsilon) \asymp \Big\{ m \ge 1/\varepsilon^2, n \ge n_{\mathsf{GoF}}(\varepsilon), mn \ge n_{\mathsf{GoF}}(\varepsilon)^2 \Big\},\$$

where the implied constants do not depend on k (in the case of  $\mathcal{P}_{Db}$ ) or  $\varepsilon$ .

For each class  $\mathcal{P}$  in Theorem 1, the entire region  $\mathcal{R}_{LF}$  (within universal constant) is attained by Ingster's  $\ell^2$ distance test from Section II-C4. We remark that for the class  $\mathcal{P}_{Db}$  in the regime  $m, n \leq k$ , the results of Theorem 1 can be deduced from the discussion after Proposition VI.1. of [22].

Each corner point {A, B, C, D} of Figure 1 has a special interpretation. A corresponds to binary hypothesis testing and D can be reduced to goodness-of-fit testing. Similarly, B and C can be reduced to the well-known problems of estimation followed by robust hypothesis testing and two-sample testing respectively. In other words, (LF) allows us to naturally interpolate between multiple statistical problems. Finally, we point out a curious fact: since the product of n and m remains constant on the line segment [B, C] on the left plot of Figure 1, it follows that

$$n_{\mathsf{Est}}(\varepsilon, \mathcal{P}) \asymp n_{\mathsf{GoF}}^2(\varepsilon, \mathcal{P}) \varepsilon^2$$

for each class  $\mathcal{P}$  treated in Theorem 1. This relation between the sample complexity of estimation and goodness-of-fit testing has not been observed before to our knowledge, and the generality of this phenomenon remains open.

Turning to our results on  $\mathcal{P}_{\mathsf{D}}$  the picture is less straightforward. As first identified in [43] and fully resolved in [44], the rates of two-sample testing undergo a phase transition in the large alphabet  $(k \ge 1/\varepsilon^4)$  regime; this phase transition appears also in likelihood-free hypothesis testing.

**Theorem 2.** Let  $\alpha = 1 \vee (\frac{k}{n} \wedge \frac{k}{m})$ . Then

$$\mathcal{R}_{\mathsf{LF}}(\varepsilon, \mathcal{P}_{\mathsf{D}}(k)) \asymp_{\log(k)} \left\{ \begin{matrix} m \ge 1/\varepsilon^2, n \ge n_{\mathsf{GoF}}(\varepsilon) \cdot \sqrt{\alpha} \\ mn \ge n_{\mathsf{GoF}}(\varepsilon)^2 \cdot \alpha \end{matrix} \right\}$$

where the equivalence is up to a logarithmic factor in the alphabet size k.



Fig. 1. Light and dark gray show  $\mathcal{R}_{LF}$  and its complement resp. (on  $\log_{1/\varepsilon}$ -scale); the striped region depicts  $\mathcal{R}_{TS} \subset \mathcal{R}_{LF}$ . Left plot is valid for  $\mathcal{P} \in \{\mathcal{P}_{H}, \mathcal{P}_{G}, \mathcal{P}_{Db}\}$  for all settings of  $\varepsilon, k$ . For  $\mathcal{P}_{D}$  the left plot applies when  $k \lesssim \varepsilon^{-4}$  and the right plot otherwise.

The log(k)-sized gap in Theorem 2 is an artifact of the proof (a union bound) and can most likely be removed. In fact, we can remove this gap in all cases except the regime  $k \gtrsim 1/\varepsilon^4$ ,  $m \lesssim n \lesssim k$ . For a comparison between  $\mathcal{R}_{LF}$  and  $\mathcal{R}_{TS}$  see Figure 1.

For the reader's convenience, Table I summarizes previously known tight results for the values of  $n_{GoF}$ ,  $n_{TS}$ ,  $\mathcal{R}_{TS}$  and  $n_{Est}$ . The fact that  $n_{HT} = \Theta(1/\varepsilon^2)$  for reasonable classes is classical, see Lemma 1. The study of goodness-of-fit

testing within a minimax framework was pioneered by Ingster [11], [33] for  $\mathcal{P}_H$ ,  $\mathcal{P}_G$ , and independently studied by the computer science community [12], [45] for  $\mathcal{P}_D$ ,  $\mathcal{P}_{Db}$  under the name *identity testing*. Two-sample testing (a.k.a. *closeness testing*) was solved in [44] for  $\mathcal{P}_D$  (with the optimal result for  $\mathcal{P}_{Db}$  implicit) and [11], [42], [46] consider  $\mathcal{P}_H$ . The study of the rate of estimation  $n_{Est}$  is older, see [47]–[50] and references for  $\mathcal{P}_H$ ,  $\mathcal{P}_G$  and [51] for  $\mathcal{P}_D$ ,  $\mathcal{P}_{Db}$ . We reiterate that we are not aware of previous literature identifying the connection between (GoF) and (Est) as

	$n_{HT}$	$n_{GoF}$	$\mathcal{R}_{TS}$	$n_{Est}$
$\mathcal{P}_{G}$	$1/\varepsilon^2$	$1/\varepsilon^{(2s+1/2)/s}$	$n \wedge m \geq n_{GoF}$	$\varepsilon^2 n_{GoF}^2$
$\mathcal{P}_{H}$	$1/\varepsilon^2$	$1/\varepsilon^{(2\beta+d/2)/\beta}$	$n \wedge m \geq n_{GoF}$	$\varepsilon^2 n_{GoF}^2$
$\mathcal{P}_{Db}$	$1/\varepsilon^2$	$\sqrt{k}/arepsilon^2$	$n \wedge m \geq n_{GoF}$	$\varepsilon^2 n_{GoF}^2$
$\mathcal{P}_{D}$	$1/\varepsilon^2$	$\sqrt{k}/arepsilon^2$	$n \vee m \geq \frac{\sqrt{k}}{\varepsilon^2} \vee \frac{k^{2/3}}{\varepsilon^{4/3}} \asymp n_{TS}, \ n \wedge m \geq n_{GoF} \sqrt{\alpha}$	$\varepsilon^2  n_{\rm GoF}^2$

TABLE I PRIOR RESULTS ON TESTING AND ESTIMATION

shown in the last column of the table, which we regard as one of the main contributions of this paper.

## C. L<sup>2</sup>-robust likelihood-free hypothesis testing

Even before seeing Theorems 1 and 2 one might guess that estimation in TV followed by a robust hypothesis test should work whenever  $n \ge n_{\text{Est}}(c\varepsilon)$  for a small enough constant c and  $m \ge 1/\varepsilon^2$ . This strategy does indeed succeed, as can be deduced from the work of Huber and Birgé [36], [40] for the measures of separation  $d \in \{\text{TV}, \text{H}\}$  (see Remark 5 and Section II-C3). When d = TV, Scheffé's test also succeeds (see Section II-C1) as can be seen by a simple application of Chebyshev's inequality. An advantage of this approach is that it provides a solution to (LF) at the corner point B on Figure 1 that is robust to model misspecification with respect to d, naturally leading us to the question of robust likelihood-free hypothesis testing. As for (LF), suppose we observe samples X, Y, Z of size n, n, m from distributions belonging to the class  $\mathcal{P}$  with densities f, g, h with respect to some base measure  $\mu$ . Given any  $u \in \mathcal{P}$ , let  $B_u(\varepsilon, \mathcal{P}) \subseteq \mathcal{P}$  denote a region around u against which we wish to be robust. We compare the hypotheses

$$H_0: h \in \mathsf{B}_f(\varepsilon, \mathcal{P}), \mathsf{TV}(f, g) \ge \varepsilon \qquad \text{versus} \qquad H_1: h \in \mathsf{B}_q(\varepsilon, \mathcal{P}), \mathsf{TV}(f, g) \ge \varepsilon, \qquad (\mathsf{rLF})$$

and write  $\mathcal{R}_{\mathsf{rLF}}(\varepsilon, \mathcal{P}, \mathsf{B}_{\cdot})$  for the region of (n, m)-values for which  $(\mathsf{rLF})$  can be performed successfully, defined analogously to  $\mathcal{R}_{\mathsf{LF}}(\varepsilon, \mathcal{P})$ , noting that  $\mathcal{R}_{\mathsf{rLF}} \subseteq \mathcal{R}_{\mathsf{LF}}$  provided  $u \in \mathsf{B}_u$  for all  $u \in \mathcal{P}$ .

**Theorem 3.** There exists a universal constant c > 0 such that the equivalence

$$\mathcal{R}_{\mathsf{LF}}(\varepsilon, \mathcal{P}) \asymp \mathcal{R}_{\mathsf{rLF}}(\varepsilon, \mathcal{P}, \mathsf{B}_{\cdot})$$

holds for

(i)  $\mathcal{P} = \mathcal{P}_{\mathsf{H}}$  and  $\mathsf{B}_{u} = \{v : \|u - v\|_{2} \le c\varepsilon\}$ (ii)  $\mathcal{P} = \mathcal{P}_{\mathsf{G}}$  and  $\mathsf{B}_{\mu\theta} = \{\mu_{\theta'} : \theta' \in \mathcal{E}(s, C_{\mathsf{G}}), \|\theta - \theta'\|_{2} \le c\varepsilon\}$  and (iii)  $\mathcal{P} = \mathcal{P}_{\mathsf{Db}}$  and  $\mathsf{B}_{u} = \{v : \|u - v\|_{2} \le c\varepsilon/\sqrt{k}\}$  and (iv)  $\mathcal{P} = \mathcal{P}_{\mathsf{D}}$  and  $\mathsf{B}_{u} = \{v : \|u - v\|_{2} \le c\varepsilon/\sqrt{k}\}$ , up to  $\log(k)$ -factors.

#### D. Beyond total variation

Recall from Remark 5 the notation  $n_{GoF}(\varepsilon, d, P)$  etc. where separation is measured with respect to a general measure of discrepancy d instead of TV.

In recent work [52, Theorem 1] and [53, Corollary 3.4] it is shown that any test that first quantizes the data by a map  $\Phi : \mathcal{X} \to \{1, 2, \dots, M\}$  for some  $M \ge 2$  must decrease the Hellinger distance between the two hypotheses by a log factor in the worst case. This implies that for every class  $\mathcal{P}$  rich enough to contain such worst case examples,

a quantizing test (such as Scheffé's) can hope to achieve  $m \approx \log(1/\varepsilon)/\varepsilon^2$  at best, as opposed to the optimal  $m \approx 1/\varepsilon^2$ . Thus, if separation is assumed with respect to Hellinger distance, Scheffé's test should be avoided. This example shows that d can have nontrivial effects on the sample complexity of a specific test. Therefore, understanding the sample complexity of (LF) for d other than TV might lead to new algorithms and insights.

This leads us to the question: does a trade-off analogous to that identified in Theorem 1 hold for other choices of d, and H in particular? In the case of  $\mathcal{P}_G$  we obtain a simple, almost vacuous answer. From Lemma 2 it follows immediately that the results of Table I and Theorem 1 continue to hold for  $\mathcal{P}_G$  for any of  $d \in \{H, \sqrt{KL}, \sqrt{\chi^2}\}$ , to name a few.

**Lemma 2.** Let C > 0 be a constant. For any  $\theta \in \ell^2$  with  $\|\theta\|_2 \leq C$ 

$$\mathsf{FV}(\mu_{\theta},\mu_{0}) \asymp \mathsf{H}(\mu_{\theta},\mu_{0}) \asymp \sqrt{\mathsf{KL}(\mu_{\theta}\|\mu_{0})} \asymp \sqrt{\chi^{2}(\mu_{\theta}\|\mu_{0})} \asymp \|\theta\|_{2},$$

where  $\mu_{\theta} \triangleq \bigotimes_{i=1}^{\infty} \mathcal{N}(\theta_i, 1)$  and the implied constant depends on C.

*Proof.* By standard inequalities between divergences (see e.g. Lemma 6), omitting the argument  $(\mu_{\theta}, \mu_0)$  for simplicity we have

$$\mathsf{TV} \le \mathsf{H} \le \sqrt{\mathsf{KL}} \le \sqrt{\chi^2} = \sqrt{\exp(\|\theta\|_2^2) - 1} \lesssim \|\theta\|_2.$$

For the lower bound we obtain  $\mathsf{TV}(\mu_{\theta}, \mu_0) \ge 1 \land \|\theta\|_2 / 200 \gtrsim \|\theta\|_2$  by [54, Theorem 1.2].

The case of  $\mathcal{P}_{\mathsf{D}}$  is more intricate. Substantial recent progress [24], [51], [55], [56] has been made, where among others, the complexities  $n_{\mathsf{GoF}}$ ,  $n_{\mathsf{TS}}$ ,  $n_{\mathsf{Est}}$  for  $\mathsf{d} = \mathsf{H}$  are identified. Since our algorithm for (LF) is  $\|\cdot\|_2$ -based, we

$$\label{eq:product} \begin{array}{c|c|c|c|c|c|} \hline n_{\text{HT}} & n_{\text{GoF}} & n_{\text{TS}} & n_{\text{Est}} \\ \hline \hline \mathcal{P}_{\text{D}} & 1/\varepsilon^2 & \sqrt{k}/\varepsilon^2 & k^{2/3}/\varepsilon^{8/3} \wedge k^{3/4}/\varepsilon^2 & n_{\text{Gof}}^2 \varepsilon^2 \\ \hline \mathcal{P}_{\text{H}} & 1/\varepsilon^2 & ? & 1/\varepsilon^{2(\beta+d)/\beta} \\ \hline & \text{TABLE II} \\ \hline & \text{PRIOR RESULTS FOR d} = \text{H}. \end{array}$$

could immediately derive achievability bounds for  $\mathcal{R}_{LF}(\varepsilon, H, \mathcal{P}_D)$  via the inequality  $\|\cdot\|_2 \ge H^2/\sqrt{k}$ , however such a naive technique yields suboptimal results, and thus we omit it. Studying (LF) under Hellinger separation for  $\mathcal{P}_D$  and  $\mathcal{P}_{Db}$  is beyond the scope of this work.

Finally, we turn to  $\mathcal{P}_{\mathsf{H}}$ . Due to the nature of our proofs, the results of Theorem 1 easily generalize to  $\mathsf{d} = \|\cdot\|_p$  for any  $p \in [1, 2]$ . The simple reason for this is that (i) our algorithm is  $\|\cdot\|_2$ -based and  $\|\cdot\|_2 \ge \|\cdot\|_p$  by Jensen's inequality and (ii) the lower bound construction involves perturbations near 1, where all said norms are equivalent. In the important case  $\mathsf{d} = \mathsf{H}$  the estimation rate  $n_{\mathsf{Est}}(\varepsilon, \mathsf{H}, \mathcal{P}_{\mathsf{H}}) \asymp 1/\varepsilon^{2(\beta+d)/\beta}$  was obtained by Birgé [57], our contribution here is the study of  $n_{\mathsf{GoF}}$ .

**Theorem 4.** Let  $\mathcal{P} = \mathcal{P}_{\mathsf{H}}(\beta, d, C_{\mathsf{H}})$ . Then

$$n_{\mathsf{GoE}}(\varepsilon,\mathsf{H},\mathcal{P}) \geq 1/\varepsilon^{2(\beta+d/2)/\beta}$$

If in addition we assume that  $\beta \in (0, 1]$ , then

$$n_{\mathsf{GoF}}(\varepsilon,\mathsf{H},\mathcal{P}) \lesssim 1/\varepsilon^{2(\beta+d/2)/\beta}$$

and in particular,  $n_{\mathsf{Est}} \asymp n_{\mathsf{GoF}}^2 \varepsilon^2$ .

## IV. SKETCH PROOF OF MAIN RESULTS

In this section we briefly sketch the proofs of the main results of the paper.

#### A. Upper bounds for Theorems 1 to 4

Consider first the case when  $\mathbb{P}_X$  and  $\mathbb{P}_Y$  are supported on the discrete alphabet [k]. Let  $\hat{p}_X, \hat{p}_Y, \hat{p}_Z$  denote empirical probability mass functions based on the samples X, Y, Z of size n, n, m from  $\mathbb{P}_X, \mathbb{P}_Y, \mathbb{P}_Z$  respectively. Define the test statistic

$$T_{\mathsf{LF}} = \|\hat{p}_{\mathsf{X}} - \hat{p}_{\mathsf{Z}}\|_{2}^{2} - \|\hat{p}_{\mathsf{Y}} - \hat{p}_{\mathsf{Z}}\|_{2}^{2}$$

and the corresponding test  $\psi(X, Y, Z) = \mathbb{1}\{T_{\mathsf{LF}} \ge 0\}$ . The proof of Theorems 1 and 2 hinge on the precise calculation of the mean and variance of  $T_{\mathsf{LF}}$ . Due to symmetry it is enough to compute these under the null. The proof of the upper bound is then completed via Chebyshev's inequality: if n, m are such that  $(\mathbb{E}T_{\mathsf{LF}})^2 \gtrsim \operatorname{var}(T_{\mathsf{LF}})$  for large enough implied constant on the right then  $\psi$  tests (LF) successfully in the sense of (II.1).

#### **Proposition 2.** We have

$$\mathbb{E}T_{\mathsf{LF}} = \|p_{\mathsf{X}} - p_{\mathsf{Z}}\|_{2}^{2} - \|p_{\mathsf{Y}} - p_{\mathsf{Z}}\|_{2}^{2} + \frac{1}{n} \left(\|p_{\mathsf{Y}}\|_{2}^{2} - \|p_{\mathsf{X}}\|_{2}^{2}\right) \text{ and}$$
$$\operatorname{var}(T_{\mathsf{LF}}) \leq \left(\frac{1}{n} + \frac{1}{m}\right) \|(p_{\mathsf{X}} + p_{\mathsf{Y}} + p_{\mathsf{Z}})(p_{\mathsf{X}} - p_{\mathsf{Y}})^{2}\|_{1} + \left(\frac{1}{n^{2}} + \frac{1}{nm}\right) \|p_{\mathsf{X}} + p_{\mathsf{Y}} + p_{\mathsf{Z}}\|_{2}^{2}.$$

*Proof.* See Proposition 4 for a more general version of the result along with its proof.

Assuming that  $||p_X||_{\infty} \vee ||p_Y||_{\infty} \vee ||p_Z||_{\infty} \leq c_{\infty}$ , we obtain the bound

$$\operatorname{var}(T_{\mathsf{LF}}) \lesssim c_{\infty} \|p_{\mathsf{X}} - p_{\mathsf{Y}}\|_{2}^{2} \left(\frac{1}{n} + \frac{1}{m}\right) + kc_{\infty}^{2} \left(\frac{1}{n^{2}} + \frac{1}{nm}\right).$$
 (IV.1)

Crucially, there is no  $1/m^2$  term in the variance due to the cancellation of  $\|\hat{p}_Z\|_2^2$ . We also have the bound  $\|\|p_Y\|_2^2 - \|p_X\|_2^2 \le \sqrt{c_\infty} \|p_X - p_Y\|_2$  by Cauchy-Schwarz. In particular, the corresponding term in the expectation is smaller than  $\|p_X - p_Y\|_2^2$  as soon as  $n \gtrsim \sqrt{c_\infty}/\|p_X - p_Y\|_2$ , which is milder than the necessary (cf. Theorem 1) condition  $n \gtrsim n_{GoF}$ .

1) Bounded discrete distributions: We have  $c_{\infty} = \mathcal{O}(1/k)$  for the class  $\mathcal{P}_{Db}$  by definition. Taking  $p_Z \in \{p_X, p_Y\}$ , plugging (IV.1) into Chebyshev's inequality and the inequality  $\|\cdot\|_2 \ge \|\cdot\|_1/\sqrt{k}$  yields the minimax optimal rate for n and m for Theorem 1. The corresponding conclusion of Theorem 3 follows similarly using in addition the triangle inequality for  $\|\cdot\|_2$ , we defer the details to the appendix.

2) Smooth densities: Next we treat the class  $\mathcal{P}_{\mathsf{H}}$ . Divide  $[0,1]^d$  into into  $\kappa^d$  regular grid cells for some  $\kappa \in \mathbb{N}$ . Discretize the three samples X, Y, Z over this grid and simply apply the optimal test for  $\mathcal{P}_{\mathsf{Db}}$ . The following lemma, originally due to Ingster [11] controls the approximation error of this discretization.

**Lemma 3.** Let  $P_{\kappa}$  denote the  $L^2$  projection onto the space of functions constant on each grid cell. There exist constants c, c' > 0 depending only on  $d, \beta, C_{\mathsf{H}}$  such that for any  $u \in C(\beta, d, 2C_{\mathsf{H}})$  the following holds

$$||u||_2 \ge ||P_{\kappa}u||_2 \ge c||u||_2 - c'\kappa^{-\beta}.$$

Proof. See [42, Lemma 7.2].

Based on Lemma 3 we choose  $\kappa \simeq \varepsilon^{-1/\beta}$  which yields  $k = \kappa^d \simeq \varepsilon^{-d/\beta}$ . The resolution is chosen to ensure that the discrete approximation to any  $\beta$ -smooth density is sufficiently accurate, i.e.  $\approx \varepsilon$ -separation is maintained even after discretization. Once again, Chebyshev's inequality and Jensen's inequality  $\|\cdot\|_1 \le \|\cdot\|_2$  along with  $c_{\infty} = \Theta(1/k)$  yields the minimax optimal rates for Theorems 1 and 3.

Our proof of the upper bound in Theorem 4 follows by a reduction to goodness-of-fit testing for discrete distributions [56] under Hellinger separation, where it is known that  $n_{GoF}(\varepsilon, H, \mathcal{P}_D) \simeq \sqrt{k}/\varepsilon^2$ . The key step is to prove a result similar to Lemma 3 but for H instead of  $\|\cdot\|_2$ .

**Proposition 3.** Let 
$$f, g \in \mathcal{P}_{\mathsf{H}}(\beta, d, C_{\mathsf{H}})$$
 with  $\beta \in (0, 1]$  and suppose that  $\mathsf{H}(f, g) \gtrsim \varepsilon$ . Then

$$\mathsf{H}(f,g) \lesssim \mathsf{H}(P_{\kappa}f,P_{\kappa}g) \leq \mathsf{H}(f,g)$$

for  $\kappa \simeq \varepsilon^{-2/\beta}$  where the constants depend only on  $\beta, d, C_{\mathsf{H}}$ .

3) Gaussian sequence model: The case of  $\mathcal{P}_{G}$  is slightly different so we refer for details to the appendix.

4) Discrete distributions: Finally, we comment on  $\mathcal{P}_{D}$ . Here we can no longer assume that  $c_{\infty} = \mathcal{O}(1/k)$ , in fact  $c_{\infty} = \Omega(1)$  is possible. We get around this by utilizing the reduction based approach of [24], [34]. We take the first half of the data and compute  $B_i = 1 + \#\{j \in [(k \wedge n)/2] : X_j = i\} + \#\{j \in [(k \wedge n)/2] : Y_j = i\} + \#\{j \in [(k \wedge m)/2] : Z_j = i\}$  for each  $i \in [k]$ . Then, we divide bin *i* into  $B_i$  bins uniformly. This transformation preserves pairwise total variation, but reduces the  $\ell^{\infty}$ -norms of  $p_X, p_Y, p_Z$  with high probability, to order  $1/(k \wedge (n \vee m))$  (after an additional step that we omit here). We can then perform the usual test for the 'flattened' distributions, which we denote  $\tilde{p}_X, \tilde{p}_Y, \tilde{p}_Z$ , using the untouched half of the data. Chebyshev's inequality with a refined analysis of the variance yields the upper bound in Theorem 2.

Alternatively, for part of the trade-off (namely the regime  $n \le m$ ) we can use the reduction Proposition 1 to two-sample testing with unequal sample size to get the optimal upper (and lower) bound.

It is insightful to interpret the 'flattening' procedure followed by  $\ell^2$ -distance comparison as a one-step procedure that simply compares a different divergence of the empirical measures. Intuitively, in contrast to the regular classes, one needs to mitigate the effect of potentially massive differences in the empirical counts on bins  $i \in [k]$  where both  $p_X(i)$  and  $p_Y(i)$  are large but their difference  $|p_X(i) - p_Y(i)|$  is moderate. Let  $LC_\lambda$  be the 'weighted Le-Cam divergence' which we define as  $LC_\lambda(p||q) = \sum_i (p_i - q_i)^2/(p_i + (1 + \lambda)q_i)$  for two pmfs p, q. Taking expectation with respect to  $B = (B_1, \ldots, B_k)$  of the flattened measures  $\tilde{p}_X, \tilde{p}_Y$  we have (heuristically)

$$\begin{split} \mathbb{E}_B \| \tilde{p}_{\mathsf{X}} - \tilde{p}_{\mathsf{Z}} \|_2^2 &= \begin{cases} 0 & \text{if } p_{\mathsf{X}} = p_{\mathsf{Z}} \\ \mathbb{E}_B \sum_{i \in [k]} (p_{\mathsf{X}}(i) - p_{\mathsf{Z}}(i))^2 / B_i & \text{if } p_{\mathsf{Y}} = p_{\mathsf{Z}} \\ &\approx \frac{1}{n \wedge k} \mathsf{LC}_\lambda(p_{\mathsf{X}} \| p_{\mathsf{Z}}), \end{split}$$

where  $\lambda = \frac{m \wedge k}{n \wedge k}$ . A similar expression holds for  $\mathbb{E}_B \| \tilde{p}_{\mathsf{Y}} - \tilde{p}_{\mathsf{Z}} \|_2^2$ . Therefore, on average, the statistic  $T_{\mathsf{LF}}$  after flattening can be thought of as

$$T_{\mathsf{LF}} \approx \frac{1}{n \wedge k} \Big( \mathsf{LC}_{\lambda}(\hat{p}_{\mathsf{X}} \| \hat{p}_{\mathsf{Z}}) - \mathsf{LC}_{\lambda}(\hat{p}_{\mathsf{Y}} \| \hat{p}_{\mathsf{Z}}) \Big).$$
(IV.2)

Performing the test in two steps (flattening first and comparing  $\ell^2$  distances) is a proof device, and we expect the test that directly compares, say, the Le-Cam divergence of the empirical pmfs to have the same minimax optimal sample complexity. Such a one-shot approach is used for example in the paper [44] for two-sample testing. While Ingster [11] only considers goodness-of-fit testing to the uniform distribution, his notation also suggests the idea of normalizing by the bin mass under the null.

## B. Lower bounds for Theorems 1 to 4

Proposition 1 immediately yields tight lower bounds on n and m. Namely, (III.1) gives  $m \gtrsim 1/\varepsilon^2$  and (III.3) gives  $n \gtrsim n_{GoF}(\varepsilon, \mathcal{P})$ . We now turn to the more challenging task of obtaining a lower bound on the interaction term  $m \cdot n$ . For this let us first introduce some well known results used to prove minimax lower bounds. Suppose that we have two (potentially composite) hypotheses  $H_0, H_1$  that we test against each other. Our strategy for proving lower bounds relies on the method of two fuzzy hypotheses [48], which is a generalization of le-Cam's two point method. Write  $\mathcal{M}(\mathcal{X})$  for the set of probability measures on the set  $\mathcal{X}$ .

**Lemma 4.** Take two hypotheses  $H_i \subseteq \mathcal{M}(\mathcal{X})$  and random  $P_i \in \mathcal{M}(\mathcal{X})$  with  $\mathbb{P}(P_i \in H_i) > 0$ . Then

$$2 \inf_{\psi} \max_{i=0,1} \sup_{P \in H_i} P(\psi \neq i) \ge 1 - \mathsf{TV}(\mathbb{E}P_0, \mathbb{E}P_1) - \sum_i \mathbb{P}(P_i \notin H_i),$$

where the infimum is over all tests  $\psi : \mathcal{X} \to \{0, 1\}$ .

*Proof.* Let  $\tilde{P}_i$  be distributed as  $P_i | \{P_i \in H_i\}$ . Then for any set  $A \subset \mathcal{X}$  we have

$$\left| \mathbb{E}\tilde{P}_i(A) - \mathbb{E}P_i(A) \right| = \mathbb{P}(P_i \notin H_i) \left| \mathbb{E}[P_i(A)|P_i \in H_i] - \mathbb{E}[P_i(A)|P_i \notin H_i] \right| \le \mathbb{P}(P_i \notin H_i)$$

In particular,  $\mathsf{TV}(\mathbb{E}\tilde{P}_0,\mathbb{E}\tilde{P}_1) \leq \mathsf{TV}(\mathbb{E}P_0,\mathbb{E}P_1) + \sum_i \mathbb{P}(P_i \notin H_i)$ . Therefore, for any  $\psi$ 

$$\max_{i=0,1} \sup_{\mathbb{P}_i \in H_i} \mathbb{P}_i(\psi \neq i) \ge \frac{1}{2} (1 - \mathsf{TV}(\mathbb{E}\tilde{P}_0, \mathbb{E}\tilde{P}_1)) \ge \frac{1}{2} \left( 1 - \mathsf{TV}(\mathbb{E}P_0, \mathbb{E}P_1) - \sum_i \mathbb{P}(P_i \notin H_i) \right).$$

For clarity, we formally state (LF) as testing between the hypotheses

$$H_{0} = \{ \mathbb{P}_{\mathsf{X}}^{\otimes n} \otimes \mathbb{P}_{\mathsf{Y}}^{\otimes n} \otimes \mathbb{P}_{\mathsf{X}}^{\otimes m} : \mathbb{P}_{\mathsf{X}}, \mathbb{P}_{\mathsf{Y}} \in \mathcal{P}, \, \mathsf{TV}(\mathbb{P}_{\mathsf{X}}, \mathbb{P}_{\mathsf{Y}}) \ge \varepsilon \}$$
  
versus  
$$H_{1} = \{ \mathbb{P}_{\mathsf{X}}^{\otimes n} \otimes \mathbb{P}_{\mathsf{Y}}^{\otimes n} \otimes P_{\mathsf{Y}}^{\otimes m} : \mathbb{P}_{\mathsf{X}}, \mathbb{P}_{\mathsf{Y}} \in \mathcal{P}, \, \mathsf{TV}(\mathbb{P}_{\mathsf{X}}, \mathbb{P}_{\mathsf{Y}}) \ge \varepsilon \}.$$
(IV.3)

The lower bounds of Theorem 3 follow from those for Theorems 1 and 2 so we only focus on the latter case. 1) Smooth densities: For concreteness let us focus on the case of  $\mathcal{P} = \mathcal{P}_{\mathsf{H}}$ . We take  $\mathbb{P}_0$  to be uniform on  $[0,1]^d$ and  $\mathbb{P}_{\eta}$  to have density

$$p_{\eta} = 1 + \sum_{j \in [\kappa]^d} \eta_j h_j \tag{IV.4}$$

with respect to  $\mathbb{P}_0$ . Here  $\kappa \in \mathbb{N}$ , each  $\eta \in \{\pm 1\}^{\kappa^d}$  is uniform and  $h_j$  is a bump function supported on the j'th cell of the regular grid of size  $\kappa^d$  on  $[0,1]^d$ . The parameters  $\kappa, h_j$  of the construction are set in a way to ensure  $\mathbb{P}_\eta \in \mathcal{P}_{\mathsf{H}}$  and  $\mathsf{TV}(\mathbb{P}_0, \mathbb{P}_\eta) \geq \varepsilon$  with probability one over  $\eta$ . We have

$$1 + \chi^{2}(\mathbb{E}_{\eta}\mathbb{P}_{\eta}^{\otimes m} \| \mathbb{P}_{0}^{\otimes m}) = \int_{[0,1]^{dm}} \left( \mathbb{E}_{\eta} \prod_{i=1}^{n} p_{\eta}(x_{i}) \right)^{2} \mathrm{d}x_{1} \dots \mathrm{d}x_{m}$$
$$= \mathbb{E}_{\eta\eta'} \langle p_{\eta}, p_{\eta'} \rangle_{L^{2}}^{m}$$
$$= \mathbb{E}(1 + \|h_{1}\|_{2}^{2} \langle \eta, \eta' \rangle)^{m}$$
$$\leq \exp(m^{2} \|h_{1}\|_{2}^{4} \kappa^{d}),$$
(IV.5)

where  $\eta, \eta'$  are i.i.d. uniform and we assume  $||h_1||_2 = ||h_j||_2$  for all  $j \in [\kappa]^d$ . The above approach is what Ingster used in his seminal paper [11] on goodness-of-fit testing, which we adapt to likelihood-free hypothesis testing (IV.3). Take  $P_0 = \mathbb{P}_{\eta}^{\otimes n} \otimes \mathbb{P}_0^{\otimes m} \otimes \mathbb{P}_{\eta}^{\otimes m}$  and  $P_1 = \mathbb{P}_{\eta}^{\otimes n} \otimes \mathbb{P}_0^{\otimes m} \otimes \mathbb{P}_0^{\otimes m}$  in Lemma 4. Bounding  $\mathsf{TV}(\mathbb{E}P_0, \mathbb{E}P_1)$  proceeds in multiple steps: first, we drop the Y-sample using the data-processing inequality. Then, we use Pinsker's inequality and the chain rule to bound TV by the KL divergence of Z conditioned on X. We bound KL by  $\chi^2$ , arriving at the same equation (IV.5). However, the mixing parameters  $\eta, \eta'$  are no longer independent, instead, given X they're independent from the posterior. In the remaining steps we use the fact that the posterior factorizes over the bins and the calculation is reduced to just a single bin where it can be done explicitly.

Let us now turn to the lower bound in Theorem 4. The difference in the rate is a consequence of the fact that H and TV behave differently for densities near zero. Inspired by this, we slightly modify the construction (IV.4) by putting the perturbations at density level  $\varepsilon^2$  as opposed to 1. Bounding TV then proceeds analogously to the steps outlined above and [11].

2) Bounded discrete distributions: The construction is entirely analogous to the case of  $\mathcal{P}_{H}$  and we refer to the appendix for details. In the computer science community the construction of  $p_{\eta}$  is attributed to Paninski [58].

3) Gaussian sequence model: As for the upper bounds, the case of  $\mathcal{P}_{\mathsf{G}}$  is somewhat different from the others. Here the null distribution  $\mathbb{P}_0$  is the no signal case  $\bigotimes_{i=1}^{\infty} \mathcal{N}(0,1)$  while the alternative is  $\mathbb{P}_{\theta} = \bigotimes_{i=1}^{\infty} \mathcal{N}(\theta_i,1)$  where  $\theta$  has prior distribution  $\bigotimes_{i=1}^{\infty} \mathcal{N}(0,\gamma_i)$  for an appropriate sequence  $\gamma \in \mathbb{R}^{\mathbb{N}}$ . We refer to the appendix for more details. 4) Discrete distributions: Once again, the irregular case  $\mathcal{P}_{\mathsf{D}}$  requires special consideration. Clearly the lower bound for  $\mathcal{P}_{\mathsf{Db}}$  carries over. However, in the large alphabet regime  $k \gtrsim 1/\varepsilon^4$  said lower bound becomes suboptimal, and we need a new construction, for which we utilize the moment-matching based approach of Valiant [59] as a black-box. The adversarial construction is derived from that used for two-sample testing by Valiant, namely the pair  $(\mathbb{P}_{\mathsf{X}}, \mathbb{P}_{\mathsf{Y}})$  is chosen uniformly at random from  $\{(p \circ \pi, q \circ \pi)\}_{\pi \in S_k}$ . Here we write  $S_k$  for the symmetric group on [k] and

$$p(i) = \begin{cases} \frac{1-\varepsilon}{n} & \text{for } i \in [n] \\ \frac{4\varepsilon}{k} & \text{for } i \in [\frac{k}{2}, \frac{3k}{4}] \\ 0 & \text{otherwise,} \end{cases}$$

where we assume that  $m \le n \le k/2$  and define q(i) = p(i) for  $i \in [k/2-1]$  and q(i) = p(3k/2-i) for  $i \in [k/2, k]$ . This construction gives a lower bound matching our upper bound in the regime  $m \le n \le k$ . The final piece of the puzzle follows by the reduction from two-sample testing with unequal sample size (III.6), as this shows that likelihood-free hypothesis testing is at least as hard as two-sample testing in the  $n \le m$  regime, and known lower bounds on the sample complexity of two-sample testing [23] (see also Table I) let us conclude.

## V. OPEN PROBLEMS

A natural follow-up direction to the present paper would be to study multiple hypothesis testing where  $\mathbb{P}_X$  and  $\mathbb{P}_Y$  are replaced by  $\mathbb{P}_{X_1}, \ldots, \mathbb{P}_{X_M}$  with corresponding hypotheses  $H_1, \ldots, H_M$ . The geometry of the family  $\{\mathbb{P}_{X_j}\}_{j \in [M]}$  might have interesting effects on the sample complexities.

## **Open problem 1.** Study the dependence on M > 2 of likelihood-free testing with M hypotheses.

Another possible avenue of research is the study of local minimax/instance optimal rates, which is the focus of recent work [45], [60]–[63] in the case of goodness-of-fit and two-sample testing.

#### **Open problem 2.** Define and study the local minimax rates of likelihood-free hypothesis testing.

Our discussion of the Hellinger case in Section III-D is quite limited, natural open problems in this direction include the following.

**Open problem 3.** Let  $\mathcal{P} \in \{\mathcal{P}_{\mathsf{H}}(\beta, d, C_{\mathsf{H}}), \mathcal{P}_{\mathsf{Db}}(k, C_{\mathsf{Db}}), \mathcal{P}_{\mathsf{D}}(k)\}.$ 

- (i) Study  $n_{GoF}$  and  $n_{TS}$  for  $\mathcal{P}$  under Hellinger separation.
- (ii) Determine the trade-off  $\mathcal{R}_{\mathsf{LF}}$  for  $\mathcal{P}$  under Hellinger separation.

More ambitiously, one might ask for a characterization of 'regular' models  $(\mathcal{P}, d)$  for which goodness-of-fit testing and two-sample testing are equally hard and the region  $\mathcal{R}_{LF}$  is given by the trade-off in Theorem 1.

**Open problem 4.** Find a general family of 'regular' models  $(\mathcal{P}, d)$  for which

$$\begin{split} n_{\mathsf{GoF}}(\varepsilon,\mathsf{d},\mathcal{P}) &\asymp n_{\mathsf{TS}}(\varepsilon,\mathsf{d},\mathcal{P}) \text{ and } \\ \mathcal{R}_{\mathsf{LF}}(\varepsilon,\mathsf{d},\mathcal{P}) &\asymp \{m \geq 1/\varepsilon^2, n \geq n_{\mathsf{GoF}}(\varepsilon,\mathsf{d},\mathcal{P}), mn \geq n_{\mathsf{GoF}}^2(\varepsilon,\mathsf{d},\mathcal{P})\} \end{split}$$

Recent follow-up work [32] showed that Scheffé's test is also minimax optimal and achieves the entire trade-off in Figure 1. It appears that the optimality of Scheffé's test is a consequence of the minimax point of view. Basically, in the worst-case the log-likelihood ratio between the hypotheses is close to being binary, hence quantizing it to  $\{0, 1\}$  does not lose optimality. Consequently, an important future direction is to better understand the competitive properties of various tests and studying some notion of regret, see [64] for prior related work.

**Open problem 5.** Study the competitive optimality of likelihood-free hypothesis testing algorithms, and Scheffé's test in particular.

## APPENDIX A UPPER BOUNDS OF THEOREM 1 AND 2

Let  $\mu$  be a measure on the measurable space  $\mathcal{X}$  and let  $\{\phi_i\}_{i \in [r]}$  be a sequence of orthonormal functions in  $L^2(\mu)$ . For  $f \in L^2(\mu)$  define its projection onto the span of  $\{\phi_1, \ldots, \phi_r\}$  as

$$P_r(f) \triangleq \sum_{i \in [r]} \langle f \phi_i \rangle \phi_i,$$

where we write  $\langle \cdot \rangle$  for integration with respect to  $\mu$  and  $\|\cdot\|_p$  for  $\|\cdot\|_{L^p(\mu)}$ . Given an i.i.d. sample  $X = (X_1, \ldots, X_n)$  from some density f, define its empirical projection as

$$\hat{P}_r[X] \triangleq \sum_{i \in [r]} \left( \frac{1}{n} \sum_{j=1}^n \phi_i(X_j) \right) \phi_i.$$

Then, our statistic reads

$$T_{\mathsf{LF}} = \|\hat{P}_r[X] - \hat{P}_r[Z]\|_2^2 - \|\hat{P}_r[Y] - \hat{P}_r[Z]\|_2^2,$$
(A.1)

for an appropriate choice of  $\mu$  and  $\{\phi_j\}_{j\geq 1}$  depending on the class  $\mathcal{P}$ . Before calculating the mean and variance, we separate out the diagonal terms in  $T_{\mathsf{LF}}$  thereby decomposing the statistic into two terms:

$$T_{\mathsf{LF}} \triangleq T_{\mathsf{LF}}^{-\mathsf{d}} + \underbrace{\frac{1}{n^2} \sum_{i \in [r]} \sum_{j \in [n]} \left( \phi_i^2(X_j) - \phi_i^2(Y_j) \right)}_{\triangleq D}.$$
(A.2)

To ease notation in the results below, we define the quantities

$$A_{fgh} = \langle f [P_r(g-h)]^2 \rangle$$
  

$$B_{fg} = \sum_{i=1}^r \langle f \phi_i P_r(g \phi_i) \rangle$$
(A.3)

for  $f, g, h \in L^2(\mu)$ , assuming the quantities involved are well-defined. We are ready to state our meta-result from which we derive all our likelihood-free hypothesis testing upper bounds.

**Proposition 4.** Let f, g, h denote probability densities on  $\mathcal{X}$  with respect to  $\mu$ , and suppose we observe independent samples X, Y, Z of size n, n, m from f, g, h respectively. Recall the test statistic

$$T_{\mathsf{LF}}^{-\mathsf{d}} = \sum_{i=1}^{r} \left\{ \frac{1}{n^2} \sum_{j \neq j'}^{n} \phi_i(X_j) \phi_i(X_{j'}) - \frac{1}{n^2} \sum_{j \neq j'}^{n} \phi_i(Y_j) \phi_i(Y_{j'}) - \frac{2}{nm} \sum_{j=1}^{n} \sum_{u=1}^{m} \phi_i(X_j) \phi_i(Z_u) + \frac{2}{nm} \sum_{j=1}^{n} \sum_{u=1}^{m} \phi_i(Y_j) \phi_i(Z_u) \right\}.$$

Then

$$\begin{split} \mathbb{E}T_{\mathsf{LF}}^{-\mathsf{d}} &= \|P_r(f-h)\|_2^2 - \|P_r(g-h)\|_2^2 + \frac{1}{n}(\|P_r[g]\|_2^2 - \|P_r[f]\|_2^2) \\ \operatorname{var}(T_{\mathsf{LF}}^{-\mathsf{d}}) &\lesssim \frac{A_{ffh} + A_{ggh}}{n} + \frac{A_{hfg}}{m} \\ &+ \frac{\|f+g+h\|_2^4 + |B_{fh}| + |B_{gh}|}{nm} + \frac{|B_{ff}| + |B_{gg}| + \|f+g+h\|_2^3(1 + \|f+g+h\|)}{n^2}, \end{split}$$

1

where the implied constant is universal.

*Remark* 8. Proposition 4 is applied to (LF) by considering the test  $\mathbb{1}(T_{LF}^{-d} \ge 0)$ . To prove that the test performs well we show that  $T_{LF}^{-d}$  concentrates around its mean by Chebyshev's inequality. For this we find sufficient conditions on the sample sizes n, m so that  $(\mathbb{E}T_{LF}^{-d})^2 \ge \operatorname{var}(T_{LF}^{-d})$  for a small enough implied constant on the left.

*Remark* 9. While Proposition 4 is enough to conclude the proof of our main theorems, notice that it uses the statistic  $T_{LF}^{-d}$  which has the diagonal terms removed. For completeness we show that the test  $\mathbb{1}\{T_{LF} \ge 0\}$  is also minimax optimal, i.e. the diagonal terms (*D* in (A.2)) can be included without degrading performance.

## A. The class $\mathcal{P}_{\mathsf{Db}}$

**Proposition 5.** For constant  $c, c_r > 0$  independent of  $\varepsilon$  and k,

$$\mathcal{R}_{\mathsf{rLF}}(\varepsilon, \mathcal{P}_{\mathsf{Db}}(k, C_{\mathsf{Db}}), \mathsf{B}_{\cdot}) \supseteq c\left\{m \ge 1/\varepsilon^2, n \ge \sqrt{k}/\varepsilon^2, mn \ge k/\varepsilon^4\right\},$$

where  $\mathsf{B}_u = \{ u \in \mathcal{P}_{\mathsf{Db}} : \|u - v\|_2 \le c_{\mathsf{r}} \varepsilon / \sqrt{k} \}.$ 

*Proof.* Choice of  $\mu$  and  $\phi$ . Take  $\mathcal{X} = [k]$  and let  $\mu = \sum_{i=1}^{k} \delta_i$  be the counting measure. Let  $\phi_i(j) = \mathbb{1}_{\{i=j\}}$  and choose r = k so that  $P_r = P_k$  is the identity. By the Cauchy-Schwarz inequality  $\|h\|_1 \leq \sqrt{k} \|h\|_2$  for all  $h \in \mathbb{R}^k$ .

**Applying Proposition 4.** Recall the notation of Proposition 4, so that f, g, h are the pmfs of  $\mathbb{P}_X, \mathbb{P}_Y, \mathbb{P}_Z$  respectively. We analyse the performance of the test  $\mathbb{1}\{T_{\mathsf{LF}}^{-\mathsf{d}} \ge 0\}$  under the null hypothesis, the proof under the alternative is analogous. Choosing the radius of robustness as  $c_r < 1$ , the inequality  $||f - h||_2 \le \frac{c_r}{2} ||f - g||_2$  along with the reverse triangle inequality gives us

$$\begin{aligned} \|g-h\|_{2}^{2} - \|f-h\|_{2}^{2} &\geq (\|f-g\|_{2} - \|f-h\|_{2})^{2} - \|f-h\|_{2}^{2} \\ &\geq \|f-g\|_{2}^{2}(1-c_{\mathsf{r}}). \end{aligned}$$

Notice that now  $-\mathbb{E}T_{LF}^{-d} \ge (1-c_r)||f-g||_1/k + R$ , where the residual term R can be bounded as

$$\begin{split} |R| &= \bigg| \frac{\|f\|_2^2 - \|g\|_2^2}{n} \\ &\lesssim \frac{\|f - g\|_2}{n\sqrt{k}}. \end{split}$$

Thus,  $|R| \leq (1 - c_r) ||f - g||_2^2/2$  provided  $n \gtrsim 1/(||f - g||_2\sqrt{k})$  which in turn is implied by  $n \gtrsim 1/\varepsilon$ . Therefore, under the assumption that  $n \gtrsim 1/\varepsilon$ , we obtain the lower bound  $-\mathbb{E}T_{\mathsf{LF}}^{-\mathsf{d}} \gtrsim ||f - g||_2^2$ . Turning towards the variance, we apply Proposition 4 to see that

$$\operatorname{var}(T_{\mathsf{LF}}^{-\mathsf{d}}) \lesssim \frac{\|f - g\|_2^2}{k} \left(\frac{1}{n} + \frac{1}{m}\right) + \frac{1}{k} \left(\frac{1}{n^2} + \frac{1}{nm}\right),$$
 (A.4)

where we use the trivial bounds

$$\|f + g + h\|_{2} \lesssim 1/\sqrt{k}$$
$$|B_{ff}| + |B_{gg}| + |B_{fh}| + |B_{gh}| \lesssim \frac{1}{k}$$
$$A_{ffh} + A_{ggh} + A_{hfg} \lesssim \frac{1}{k} \|f - g\|_{2}^{2}$$

Applying Chebyshev's inequality and looking at each term separately in (A.4) yields the desired bounds on n, m.

**The diagonal.** While the above test using  $T_{LF}^{-d}$  already achieves the minimax optimal sample complexity, here we show for completeness that the diagonal D (cf. (A.2)) can be included without degrading the test's performance. Indeed, we have

$$D = \frac{1}{n^2} \sum_{i \in [r]} \sum_{j \in [n]} \left( \mathbb{1}\{X_j = i\}^2 - \mathbb{1}\{Y_j = i\}^2 \right)$$
  
= 0

almost surely. Therefore, trivially, the test  $\mathbb{1}\{T_{\mathsf{LF}} \geq 0\}$  is minimax optimal.

## B. The class $\mathcal{P}_{H}$

**Proposition 6.** For constants  $c, c_r > 0$  independent of  $\varepsilon$ ,

$$\mathcal{R}_{\mathsf{rLF}}(\varepsilon, \mathcal{P}_{\mathsf{H}}(\beta, d, C_{\mathsf{H}}), \mathsf{B}_{\cdot}) \supseteq c\left\{ m \ge 1/\varepsilon^2, n \ge 1/\varepsilon^{(2\beta + d/2)/\beta}, mn \ge 1/\varepsilon^{2(2\beta + d/2)/\beta} \right\},\tag{A.5}$$

where  $B_u = \{v \in \mathcal{P}_{\mathsf{H}} : \|v - u\|_2 \le c_{\mathsf{r}}\varepsilon\}.$ 

*Proof.* Choice of  $\mu$  and  $\phi$ . Take  $\mathcal{X} = [0, 1]^d$ , let  $\mu$  the Lebesgue measure on  $\mathcal{X}$ . Let  $\{\phi_i\}_{1 \le i \le \kappa^d}$  be the indicators of the cells of the regular grid with  $\kappa^d$  bins, normalized to have  $L^2(\mu)$ -norm equal to 1 (i.e. the indicator is multiplied by  $\kappa^d$ , one over the volume of one grid cell). By [42, Lemma 7.2] for any resolution  $r = \kappa^d$  and  $u \in \mathcal{C}(\beta, d, 2C_{\mathsf{H}})$  we have

$$||P_r(u)||_2 \ge c_1 ||u||_2 - c_2 \kappa^{-\beta}$$
(A.6)

for constants  $c_1, c_2 > 0$  that don't depend on r. In particular, if  $||u||_1 \ge 2\varepsilon$  then taking  $\kappa^{-\beta} = c_1 \varepsilon / c_2$  in (A.6) ensures with the help of Jensen's inequality that  $||P_r(u)||_2 \ge c_1 \varepsilon$ .

Applying Proposition 4. Recall the notation of Proposition 4 so that f, g, h are the  $\mu$ -densities of  $\mathbb{P}_X, \mathbb{P}_Y, \mathbb{P}_Z$ . We analyse the performance of the test  $\mathbb{1}\{T_{\mathsf{LF}}^{-\mathsf{d}} \ge 0\}$  under the null hypothesis, the proof under the alternative is analogous. Choosing the radius of robustness  $c_{\mathsf{r}} < c_1/2$ , applying the inequality  $\|P_r(f-h)\|_2 \le \frac{c_r}{c_1}\|P_r(f-g)\|_2$  (by taking u = f - g in (A.6)) we obtain

$$||P_r(g-h)||_2^2 - ||P_r(f-h)||_2^2 \ge ||P_r(f-g)||_2^2(1-2\frac{c_r}{c_1})$$

Thus,  $-\mathbb{E}T_{\mathsf{LF}}^{-\mathsf{d}} \ge (1 - 2c_{\mathsf{r}}/c_1) \|P_r(f - g)\|_2^2 + R$  where the residual term R can be bounded as

$$\begin{split} |R| &= \bigg| \frac{\|f\|_2^2 - \|g\|_2^2}{n} \\ &\lesssim \frac{\|f - g\|_2}{n}. \end{split}$$

Thus,  $|R| \leq (1 - 2c_r/c_1) ||f - g||_2^2/2$  provided  $n \gtrsim 1/||f - g||_2$  which in turn is implied by  $n \gtrsim 1/\varepsilon$ . Therefore, under the assumption that  $n \gtrsim 1/\varepsilon$ , we may assume that  $-\mathbb{E}T_{\mathsf{LF}}^{-\mathsf{d}} \gtrsim ||P_r(f - g)||_2^2$ . Turning to the variance, using Proposition 4 we obtain

$$\operatorname{var}(T_{\mathsf{LF}}^{-\mathsf{d}}) \lesssim \|P_r(f-g)\|_2^2 \left(\frac{1}{n} + \frac{1}{m}\right) + \varepsilon^{-d/\beta} \left(\frac{1}{n^2} + \frac{1}{nm}\right),\tag{A.7}$$

where we apply the trivial inequalities

$$||f + g + h||_2 \lesssim 1$$
  
$$|B_{ff}| + |B_{gg}| + |B_{fh}| + |B_{gh}| \lesssim r \approx \varepsilon^{-d/\beta}$$
  
$$A_{ffh} + A_{ggh} + A_{hfg} \lesssim ||P_r(f - g)||_2^2.$$

Applying Chebyshev's inequality and looking at each term separately in (A.7) yields the desired bounds on n, m.

The diagonal. While the above test using  $T_{LF}^{-d}$  already achieves the minimax optimal sample complexity, for completeness we also note that including the diagonal terms D (cf. (A.2)) doesn't degrade performance. The fact that D = 0 almost surely follows analogously to the case of  $\mathcal{P}_{Db}$ .

C. The class  $\mathcal{P}_{G}$ 

**Proposition 7.** For constants  $c, c_r > 0$  independent of  $\varepsilon$ ,

$$\mathcal{R}_{\mathsf{rLF}}(\varepsilon, \mathcal{P}_{\mathsf{G}}(s, C_{\mathsf{G}}), \mathsf{B}_{\cdot}) \supseteq c\left\{ m \ge 1/\varepsilon^2, n \ge 1/\varepsilon^{(2s+1/2)/s}, mn \ge 1/\varepsilon^{2(2s+1/2)/s} \right\},$$

where  $\mathsf{B}_{\mu_{\theta}} = \{\mu_{\theta'} : \theta' \in \mathcal{E}(s, C_{\mathsf{G}}), \|\theta - \theta'\|_2 \le c_{\mathsf{r}}\varepsilon\}.$ 

*Proof.* Choosing  $\mu$  and  $\phi$ . Let  $\mathcal{X} = \mathbb{R}^{\mathbb{N}}$  be the set of infinite sequences and take as the base measure  $\mu = \bigotimes_{d=1}^{\infty} \mathcal{N}(0, 1)$ , the infinite dimensional standard Gaussian. For  $\theta \in \ell^2$  write  $\mu_{\theta} = \bigotimes_{d=1}^{\infty} \mathcal{N}(\theta_i, 1)$  so that  $\mu_0 = \mu$ . Take the orthonormal functions  $\phi_i(x) = x_i$  for  $i \geq 1$ , so that

$$P_r\left(\frac{\mathrm{d}\mu_\theta}{\mathrm{d}\mu}\right) = \sum_{i=1}^r x_i \theta_i.$$

Let  $\theta, \theta' \in \mathcal{E}(s, C_{\mathsf{G}})$  with  $\mathsf{TV}(\mu_{\theta}, \mu_{\theta'}) \geq \varepsilon$ . By Pinsker's inequality  $\varepsilon \leq \|\theta - \theta'\|_2$  and the following holds:

$$\|P_r\left(\frac{\mathrm{d}\mu_{\theta}}{\mathrm{d}\mu} - \frac{\mathrm{d}\mu_{\theta'}}{\mathrm{d}\mu}\right)\|_2^2 = \sum_{i=1}^r (\theta_i - \theta_i')^2 \ge \varepsilon^2 - r^{-2s} \sum_{i>r} (\theta_i - \theta_i')^2 i^{2s} \ge \varepsilon^2 - 4C_{\mathsf{G}}^2 r^{-2s}.$$
 (A.8)

In particular, taking  $r \simeq \varepsilon^{-1/s}$  for a constant independent of  $\varepsilon$ , the above is lower bounded by  $\varepsilon^2/4$ .

**Applying Proposition 4.** Recall the notation of Proposition 4, and let f, g, h be the  $\mu$ -densities of  $\mathbb{P}_{\mathsf{X}} = \mu_{\theta}, \mathbb{P}_{\mathsf{Y}} = \mu_{\theta'}, \mathbb{P}_{\mathsf{Z}} = \mu_{\theta''}$  respectively. We analyse the test  $\mathbb{1}\{T_{\mathsf{LF}}^{-\mathsf{d}}$  only under the null hypothesis, as the analysis under the alternative is analogous. Taking the radius of robustness  $c_{\mathsf{r}} < 1/4$ , using the inequality  $||P_r(f-h)||_2 \le 2c_{\mathsf{r}}||P_r(f-g)||_2$  we see that

$$\|P_r(g-h)\|_2^2 - \|P_r(f-h)\|_2^2 \ge (1-4c_r)\|P_r(f-g)\|_2^2.$$
(A.9)

Notice that  $\mathbb{E}T_{\mathsf{LF}}^{-\mathsf{d}} \ge (1 - 4c_{\mathsf{r}}) \|P_r(f - g)\|_2^2 + R$ , where the residual term R can be bounded as

$$|R| = \left| \frac{\|P_r(f)\|_2^2 - \|P_r(g)\|_2^2}{n} \right| \le \frac{\|P_r(f-g)\|_2}{n}.$$

Thus,  $|R| \leq (1 - 4c_r) \|P_r(f - g)\|_2^2/2$  provided  $n \gtrsim 1/\|P_r(f - g)\|_2$  which in turn is implied by  $n \gtrsim 1/\varepsilon$ . Therefore, under the assumption that  $n \gtrsim 1/\varepsilon$ , we obtain  $-\mathbb{E}T_{\mathsf{LF}}^{-\mathsf{d}} \gtrsim \|P_r(f - g)\|_2^2$ . Straightforward calculations involving Gaussian random variables produce

$$\begin{aligned} A_{fgh} &= \sum_{ij}^{r} (\mathbb{1}(i=j) + \theta_{i}\theta_{j})(\theta_{i}' - \theta_{i}'')(\theta_{j}' - \theta_{j}'') \leq (1 + C_{\mathsf{G}}^{2}) \|P_{r}(g-h)\|_{2}^{2} \\ &\lesssim \|P_{r}(g-h)\|_{2}^{2} \\ \|f\|_{2} &= \exp\left(\frac{1}{2}\|\theta\|_{2}^{2}\right) \leq \exp(C_{\mathsf{G}}^{2}/2) \\ &\lesssim 1 \\ B_{fg} &= \sum_{i=1}^{r} \left(1 + \theta_{i}^{2} + {\theta_{i}'}_{i}^{2} + \theta_{i}\theta_{i}'\sum_{j=1}^{r} \theta_{j}\theta_{j}'\right) \\ &\leq r + 2C_{\mathsf{G}}^{2} + C_{\mathsf{G}}^{4} \\ &\lesssim r. \end{aligned}$$

Applying Proposition 4 tells us that

$$\operatorname{var}(T_{\mathsf{LF}}^{-\mathsf{d}}) \lesssim \|P_r(f-g)\|_2^2 \left(\frac{1}{n} + \frac{1}{m}\right) + \varepsilon^{-1/s} \left(\frac{1}{n^2} + \frac{1}{nm}\right)$$
 (A.10)

Applying Chebyshev's inequality and looking at each term separately in (A.10) yields the desired bounds on n, m.

**The diagonal.** While the above test using  $T_{LF}^{-d}$  already achieves the minimax optimal sample complexity, for completeness we show that including the diagonal terms D (cf. (A.2)) doesn't degrade performance. To this end we compute

$$\mathbb{E}D = \frac{1}{n^2} \sum_{i \in [r]} \sum_{j \in [n]} \left( \phi_i^2(X_j) - \phi_i^2(Y_j) \right)$$
  
=  $\frac{1}{n} \sum_{i \in [r]} (\theta_i^2 - {\theta_i'}^2)$   
 $\leq \frac{1}{n} \|\theta + \theta'\|_2 \sqrt{\sum_{i \in [r]} (\theta_i - {\theta_i'}^2)}$   
 $\lesssim \frac{\|P_r(f - g)\|_2}{n}.$ 

We see that  $|\mathbb{E}T_{\mathsf{LF}}^{-\mathsf{d}}| \gtrsim |\mathbb{E}D|$  as soon as  $n \gtrsim 1/\varepsilon$ , which is weaker than the requirement that  $n \gtrsim n_{\mathsf{GoF}}$  and thus doesn't degrade the sample complexity. Turning to the variance, we have

$$\operatorname{var}(D) = \frac{1}{n^3} \sum_{i \in [r]} \left( \operatorname{var}(\phi_i^2(X_1)) + \operatorname{var}(\phi_i^2(Y_1)) \right)$$
$$\lesssim \frac{r}{n^3}.$$

Once again, this doesn't impose any new restrictions on n or m and thus the sample complexity is unchanged and the test  $\mathbb{1}\{T_{\mathsf{LF}} \ge 0\}$  is minimax optimal.

## D. The class $\mathcal{P}_{\mathsf{D}}$ **Proposition 8.** Let $\alpha = 1 \lor \left(\frac{k}{n} \land \frac{k}{m}\right)$ . For a constant c > 0 independent of $\varepsilon$ and k,

$$\mathcal{R}_{\mathsf{rLF}}(\varepsilon, \mathcal{P}_{\mathsf{D}}(k), \mathsf{B}_{\cdot}) \supseteq c\left\{m \ge 1/\varepsilon^2, n \ge \sqrt{k\alpha}/\varepsilon^2, mn \ge \log(k)k\alpha/\varepsilon^4\right\},\$$

where  $B_u = \{v : \|u - v\|_2 \le c_r \varepsilon / \sqrt{k}, \|v/u\|_\infty \le c'\}$  for universal constants  $c_r, c' > 0$ .

*Proof.* Choosing  $\mu$  and  $\phi$ . As for  $\mathcal{P}_{Db}$ , we take  $\mathcal{X} = [k]$ ,  $\mu = \sum_{i=1}^{k} \delta_i$ ,  $\phi_i(j) = \mathbb{1}_{\{i=j\}}$  and r = k. By the Cauchy-Schwarz inequality  $\|h\|_1 \leq \sqrt{k} \|h\|_2$  for all  $h \in \mathbb{R}^k$ .

Reducing to the small-norm case. Before applying Proposition 4 we need to 'pre-process' our distributions. For an in-depth explanation of this technique see [24], [34]. Recall that we write f, g, h for the probability mass functions of  $\mathbb{P}_X$ ,  $\mathbb{P}_Y$ ,  $\mathbb{P}_Z$  respectively, from which we observe the samples X, Y, Z of size n, n, m respectively. Recall also that the null hypothesis is that  $||f - h||_2 \leq c_r \varepsilon/\sqrt{k}$  while the alternative says that  $||g - h||_2 \leq c_r \varepsilon/\sqrt{k}$ , with  $||f - g||_2 \geq 2\varepsilon/\sqrt{k}$  guaranteed under both. In the following section we use the standard inequality  $\mathbb{P}(\lambda - x \geq \text{Poi}(\lambda)) \leq \exp(-\frac{x^2}{2(\lambda+x)})$  valid for all  $x \geq 0$  repeatedly. We also utilize the identity

$$\mathbb{E}\left[\frac{1}{\operatorname{Poi}(\lambda)+1}\right] = \begin{cases} 1 & \text{if } \lambda = 0\\ \frac{1-e^{-\lambda}}{\lambda} & \text{if } \lambda > 0, \end{cases}$$
(A.11)

which is easily verified by direct calculation. Finally, the following Lemma will come handy.

**Proposition 9.** [34, Corollary 11.6] Given t samples from an unknown discrete distribution p, there exists an algorithm that produces an estimate  $\|p\|_2^2$  with the property

$$\mathbb{P}(\widehat{\|p\|_{2}^{2}} \notin (\frac{1}{2} \|p\|_{2}^{2}, \frac{3}{2} \|p\|_{2}^{2})) \lesssim \frac{1}{\|p\|_{2}t}$$

where the implied constant is universal.

First we describe a random 'filter'  $F : \mathcal{P}_{\mathsf{D}}(k) \to \mathcal{P}_{\mathsf{D}}(K)$  that maps distributions on [k] to distributions on the inflated alphabet [K]. Let  $(n_{\mathsf{X}}, n_{\mathsf{Y}}, n_{\mathsf{Z}}) = \frac{1}{2}(n \land k, n \land k, m \land k)$  and let  $N^{\mathsf{X}} \sim \operatorname{Poi}(n_{\mathsf{X}}/2)$  independently of all other randomness, and define  $N^{\mathsf{Y}}, N^{\mathsf{Z}}$  similarly. We take the first  $N^{\mathsf{X}}, N^{\mathsf{Y}}, N^{\mathsf{Z}}$  samples from the data sets X, Y, Z respectively. In the event  $N^{\mathsf{X}} \lor N^{\mathsf{Y}} > n$  or  $N^{\mathsf{Z}} > m$  let our output to the likelihood-free hypothesis test be arbitrary, this happens with exponentially small probability. Let  $N_i^{\mathsf{X}}$  be the number of the samples  $X_1, \ldots, X_{N^{\mathsf{X}}}$  falling in bin *i*, so that  $N_i^{\mathsf{X}} \sim \operatorname{Poi}(n_{\mathsf{X}}f_i/2)$  independently for each  $i \in [k]$ , and define  $N_i^{\mathsf{Y}}, N_i^{\mathsf{Z}}$  analogously. The filter *F* is defined as follows: divide each bin  $i \in [k]$  uniformly into  $1 + N_i^{\mathsf{X}} + N_i^{\mathsf{Y}} + N_i^{\mathsf{Z}}$  bins. This filter has the following properties trivially.

- 1) The construction succeeds with probability at least  $1 3 \exp(-n \wedge m \wedge k/16)$ , we focus on this event from here on.
- 2) The construction uses at most  $n_X, n_Y, n_Z$  samples from X, Y, Z respectively and satisfies  $K \leq 5k/2$ .
- 3) For any  $u, v \in \mathcal{P}_{\mathsf{D}}(k)$  we have  $\mathsf{TV}(F(u), F(v)) = \mathsf{TV}(u, v)$  and  $||F(u) F(v)||_2 \le ||u v||_2$ .
- 4) Given a sample from an unknown  $u \in \mathcal{P}_{\mathsf{D}}(k)$  we can generate a sample from F(u) and vice-versa.

Let  $\tilde{f} \triangleq F(f)$  be the probability mass function after processing and define  $\tilde{g}, \tilde{h}$  analogously. By properties 1-2 of the filter, we may assume with probability 99% that the new alphabet's size is at most 5k/2 and that we used at most half of our samples X, Y, Z. We immediately get  $2\varepsilon \leq ||f - g||_1 = ||\tilde{f} - \tilde{g}||_1 \leq \sqrt{5k/2}||\tilde{f} - \tilde{g}||_2$  and  $||\tilde{f} - \tilde{h}||_2 \leq ||f - h||_2, ||\tilde{g} - \tilde{h}||_2 \leq ||g - h||_2$ . Adopting the convention 0/0 = 1 and using (A.11) we can bound inner products between the mass functions as

$$\mathbb{E}\left[B_{\tilde{f}\tilde{h}} + B_{\tilde{g}\tilde{h}}\right] = \mathbb{E}\left[\langle \tilde{f}\tilde{h} \rangle + \langle \tilde{g}\tilde{h} \rangle\right] \le 4\sum_{i \in [k]} \frac{f_i h_i + g_i h_i}{(n \wedge k)(f_i + g_i) + (m \wedge k)h_i} \le \frac{8}{(n \vee m) \wedge k}$$
$$\mathbb{E}\left[B_{\tilde{f}\tilde{f}} + B_{\tilde{g}\tilde{g}}\right] = \mathbb{E}\left[\|\tilde{f}\|_2^2 + \|\tilde{g}\|_2^2\right] \le 4\sum_{i \in [k]} \frac{f_i^2 + g_i^2}{(n \wedge k)(f_i + g_i) + (m \wedge k)h_i} \le \frac{8}{n \wedge k}$$
$$\mathbb{E}\|\tilde{h}\|_2^2 \le 4\sum_{i \in [k]} \frac{h_i^2}{(n \wedge k)(f_i + g_i) + (m \wedge k)h_i} \le \frac{4}{m \wedge k}.$$

By Markov's inequality we may assume that the above inequalities hold not only in expectation but with 99% probability overall with universal constants. Notice that under the null hypothesis  $\|\tilde{f} - \tilde{h}\|_2 \leq c_r \varepsilon / \sqrt{k}$  and thus  $\|\tilde{f}\|_2 \leq \|\tilde{h}\|_2 + c_r \varepsilon / \sqrt{k} \leq \|\tilde{f}\|_2 + 2c_r \varepsilon / \sqrt{k}$ , and similarly with  $\tilde{f}$  replaced by  $\tilde{g}$  under the alternative. We restrict our attention to  $c_r \leq 1$  so that  $c_r$  is treated as a constant where appropriate. Notice that  $\varepsilon / \sqrt{k} \leq 1 / \sqrt{(n \vee m) \wedge k}$  holds trivially. Thus, we obtain  $\|\tilde{f}\|_2 \vee \|\tilde{h}\|_2 \leq c / \sqrt{(m \vee n) \wedge k}$  under the null and  $\|\tilde{g}\|_2 \vee \|\tilde{h}\|_2 \leq c / \sqrt{(n \vee m) \wedge k}$  under the alternative for a universal constant c. We would like to ensure that

$$\|\tilde{f}\|_2 \vee \|\tilde{g}\|_2 \vee \|\tilde{h}\|_2 \lesssim \frac{1}{\sqrt{(m \vee n) \wedge k}}.$$
(A.12)

To this end we apply Proposition 9 using (n/4, n/4) of the remaining (transformed) X, Y samples. Let  $\|\tilde{f}\|_2^2, \|\tilde{g}\|_2^2$  denote the estimates, which lie in  $(\frac{1}{2}\|\tilde{f}\|_2^2, \frac{3}{2}\|\tilde{f}\|_2^2)$  and  $(\frac{1}{2}\|\tilde{g}\|_2^2, \frac{3}{2}\|\tilde{g}\|_2^2)$  respectively, with probability at least  $1 - \mathcal{O}((|\tilde{f}|_2^{-1} + \|\tilde{g}\|_2^{-1})/n) \ge 1 - \mathcal{O}(\sqrt{k}/n)$ , since  $\|\tilde{f}\|_2 \wedge \|\tilde{g}\|_2 \ge \sqrt{2/(5k)}$  by the Cauchy-Schwarz inequality. Assuming that  $n \ge \sqrt{k}$  this probability can be taken to be arbitrarily high, say 99%. Now we perform the following procedure: if  $\|\tilde{f}\|_2^2 > \frac{3}{2}c^2/((n \lor m) \land k)$  reject the null hypothesis, otherwise if  $\|\tilde{g}\|_2^2 > \frac{3}{2}c^2/((n \lor m) \land k)$  accept the null hypothesis, otherwise proceed with the assumption that (A.12) holds. By design this process, on our 97\% \le probability event of interest, correctly identifies the hypothesis or correctly concludes that (A.12) holds. The last step of the reduction is ensuring that the quantities  $A_{\tilde{f}\tilde{f}\tilde{h}}, A_{\tilde{q}\tilde{q}\tilde{h}}, A_{\tilde{h}\tilde{f}\tilde{q}}$  are small. The first two may be bounded

easily as

$$A_{\tilde{f}\tilde{f}\tilde{h}} + A_{\tilde{g}\tilde{g}\tilde{h}} = \langle \tilde{f}(\tilde{f} - \tilde{h})^2 \rangle + \langle \tilde{g}(\tilde{g} - \tilde{h})^2 \rangle$$

$$\leq \|\tilde{f}\|_2 \|\tilde{f} - \tilde{h}\|_4^2 + \|\tilde{g}\|_2 \|\tilde{g} - \tilde{h}\|_4^2$$

$$\lesssim \frac{\|\tilde{f} - \tilde{h}\|_2^2 + \|\tilde{g} - \tilde{h}\|_2^2}{\sqrt{(n \lor m) \land k}}$$

$$\leq \frac{\|\tilde{f} - \tilde{g}\|_2^2 + c_r^2 \varepsilon^2 / k}{\sqrt{(n \lor m) \land k}} \lesssim \frac{\|\tilde{f} - \tilde{g}\|_2^2}{\sqrt{(n \lor m) \land k}}.$$
(A.13)

To bound  $A_{\tilde{h}\tilde{f}\tilde{g}}$  we need a more sophisticated method. Recall that by definition

$$A_{\tilde{h}\tilde{f}\tilde{g}} = \sum_{i \in [k]} \frac{h_i(f_i - g_i)^2}{(1 + N_i^{\mathsf{X}} + N_i^{\mathsf{Y}} + N_i^{\mathsf{Z}})^2}.$$

Fix an  $i \in [k]$  and let  $P \triangleq N_i^{\mathsf{X}} + N_i^{\mathsf{Y}} + N_i^{\mathsf{Z}} \sim \operatorname{Poi}((n \wedge k)(f_i + g_i)/4 + (m \wedge k)h_i/4)$  and take a constant c > 0 to be specified. We have

$$\mathbb{P}\left(\frac{1}{1+P} > c\log(k)\frac{1}{\mathbb{E}P}\right) = \begin{cases} 0 & \text{if } \mathbb{E}P \le c\log(k) \\ \mathbb{P}\left(\mathbb{E}P - \left(\mathbb{E}P\left(1 - \frac{1}{c\log(k)}\right) + 1\right) > P\right) & \text{if } \mathbb{E}P > c\log(k) \end{cases}$$

Assuming that i is such that  $\mathbb{E}P \ge c \log(k)$  and taking k large enough so that  $c \log(k) \ge 2$ , we can proceed as

$$\begin{split} \mathbb{P}\left(\mathbb{E}P - \left(\mathbb{E}P\left(1 - \frac{1}{c\log(k)}\right) + 1\right) > P\right) &\leq \exp\left(-\frac{1}{2}\frac{\left(\mathbb{E}P(1 - \frac{1}{c\log(k)}) + 1\right)^2}{\mathbb{E}P(2 - \frac{1}{c\log(k)}) + 1}\right) \\ &\leq \exp\left(-\frac{1}{16}\mathbb{E}P\right) \\ &\leq \frac{1}{k^{c/16}}. \end{split}$$

Choosing c = 32, the inequality

$$A_{\tilde{h}\tilde{f}\tilde{g}} \lesssim \frac{\log(k)}{m \wedge k} \sum_{i \in [k]} \frac{(f_i - g_i)^2}{1 + N_i^{\mathsf{X}} + N_i^{\mathsf{Y}} + N_i^{\mathsf{Z}}} \asymp \frac{\log(k)}{m \wedge k} \|\tilde{f} - \tilde{g}\|_2^2$$

holds with probability at least 1 - 1/k. Using that  $\|h/f\|_{\infty} \wedge \|h/g\|_{\infty} \lesssim 1$  for both (LF) and (rLF), we obtain  $A_{\tilde{h}\tilde{f}\tilde{g}} \lesssim \frac{\log(k)}{n \wedge k} \|\tilde{f} - \tilde{g}\|_2^2$  similarly. Combining the two bounds yields

$$A_{\tilde{h}\tilde{f}\tilde{g}} \lesssim \frac{\log(k)}{(m \lor n) \land k} \|\tilde{f} - \tilde{g}\|_2^2.$$
(A.14)

To summarize, under the assumptions that  $n \gtrsim \sqrt{k}$ , and at the cost of inflating the alphabet size to at most  $\frac{5}{2}k$  and a probability of error at most  $3\% + \frac{1}{k}$ , we may assume that the inequalities (A.12), (A.13) and (A.14) hold with universal constants.

Applying Proposition 4. We only analyse the type-I error, as the type-II error follows analogously. As explained earlier, we now apply the test  $\mathbb{1}\{T_{\mathsf{LF}}^{-\mathsf{d}} \ge 0\}$  to the transformed samples with pmfs  $\tilde{f}, \tilde{g}, \tilde{h}$ . Note that for  $c_{\mathsf{r}} < \sqrt{2/5}$  we have

$$\|\tilde{g} - \tilde{h}\|_2^2 - \|\tilde{f} - \tilde{h}\|_2^2 \gtrsim \|\tilde{f} - \tilde{g}\|_2^2$$

for universal constants. Therefore we see that  $-\mathbb{E}T_{\mathsf{LF}}^{-\mathsf{d}} \ge c \|\tilde{f} - \tilde{g}\|_2^2 + R$  for some universal constant c > 0, where the residual term R can be bounded as

$$\begin{aligned} |R| &= \left| \frac{\|\tilde{f}\|_2^2 - \|\tilde{g}\|_2^2}{n} \right| \\ &\lesssim \frac{\|\tilde{f} - \tilde{g}\|_2}{n\sqrt{k \wedge (m \vee n)}}, \end{aligned}$$

where we used (A.12). Let  $\alpha = (\frac{k}{n} \wedge \frac{k}{m}) \vee 1$ . We have  $-\mathbb{E}T_{\mathsf{LF}}^{-\mathsf{d}} \gtrsim \|\tilde{f} - \tilde{g}\|_2^2$  provided  $n \gtrsim 1/(\|\tilde{f} - \tilde{g}\|_2 \sqrt{k \wedge (m \vee n)}) \approx \sqrt{\alpha}/\varepsilon$ , which we assume from here on. Plugging in the bounds derived above, the test  $\mathbb{1}(T_{\mathsf{LF}} \ge 0)$  on the transformed observations has type-I probability of error bounded by 1/3 provided

$$\|\tilde{f} - \tilde{g}\|_{2}^{4} \gtrsim \frac{1}{n} \frac{\sqrt{\alpha}}{\sqrt{k}} \|\tilde{f} - \tilde{g}\|_{2}^{2} + \frac{1}{m} \frac{\log(k)\alpha}{k} \|\tilde{f} - \tilde{g}\|_{2}^{2} + \frac{\alpha}{k} \left(\frac{1}{nm} + \frac{1}{n^{2}}\right)$$

for a small enough implied constant on the left. Looking at each term separately yields the sufficient conditions

$$m \gtrsim \frac{\log(k)\alpha}{\varepsilon^2}$$
 and  $n \gtrsim \frac{\sqrt{k\alpha}}{\varepsilon^2}$  and  $mn \gtrsim \frac{k\alpha}{\varepsilon^4}$ .

The diagonal. See the discussion at the end of the proof for  $\mathcal{P}_{Db}$ .

## APPENDIX B

#### LOWER BOUNDS OF THEOREM 1 AND 2

Let  $\mathcal{M}(\mathcal{X})$  be the set of all probability measures on some space  $\mathcal{X}$ , and  $\mathcal{P} \subseteq \mathcal{M}(\mathcal{X})$  be some family of distributions. In this section we prove lower bounds for likelihood-free hypothesis testing problems. For clarity, let us formally state the problem as testing between the hypotheses

$$H_{0} = \{ \mathbb{P}_{\mathsf{X}}^{\otimes n} \otimes \mathbb{P}_{\mathsf{Y}}^{\otimes n} \otimes \mathbb{P}_{\mathsf{X}}^{\otimes m} : \mathbb{P}_{\mathsf{X}}, \mathbb{P}_{\mathsf{Y}} \in \mathcal{P}, \, \mathsf{TV}(\mathbb{P}_{\mathsf{X}}, \mathbb{P}_{\mathsf{Y}}) \ge \varepsilon \}$$
  
versus  
$$H_{1} = \{ \mathbb{P}_{\mathsf{X}}^{\otimes n} \otimes \mathbb{P}_{\mathsf{Y}}^{\otimes n} \otimes P_{\mathsf{Y}}^{\otimes m} : \mathbb{P}_{\mathsf{X}}, \mathbb{P}_{\mathsf{Y}} \in \mathcal{P}, \, \mathsf{TV}(\mathbb{P}_{\mathsf{X}}, \mathbb{P}_{\mathsf{Y}}) \ge \varepsilon \}.$$
(B.1)

Our strategy for proving lower bounds relies on the following well known result proved in the main text.

**Lemma 5.** Take hypotheses  $H_0, H_1 \subseteq \mathcal{M}(\mathcal{X})$  and  $P_0, P_1 \in \mathcal{M}(\mathcal{X})$  random. Then

$$\inf_{\psi} \max_{i=0,1} \sup_{P \in H_i} P(\psi \neq i) \geq \frac{1}{2} \left( 1 - \mathsf{TV}(\mathbb{E}P_0, \mathbb{E}P_1) \right) - \sum_i \mathbb{P}(P_i \notin H_i),$$

where the infimum is over all tests  $\psi : \mathcal{X} \to \{0, 1\}$ .

The following will also be used multiple times throughout:

**Lemma 6** ([48, Lemmas 2.3 and 2.4]). For any probability measures  $\mathbb{P}_0, \mathbb{P}_1$ ,

$$\frac{1}{4}\mathsf{H}^4(\mathbb{P}_0,\mathbb{P}_1) \le \mathsf{TV}^2(\mathbb{P}_0,\mathbb{P}_1) \le \mathsf{H}^2(\mathbb{P}_0,\mathbb{P}_1) \le \mathsf{KL}(\mathbb{P}_0\|\mathbb{P}_1) \le \chi^2(\mathbb{P}_0\|\mathbb{P}_1).$$

The inequalities between TV and H are attributed to Le Cam, while the bound TV  $\leq \sqrt{KL/2}$  is due to Pinsker. The use of the  $\chi^2$ -divergence for bounding the total variation distance between mixtures of products was pioneered by Ingster [65], and is sometimes referred to as the *Ingster-trick*.

Recall that the necessity of  $m \gtrsim n_{\text{HT}}(\varepsilon, \mathcal{P})$  and  $n \gtrsim n_{\text{GoF}}(\varepsilon, \mathcal{P})$  were shown in Proposition 1. Thus, most of work lies in obtaining the lower bound on  $m \cdot n$ .

A. The class  $\mathcal{P}_{H}$ 

**Proposition 10.** For a constant c > 0 independent of  $\varepsilon$ ,

$$c\{m \ge 1/\varepsilon^2, n \ge \varepsilon^{-(2\beta+d/2)/\beta}, mn \ge \varepsilon^{-2(2\beta+d/2)/\beta}\} \supseteq \mathcal{R}_{\mathsf{LF}}(\varepsilon, \mathcal{P}_{\mathsf{H}}(\beta, d, C_{\mathsf{H}}))$$

*Proof.* Adversarial construction. Take a smooth function  $h : \mathbb{R}^d \to \mathbb{R}$  supported on of  $[0,1]^d$  with  $\int h = 0$  and  $\int h^2 = 1$ . Let  $\kappa \ge 1$  be an integer, and for  $j \in [\kappa]^d$  define the scaled and translated functions  $h_j$  as

$$h_j(x) = \kappa^{d/2} h(\kappa x - j + 1).$$

Then  $h_j$  is supported on the cube  $[(j-1)/\kappa, j/\kappa]$  and  $\int h_j^2 = 1$ , where we write  $j/\kappa = (j_1/\kappa, \dots, j_d/\kappa)$ . Let  $\rho > 0$  be small and for each  $\eta \in \{-1, 0, 1\}^{\kappa^d}$  define the function

$$f_{\eta}(x) = 1 + \rho \sum_{j \in [\kappa]^d} \eta_j h_j(x).$$

In particular,  $f_0 = 1$  is the uniform density. Clearly  $\int f_{\eta} = 1$ , and to make it positive we choose  $\rho, \kappa$  such that  $\rho \kappa^{d/2} \|h\|_{\infty} \leq 1/2$ . By [42], choosing

$$\rho \kappa^{d/2+\beta} \le C_{\mathsf{H}}/(4\|h\|_{\mathcal{C}^{\lfloor\beta\rfloor}} \vee 2\|h\|_{\mathcal{C}^{\lfloor\beta\rfloor+1}}) \tag{B.2}$$

ensures that  $f_{\eta} \in \mathcal{P}(\beta, d, C_{\mathsf{H}})$ . Note also that  $||f_{\eta} - 1||_1 = \rho \kappa^{d/2}$ . For  $\varepsilon \in (0, 1)$  we set  $\kappa \simeq \varepsilon^{-1/\beta}$  and  $\rho \simeq \varepsilon^{(2\beta+d)/(2\beta)}$ . These ensure that (B.2) and  $\mathsf{TV}(f_{\eta}, f_0) \gtrsim \varepsilon$  hold, where as usual the constants may depend on  $(\beta, d, C_{\mathsf{H}})$ . Noting that  $||\sqrt{f_{\eta}} - 1||_2 \simeq ||f_{\eta} - 1||_1 \gtrsim \varepsilon$ , we immediately obtain the lower bound  $m \gtrsim 1/\varepsilon^2$  by reduction from binary hypothesis testing (III.1). Observe also that for any  $\eta, \eta'$ ,

$$\int_{[0,1]^d} f_{\eta}(x) f_{\eta'}(x) \mathrm{d}x = 1 + \rho^2 \langle \eta, \eta' \rangle$$
(B.3)

which will be used later.

**Goodness-of-fit testing.** Let  $\eta$  be drawn uniformly at random. We show that  $\mathsf{TV}(f_0^{\otimes n}, \mathbb{E}f_{\eta}^{\otimes n})$  can be made arbitrarily small provided  $n \geq \varepsilon^{-(2\beta+d/2)/\beta}$ , which yields a lower bound on n via reduction from goodness-of-fit testing (III.3). By Lemma 6 we can focus on bounding the  $\chi^2$  divergence. Via Ingster's trick we have

$$\chi^{2}(\mathbb{E}_{\eta}[f_{\eta}^{\otimes n}], f_{0}^{\otimes n}) + 1 = \int_{[0,1]^{d} \times \dots \times [0,1]^{d}} \left( \mathbb{E}_{\eta} \prod_{i=1}^{n} f_{\eta}(x_{i}) \right)^{2} \mathrm{d}x_{1} \cdots \mathrm{d}x_{n}$$
$$= \mathbb{E}_{\eta\eta'} \prod_{i=1}^{n} \left( \int_{[0,1]^{d}} f_{\eta}(x) f_{\eta'}(x) \mathrm{d}x \right),$$

where  $\eta, \eta'$  are i.i.d.. By (B.3) and the inequalities  $1 + x \le e^x$ ,  $\cosh(x) \le \exp(x^2)$  for all  $x \in \mathbb{R}$ , we have

$$= \mathbb{E}_{\eta\eta'} \left( 1 + \rho^2 \langle \eta, \eta' \rangle \right)^n$$
  

$$\leq \mathbb{E}_{\eta\eta'} \exp(n\rho^2 \langle \eta, \eta' \rangle)$$
  

$$= \cosh(n\rho^2)^{\kappa^d}$$
  

$$\leq \exp(n^2 \rho^4 \kappa^d).$$

Thus, goodness-of-fit testing is impossible unless  $n \gtrsim \rho^{-2} \kappa^{-d/2} \asymp 1/\varepsilon^{(2\beta+d/2)/\beta}$ .

**Likelihood-free hypothesis testing.** We are now ready to show the lower bound on the interaction term mn. Once again  $\eta \in \{\pm 1\}^{\kappa^d}$  is drawn at random and we apply Lemma 5 with the choices  $P_0 = f_{\eta}^{\otimes n} \otimes f_0^{\otimes n} \otimes f_{\eta}^{\otimes m}$  against  $P_1 = f_{\eta}^{\otimes n} \otimes f_0^{\otimes n+m}$ . Let  $\mathbb{P}_{0,XYZ}$ ,  $\mathbb{P}_{1,XYZ}$  denote the joint distribution of the samples X, Y, Z under the measures  $\mathbb{E}P_0$ ,  $\mathbb{E}P_1$  respectively. By Pinsker's inequality and the chain rule we have

$$\begin{aligned} \mathsf{TV}(\mathbb{P}_{0,XYZ},\mathbb{P}_{1,XYZ})^2 &= \mathsf{TV}(\mathbb{P}_{0,XZ},\mathbb{P}_{1,XZ})^2 \\ &\leq \mathsf{KL}(\mathbb{P}_{0,XZ} \| \mathbb{P}_{1,XZ}) \\ &= \mathsf{KL}(\mathbb{P}_{0,Z|X} \| \mathbb{P}_{1,Z|X} | \mathbb{P}_{0,X}) + \underbrace{\mathsf{KL}(\mathbb{P}_{0,X} \| \mathbb{P}_{1,X})}_{=0}, \end{aligned}$$

where the last line uses that the marginal of X is equal under both measures. Clearly  $P_{1,Z|X}$  is simply  $\text{Unif}([0,1]^d)^{\otimes m}$ and  $P_{0,X}, P_{0,Z|X}$  have densities  $\mathbb{E}_{\eta} f_{\eta}^{\otimes n}$  and  $\mathbb{E}_{\eta|X} f_{\eta}^{\otimes m}$  respectively. Given X, let  $\eta'$  be an independent copy of  $\eta$ from the posterior. By Ingster's trick we have

$$\begin{aligned} \mathsf{KL}(\mathbb{P}_{0,Z|X} \| \mathbb{P}_{1,Z|X} | \mathbb{P}_{0,X}) &\leq \chi^2(\mathbb{P}_{0,Z|X} \| \mathbb{P}_{1,Z|X} | \mathbb{P}_{0,X}) \\ &= -1 + \mathbb{E}_X \int_{[0,1]^d \times \dots \times [0,1]^d} \mathbb{E}_{\eta|X} \mathbb{E}_{\eta'|X} \prod_{i=1}^m f_\eta(z_i) f_{\eta'}(z_i) \mathrm{d} z_1 \dots \mathrm{d} z_m \\ &= -1 + \mathbb{E}_{\eta\eta'} (1 + \rho^2 \langle \eta, \eta' \rangle)^m, \end{aligned}$$

where the last line uses (B.3). Let  $N = (N_1, \ldots, N_{\kappa^d})$  be the vector of counts indicating the number of  $X_i$  that fall into each bin  $\{[(j-1)/\kappa, j/\kappa]\}_{j \in [\kappa]^d}$ . Clearly  $N \stackrel{d}{\sim} \text{Multinomial}(n, (\frac{1}{\kappa^d}, \ldots, \frac{1}{\kappa^d}))$ . Using that  $\eta_j \eta'_j$  depends on only those  $X_i$  that fall in bin j and the inequality  $1 + x \leq \exp(x)$  valid for all  $x \in \mathbb{R}$ , we can write

$$\chi^{2}(\mathbb{P}_{0,Z|X}||\mathbb{P}_{1,Z|X}||\mathbb{P}_{0,X}) + 1 \leq \mathbb{E}_{N}\mathbb{E}_{\eta\eta'|N}\prod_{j\in[\kappa]^{d}}\exp(\rho^{2}m\eta_{j}\eta'_{j})$$
$$= \mathbb{E}_{N}\prod_{j\in[\kappa]^{d}}\mathbb{E}_{\eta_{j}\eta'_{j}|N_{j}}\exp(\rho^{2}m\eta_{j}\eta'_{j}).$$

We now focus on a particular bin j. Define the bin-conditional densities

$$p_{\pm} = \kappa^d (1 \pm \rho h_j) \mathbb{1}_{[(j-1)/\kappa, j/\kappa]},\tag{B.4}$$

where we drop the dependence on j in the notation. Let  $X^{(j)} \triangleq (X_{i_1}, \ldots, X_{i_{N_j}})$  be those  $X_i$  that fall in bin j. Note that  $\{i_1, \ldots, i_{N_j}\}$  is a uniformly distributed size  $N_j$  subset of [n] and given  $N_j$ , the density of  $X_{i_1}, \ldots, X_{i_{N_j}}$  is  $\frac{1}{2}(p_+^{\otimes N_j} + p_-^{\otimes N_j})$ . We can calculate

$$\begin{split} \mathbb{P}(\eta_{j}\eta_{j}' = 1|N_{j}) &= \mathbb{E}_{X^{(j)}|N_{j}} \mathbb{P}(\eta_{j}\eta_{j}' = 1|X^{(j)}) \\ &= \mathbb{E}_{X^{(j)}|N_{j}} \left[ \mathbb{P}(\eta_{j} = 1|X^{(j)})^{2} + \mathbb{P}(\eta_{j} = -1|X^{(j)})^{2} \right] \\ &= \mathbb{E}_{X^{(j)}|N_{j}} \left[ \frac{\frac{1}{4}(p_{+}^{\otimes N_{j}})^{2} + \frac{1}{4}(p_{-}^{\otimes N_{j}})^{2}}{\frac{1}{4}(p_{+}^{\otimes N_{j}} + p_{-}^{\otimes N_{j}})^{2}} \right] \\ &= \frac{1}{2} + \frac{1}{4} \left( \chi^{2}(p_{+}^{\otimes N_{j}} \| \frac{1}{2}(p_{+}^{\otimes N_{j}} + p_{-}^{\otimes N_{j}})) + \chi^{2}(p_{-}^{\otimes N_{j}} \| \frac{1}{2}(p_{+}^{\otimes N_{j}} + p_{-}^{\otimes N_{j}})) \right). \end{split}$$

By convexity of the  $\chi^2$  divergence in its arguments and tensorization, we have

$$\mathbb{P}(\eta_{j}\eta_{j}'=1|N_{j}) \leq \frac{1}{2} + \frac{1}{8} \left( \chi^{2}(p_{+}^{\otimes N_{j}} \| p_{-}^{\otimes N_{j}}) + \chi^{2}(p_{-}^{\otimes N_{j}} \| p_{+}^{\otimes N_{j}}) \right)$$
$$= \frac{1}{4} + \sum_{\omega \in \{\pm 1\}} \left( \kappa^{d} \int_{[(j-1)/\kappa, j/\kappa]} \frac{(1+\omega\rho h_{j}(x))^{2}}{1-\omega\rho h_{j}(x)} \mathrm{d}x \right)^{N_{j}}.$$

Using that  $\rho \|h_j\|_{\infty} \leq 1/2$  by construction, we have

$$\int_{[(j-1)/\kappa,j/\kappa]} \frac{(1+\rho h_j(x))^2}{1-\rho h_j(x)} \mathrm{d}x = \frac{1}{\kappa^d} + \int_{[(j-1)/\kappa,j/\kappa]} \frac{4\rho^2 h_j^2(x)}{1-\rho h_j(x)} \mathrm{d}x$$
$$\leq \frac{1}{\kappa^d} + 8\rho^2.$$

The same bound is obtained for the other integral term. We get

$$\chi^{2}(\mathbb{P}_{0,Z|X}\|\mathbb{P}_{1,Z|X}\|\mathbb{P}_{0,X}) + 1 \leq \mathbb{E}_{N} \prod_{j \in [\kappa]^{d}} \left(\frac{1}{4} \left(e^{\rho^{2}m} - e^{-\rho^{2}m}\right) \left(1 + (1 + 8\rho^{2}\kappa^{d})^{N_{j}}\right) + e^{-\rho^{2}m}\right) = (\dagger).$$

The final step is to apply Lemma 8 to pass the expectation through the product. Assuming that  $m \vee n \lesssim \rho^{-2} \approx \epsilon^{-(2\beta+d)/\beta}$  for a small enough implied constant, using the inequalities  $e^x \leq 1 + x + x^2$ ,  $1 - x \leq e^{-x} \leq 1 - x + x^2/2$  valid for all  $x \in [0, 1]$ , and Lemma 8, we obtain

$$\begin{aligned} (\dagger) &\leq (e^{-\rho^2 m} + \frac{1}{4} \left( e^{\rho^2 m} - e^{-\rho^2 m} \right) (1 + e^{8\rho^2 n}))^{\kappa^d} \\ &\leq (1 + c\rho^4 m n)^{\kappa^d} \\ &\leq \exp(c\rho^4 \kappa^d m n) \end{aligned}$$

for a universal constant c > 0. Therefore, if  $m \lor n \lesssim \varepsilon^{-(2\beta+d)/\beta}$  likelihood-free hypothesis testing is impossible unless  $mn \gtrsim \rho^{-4} \kappa^{-d} \asymp 1/\varepsilon^{2(2\beta+d/2)/\beta}$ . Combining with the previously derived bounds  $m \gtrsim 1/\varepsilon^2$  and  $n \gtrsim 1/\varepsilon^{(2\beta+d/2)/\beta}$ , we can conclude.

#### B. The class $\mathcal{P}_{G}$

**Proposition 11.** For a constant c > 0 independent of  $\varepsilon$ ,

$$c\{m \ge 1/\varepsilon^2, n \ge \varepsilon^{-(2s+1/2)/s}, mn \ge \varepsilon^{-2(2s+1/2)/s}\} \supseteq \mathcal{R}_{\mathsf{LF}}(\varepsilon, \mathcal{P}_{\mathsf{G}}(s, C_{\mathsf{G}}))$$

*Proof.* Adversarial construction. Let  $\gamma \in \ell^1$  and  $\theta \sim \bigotimes_{k=1}^{\infty} \mathcal{N}(0, \gamma_k)$ . Define the random measure  $\mathbb{P}_{\gamma} \triangleq \bigotimes_{k=1}^{\infty} \mathcal{N}(\theta_k, 1)$ . Recall our definition of the Sobolev ellipsoid  $\mathcal{E}(s, C_{\mathsf{G}})$  with associated sobolev norm  $\|\cdot\|_s$ . We have

$$(\mathbb{E}\|\theta\|_s)^2 \le \mathbb{E}\sum_{j=1}^{\infty} j^{2s} \theta_i^2$$
$$= \|\sqrt{\gamma}\|_s^2.$$

Let  $\varepsilon \in (0,1)$  be given. For our proofs we use

$$\gamma_k = \begin{cases} c_1 \varepsilon^{(2s+1)/s} & \text{for } 1 \le k \le c_2 \varepsilon^{-1/s} \\ 0 & \text{otherwise} \end{cases}$$
(B.5)

for appropriate constants  $c_1, c_2$ . We need to verify that this choice is valid, in that  $\mathbb{P}_{\gamma} \in \mathcal{P}_{\mathsf{G}}(s, C_{\mathsf{G}})$  and  $\mathsf{TV}(\mathbb{P}_{\gamma}, \mathbb{P}_0) \gtrsim \varepsilon$  with high probability. To this end, we compute

$$\|\sqrt{\gamma}\|_{s}^{2} = c_{1}\varepsilon^{(2s+1)/s} \sum_{j=1}^{c_{2}\varepsilon^{-1/s}} j^{2s} \le c_{1}c_{2}^{2s+1}$$
$$\mathsf{TV}(\mathbb{P}_{\gamma}, \mathbb{P}_{0}) \ge \frac{1 \wedge \|\theta\|_{2}}{200},$$

where the second line follows by [54, Theorem 1.2]. By standard results, the squared norm  $\|\theta\|_2^2$  concentrates around  $c_1c_2\varepsilon^2$  with exponentially high probability. Further, for sufficiently large  $c_1, c_2$  the event  $\{\mathbb{P}_{\gamma} \notin \mathcal{P}_{\mathsf{G}}(s, C_{\mathsf{G}})\}$ has probability at most, say, 0.01 by Markov's inequality. Thus  $c_1, c_2$  can be chosen independently of  $\varepsilon$  so that  $\mathbb{P}(\mathbb{P}_{\gamma} \in \mathcal{P}_{\mathsf{G}}(s, C_{\mathsf{G}}), \mathsf{TV}(\mathbb{P}_{\gamma}, \mathbb{P}_{0}) \geq \varepsilon) \geq .98. \text{ By Lemma 2 we have } \mathsf{H}(\mathbb{P}_{\gamma}, \mathbb{P}_{0}) \asymp \varepsilon \text{ with high probability, and thus } \mathbb{P}(\mathbb{P}_{\gamma}, \mathbb{P}_{0}) \geq \varepsilon \geq .98. \mathbb{P}(\mathbb{P}_{\gamma}, \mathbb{P}_{0}) \geq .98. \mathbb{P}(\mathbb{P}_{\gamma}, \mathbb{P}_{0}) \geq \varepsilon \geq .98. \mathbb{P}(\mathbb{P}_{\gamma}, \mathbb{P}_{0}) \geq \varepsilon \geq .98. \mathbb{P}(\mathbb{P}_{\gamma}, \mathbb{P}(\mathbb{P}_{\gamma}, \mathbb{P}_{0}) \geq \varepsilon \geq .98. \mathbb{P}(\mathbb{P}_{\gamma}, \mathbb{P}(\mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}(\mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}(\mathbb{P}_{1}, \mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}(\mathbb{P}_{1}, \mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}(\mathbb{P}_{1}, \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}_{1}) \geq .98. \mathbb{P}(\mathbb{P}(\mathbb{P}_{1}, \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}(\mathbb{P}_{1}, \mathbb{P}_$ the bound  $m \gtrsim 1/\varepsilon^2$  follows by reduction from hypothesis testing (III.1).

**Goodness-of-fit testing.** We show that  $\mathsf{TV}(\mathbb{P}_0^{\otimes n}, \mathbb{E}\mathbb{P}_{\gamma}^{\otimes n})$  can be made arbitrarily small as long as  $n \gtrsim 1/\varepsilon^{(2s+1/2)/s}$ , which yields a lower bound on n via reduction from goodness-of-fit testing (III.3). Let us compute the distribution  $\mathbb{EP}_{\gamma}^{\otimes n}$ . By independence clearly  $\mathbb{EP}_{\gamma}^{\otimes n} = \bigotimes_{k=1}^{\infty} \mathbb{E}_{\theta \sim \mathcal{N}(0,\gamma_k)} \mathcal{N}(\theta, 1)^{\otimes n}$ . Focusing on the inner term and and dropping the subscript k, for the density we have

$$\mathbb{E}_{\theta \sim \mathcal{N}(0,\gamma)} \left[ \frac{1}{(2\pi)^{n/2}} \exp(-\frac{1}{2} \sum_{j=1}^{n} (x_j - \theta)^2) \right] \propto \exp(-\frac{\|x\|_2^2}{2}) \mathbb{E} \exp(-\frac{n}{2} (\theta^2 - 2\theta \bar{x})),$$

where we write  $\bar{x} \triangleq \frac{1}{n} \sum_{j} x_{j}$ . Looking at just the term involving  $\theta$ , we have

$$\mathbb{E}\exp(-\frac{n}{2}(\theta^2 - 2\theta\bar{x})) \propto \int \exp(-\frac{1}{2}(\theta^2(n+\frac{1}{\gamma}) - 2\theta n\bar{x})) \mathrm{d}\theta \propto \exp(\frac{1}{2}\frac{n^2\bar{x}^2}{n+\frac{1}{\gamma}}).$$

Putting everything together, we see that  $\mathbb{EP}_{\gamma}^{\otimes n} = \bigotimes_{k=1}^{\infty} \mathcal{N}(0, (\mathsf{Id}_n - \frac{\gamma_k}{1+n\gamma_k} \mathbb{1}_n \mathbb{1}_n^{\mathsf{T}})^{-1})$ . Thus, using Lemma 6 we obtain

$$\begin{split} \mathsf{TV}^2(\mathbb{P}_0^{\otimes n}, \mathbb{E}\mathbb{P}_{\gamma}^{\otimes n}) &\leq \sum_{k=1}^{\infty} \mathsf{KL}(\mathcal{N}(0, \mathsf{Id}_n) \| \mathcal{N}(0, (\mathsf{Id}_n - \frac{\gamma_k}{1 + n\gamma_k} \mathbbm{1}_n \mathbbm{1}_n^{\mathsf{T}})^{-1})) \\ &= \frac{1}{2} \sum_{k=1}^{\infty} \left( -\frac{n\gamma_k}{n\gamma_k + 1} + \log(1 + n\gamma_k) \right) \\ &\leq \frac{1}{2} \sum_{k=1}^{\infty} \frac{n^2 \gamma_k^2}{1 + n\gamma_k} \lesssim \sum_{k=1}^{\infty} n^2 \gamma_k^2. \end{split}$$

Taking  $\gamma$  as in (B.5) gives

$$\mathsf{TV}^2(\mathbb{P}_0^{\otimes n}, \mathbb{EP}_\gamma^{\otimes n}) \lesssim n^2 \varepsilon^{2(2s+1/2)/s}.$$

Thus, goodness-of-fit testing is impossible unless  $n \gtrsim 1/\varepsilon^{(2s+1/2)/s}$  as desired. **Likelihood-free hypothesis testing.** We apply Lemma 5 with measures  $P_0 = \mathbb{P}_{\gamma}^{\otimes n} \otimes \mathbb{P}_0^{\otimes n} \otimes \mathbb{P}_{\gamma}^{\otimes m}$  and  $P_1 = \mathbb{P}_{\gamma}^{\otimes n} \otimes \mathbb{P}_0^{\otimes n} \otimes \mathbb{P}_0^{\otimes m}$ . By an analogous calculation to that in the previous part, we obtain

$$\mathbb{E}P_{0} = \otimes_{k=1}^{\infty} \mathcal{N}(0, (\mathsf{Id}_{2n+m} - \frac{1}{n+m+\frac{1}{\gamma_{k}}} \begin{pmatrix} \mathbb{1}_{n} \mathbb{1}_{n}^{\mathsf{T}} & 0 & \mathbb{1}_{n} \mathbb{1}_{m}^{\mathsf{T}} \\ 0 & 0 & 0 \\ \mathbb{1}_{m} \mathbb{1}_{n}^{\mathsf{T}} & 0 & \mathbb{1}_{m} \mathbb{1}_{m}^{\mathsf{T}} \end{pmatrix})^{-1}) \triangleq \otimes_{k=1}^{\infty} \mathcal{N}(0, \Sigma_{0k})$$
$$\mathbb{E}P_{1} = \otimes_{k=1}^{\infty} \mathcal{N}(0, (\mathsf{Id}_{2n+m} - \frac{1}{n+\frac{1}{\gamma_{k}}} \begin{pmatrix} \mathbb{1}_{n} \mathbb{1}_{n}^{\mathsf{T}} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}})^{-1}) \triangleq \otimes_{k=1}^{\infty} \mathcal{N}(0, \Sigma_{1k}).$$

By the Sherman-Morrison formula, we have

$$\Sigma_{0k} = \mathsf{Id}_{2n+m} + \gamma_k \begin{pmatrix} \mathbbm{1}_n \mathbbm{1}_n^\mathsf{T} & 0 & \mathbbm{1}_n \mathbbm{1}_n^\mathsf{T} \\ 0 & 0 & 0 \\ \mathbbm{1}_m \mathbbm{1}_n^\mathsf{T} & 0 & \mathbbm{1}_m \mathbbm{1}_n^\mathsf{T} \end{pmatrix}$$

Therefore, by Pinsker's inequality and the closed form expression for the KL-divergence between centered Gaussians, we obtain

$$\mathsf{TV}^{2}(\mathbb{E}P_{0},\mathbb{E}P_{1}) \leq \mathsf{KL}(\mathbb{E}P_{0}||\mathbb{E}P_{1})$$
$$= \frac{1}{2}\sum_{k=1}^{\infty} \left(\gamma_{k}m - \log\left(1 + \frac{\gamma_{k}m}{\gamma_{k}(n+m) + 1}\right)\right).$$

$$\mathsf{TV}^2(\mathbb{E}P_0,\mathbb{E}P_1) \lesssim \varepsilon^{-2(2s+1/2)/s}(m^2+mn).$$

Therefore, likelihood-free hypothesis testing is impossible unless  $m \gtrsim \varepsilon^{-(2s+1/2)/s}$  or  $nm \gtrsim \varepsilon^{-2(2s+1/2)/s}$ . Note that we already have the lower bound  $n \gtrsim \varepsilon^{-(2s+1/2)/s}$  by reduction from goodness-of-fit testing (III.3), so that  $m \gtrsim \varepsilon^{-(2s+1/2)/s}$  automatically implies  $nm \gtrsim \varepsilon^{-2(2s+1/2)/s}$ . Combining everything we get the desired bounds.  $\Box$ 

C. The classes  $\mathcal{P}_{\mathsf{Db}}$  and  $\mathcal{P}_{\mathsf{D}}$ 

**Proposition 12.** For a constant c > 0 independent of  $\varepsilon$  and k,

$$c\{m \ge 1/\varepsilon^2, n \ge \sqrt{k}/\varepsilon^2, mn \ge k/\varepsilon^4\} \supseteq \mathcal{R}_{\mathsf{LF}}(\varepsilon, \mathcal{P}_{\mathsf{Db}}) \supseteq \mathcal{R}_{\mathsf{LF}}(\varepsilon, \mathcal{P}_{\mathsf{D}}).$$

*Proof.* The second inclusion is trivial. For the first inclusion we proceed analogously to the case of  $\mathcal{P}_{\mathsf{H}}$ . Adversarial construction. Let k be an integer and  $\varepsilon \in (0, 1)$ . For  $\eta \in \{-1, 1\}^k$  define the distribution  $p_{\eta}$  on [2k] by

$$p_{\eta}(2j-1) = \frac{1}{2k}(1+\eta_{j}\varepsilon)$$
$$p_{\eta}(2j) = \frac{1}{2k}(1-\eta_{j}\varepsilon),$$

for  $j \in [k]$ . Clearly  $\mathsf{H}(p_{\eta}, p_0) \asymp \mathsf{TV}(p_{\eta}, p_0) = \varepsilon$ , where  $p_0 = \mathrm{Unif}[2k]$ , so that by reduction from binary hypothesis testing (III.3) we get the lower bound  $m \gtrsim 1/\varepsilon^2$ . Observe also that for any  $\eta, \eta' \in \{\pm 1\}^k$ ,

$$\sum_{j\in[2k]} p_{\eta}(j)p_{\eta'}(j) = \frac{1}{2k} \left(1 + \frac{\varepsilon^2 \langle \eta, \eta' \rangle}{k}\right).$$
(B.6)

**Goodness-of-fit testing.** Let  $\eta$  be uniformly random. We show that  $\mathsf{TV}(p_0^{\otimes n}, \mathbb{E}p_{\eta}^{\otimes n})$  can be made arbitrarily small as long as  $n \leq \sqrt{k}/\varepsilon^2$ , which yields the corresponding lower bound on n by reduction from goodness-of-fit testing (III.3). Once again, by Lemma 6 we focus on the  $\chi^2$  divergence. We have

$$\chi^{2}(\mathbb{E}p_{\eta}^{\otimes n} \| p_{0}^{\otimes n}) + 1 = (2k)^{n} \sum_{j \in [2k]^{n}} \mathbb{E}_{\eta\eta'} \prod_{i=1}^{n} p_{\eta}(j_{i}) p_{\eta'}(j_{i})$$
$$= \mathbb{E}_{\eta\eta'} (1 + \frac{\varepsilon^{2} \langle \eta, \eta' \rangle}{k})^{n}$$
$$\leq \exp(n^{2} \varepsilon^{4} / k)$$

where the penultimate line follows from (B.6) and the last line via the same argument as in B-A. Thus, goodness-of-fit testing is impossible unless  $n \gtrsim \sqrt{k}/\varepsilon^2$ . Likelihood-free hypothesis testing. We apply Lemma 5 with the two random measures  $P_0 = p_{\eta}^{\otimes n} \otimes p_0^{\otimes n} \otimes p_{\eta}^{\otimes m}$ 

Likelihood-free hypothesis testing. We apply Lemma 5 with the two random measures  $P_0 = p_{\eta}^{\otimes n} \otimes p_0^{\otimes m} \otimes p_{\eta}^{\otimes m}$ and  $P_1 = p_{\eta}^{\otimes n} \otimes p_0^{\otimes (n+m)}$ . Analogously to the case of  $\mathcal{P}_{\mathsf{H}}$ , let  $\mathbb{P}_{0,XYZ}$ ,  $\mathbb{P}_{1,XYZ}$  respectively denote the distribution of the observations X, Y, Z under  $\mathbb{E}P_0$ ,  $\mathbb{E}P_1$  respectively. As for  $\mathcal{P}_{\mathsf{H}}$ , we have

$$\begin{aligned} \mathsf{TV}^2(\mathbb{P}_{0,XYZ},\mathbb{P}_{1,XYZ}) &\leq \mathsf{KL}(\mathbb{P}_{0,XYZ} \| \mathbb{P}_{1,XYZ}) \\ &\leq \mathsf{KL}(\mathbb{P}_{0,Z|X} \| \mathbb{P}_{1,Z|X} | \mathbb{P}_{0,X}) \end{aligned}$$

For any X the distribution  $\mathbb{P}_{1,Z|X}$  is uniform, and  $\mathbb{P}_{0,Z|X}$ ,  $\mathbb{P}_{0,X}$  have pmf  $\mathbb{E}_{\eta|X}p_{\eta}^{\otimes m}$  and  $\mathbb{E}_{\eta}p_{\eta}^{\otimes n}$  respectively. Once again, by Lemma 6 we may turn our attention to the  $\chi^2$ -divergence. Given X, let  $\eta'$  have the same distribution as  $\eta$  and be independent of it. Then

$$\begin{split} \chi^2(\mathbb{P}_{0,Z|X} \| \mathbb{P}_{1,Z|X} | \mathbb{P}_{0,X}) + 1 &= (2k)^m \mathbb{E}_X \sum_{j \in [2k]^m} \mathbb{E}_{\eta|X} \mathbb{E}_{\eta'|X} \prod_{i=1}^n p_\eta(j_i) p_{\eta'}(j_i) \\ &= \mathbb{E}_{\eta\eta'} (1 + \frac{\varepsilon^2 \langle \eta, \eta' \rangle}{k})^m \\ &\leq \mathbb{E}_{\eta\eta'} \prod_{j \in [k]} \exp(\frac{\varepsilon^2 m \eta_j \eta'_j}{k}), \end{split}$$

where we used Lemma B.6. Let  $N = (N_1, \ldots, N_k)$  be the vector of counts indicating the number of the  $X_1, \ldots, X_n$  that fall into the bins  $\{2j - 1, 2j\}$  for  $j \in [k]$ . Clearly  $N \sim \text{Mult}(n, (\frac{1}{k}, \ldots, \frac{1}{k}))$ . Let us focus on a specific bin  $\{2j - 1, 2j\}$  and define the bin-conditional pmf

$$p_{\pm}(x) = \begin{cases} \frac{1}{2}(1 \pm \varepsilon) & \text{if } x = 2j - 1, \\ \frac{1}{2}(1 \mp \varepsilon) & \text{if } x = 2j \\ 0 & \text{otherwise,} \end{cases}$$

where we drop the dependence on j in the notation. Let  $X_{i_1}, \ldots, X_{i_{N_j}}$  be the  $N_j$  observations falling in  $\{2j-1, 2j\}$ . Given  $N_j$ , the pmf of  $X_{i_1}, \ldots, X_{i_{N_j}}$  is  $\frac{1}{2}(p_+^{\otimes N_j} + p_-^{\otimes N_j})$ . We have  $\eta_j \eta'_j \in \{\pm 1\}$  almost surely, and analogously to Section B-A we may compute

$$\mathbb{P}(\eta_{j}\eta_{j}' = 1|N_{j}) = \mathbb{E}_{X|N_{j}}\mathbb{P}(\eta_{j}\eta_{j}' = 1|X)$$

$$= \mathbb{E}_{X|N_{j}}\left[\mathbb{P}(\eta_{j} = 1|X)^{2} + \mathbb{P}(\eta_{j} = -1|X)^{2}\right]$$

$$= \frac{1}{2} + \frac{1}{4}\left(\chi^{2}(p_{+}^{\otimes N_{j}} \| \frac{1}{2}(p_{+}^{\otimes N_{j}} + p_{-}^{\otimes N_{j}})) + \chi^{2}(p_{-}^{\otimes N_{j}} \| \frac{1}{2}(p_{+}^{\otimes N_{j}} + p_{-}^{\otimes N_{j}})\right)$$

$$\le \frac{1}{2} + \frac{1}{8}\left(\chi^{2}(p_{-}^{\otimes N_{j}} \| p_{+}^{\otimes N_{j}}) + \chi^{2}(p_{+}^{\otimes N_{j}} \| p_{-}^{\otimes N_{j}})\right).$$

We can bound the two  $\chi^2$ -divergences by

$$\chi^2(p_{\pm}^{\otimes N_j} \| p_{\mp}^{\otimes N_j}) + 1 = \left(\frac{1 + \frac{3}{2}\varepsilon^2}{1 - \varepsilon^2}\right)^{N_j} \le (1 + 3\varepsilon^2)^{N_j},$$

provided  $\varepsilon \leq c$  for some universal constant c > 0. Using Lemma 8, we obtain the bound

$$\begin{split} \mathbb{E}_{N} \prod_{j \in [k]} \mathbb{E}_{\eta \eta' \mid N_{j}} \exp(\frac{\varepsilon^{2} m \eta_{j} \eta'_{j}}{k}) &\leq \mathbb{E}_{N} \prod_{j \in [k]} \left( \frac{1}{2} (\exp(\frac{\varepsilon^{2} m}{k}) - \exp(-\frac{\varepsilon^{2} m}{k}))(1 + (1 + 2\varepsilon^{2})^{N_{j}}) + \exp(-\frac{\varepsilon^{2} m}{k}) \right) \\ &\leq \left( \frac{1}{2} (\exp(\frac{\varepsilon^{2} m}{k}) - \exp(-\frac{\varepsilon^{2} m}{k}))(1 + \exp(\frac{2\varepsilon^{2} n}{k})) + \exp(-\frac{\varepsilon^{2} m}{k}) \right)^{k}. \end{split}$$

Now, under the assumption that  $m \lor n \lesssim k/\varepsilon^2$  for some small enough implied constant, the above can be further bounded by

$$\leq (1 + c\frac{\varepsilon^4 mn}{k^2})^k$$
$$\leq \exp(\frac{c\varepsilon^4 mn}{k}),$$

for a universal constant c > 0. In other words, for  $n \lor m \leq k/\varepsilon^2$  likelihood-free hypothesis testing is impossible unless  $mn \geq k/\varepsilon^4$ . Combining everything yields the desired bounds.

Our second lower bound, tight in the regime  $n \le m$ , follows by reduction to two-sample testing Proposition 1. **Proposition 13.** For a constant c > 0 independent of  $\varepsilon$  and k,

$$c\{m \ge 1/\varepsilon^2, n^2m \ge k^2/\varepsilon^4, n \le m\} \supseteq \mathcal{R}_{\mathsf{LF}}(\varepsilon, \mathsf{TV}, \mathcal{P}_{\mathsf{D}}) \cap \mathbb{N}^2_{n < m}$$

where  $\mathbb{N}^2_{n \leq m} = \{(n,m) \in \mathbb{N}^2 : n \leq m\}.$ 

*Proof.* Follows from (III.6) and the lower bound construction in [23].

1) Valiant's wishful thinking theorem.: For our third and final lower bound (which is tight in the regime  $m \le n \le k$ ) we apply a method developed by Valiant.

**Definition 6.** For distributions  $p_1, \ldots, p_\ell$  on [k] and  $(n_1, \ldots, n_\ell) \in \mathbb{N}^\ell$ , we define the  $(n_1, \ldots, n_\ell)$ -based moments of  $(p_1, \ldots, p_\ell)$  as

$$m(a_1, \dots, a_\ell) = \sum_{i=1}^k \prod_{j=1}^\ell (n_j p_j(i))^{a_j}$$

for  $(a_1, \ldots, a_\ell) \in \mathbb{N}^\ell$ .

Let  $p^+ = (p_1^+, \ldots, p_\ell^+)$  and  $p^- = (p_1^-, \ldots, p_\ell^-)$  be  $\ell$ -tuples of distributions on [k] and suppose we observe samples  $\{X^{(i)}\}_{i\in[\ell]}$ , where the number of observations in  $X^{(i)}$  is  $\operatorname{Poi}(n_i)$ . Let  $H^{\pm}$  denote the hypothesis that the samples came from  $p^{\pm}$ , up to an arbitrary relabeling of the alphabet [k]. It can be shown that to test  $H^+$  against  $H^-$ , we may assume without loss of generality that our test is invariant under relabeling of the support, or in other words, is a function of the *fingerprints*. The fingerprint f of a sample  $\{X^{(i)}\}_{i\in[\ell]}$  is the function  $f:\mathbb{N}^\ell \to \mathbb{N}$  which given  $(a_1, \ldots, a_\ell) \in \mathbb{N}^\ell$  counts the number of bins in [k] which have exactly  $a_i$  occurences in the sample  $X^{(i)}$ .

**Theorem 5** ([59, Wishful thinking]). Suppose that  $|p_i^{\pm}|_{\infty} \leq \eta/n_i$  for all  $i \in [\ell]$  for some  $\eta > 0$ , and let  $m^+$  and  $m^-$  denote the  $(n_1, \ldots, n_{\ell})$ -based moments of  $p^+, p^-$  respectively. Let  $f^{\pm}$  denote the distribution of the fingerprint under  $H^{\pm}$  respectively. Then

$$\mathsf{TV}(f^+, f^-) \le 2(e^{\eta \ell} - 1) + e^{\ell(\eta/2 + \log 3)} \sum_{a \in \mathbb{N}^\ell} \frac{|m^+(a) - m^-(a)|}{\sqrt{1 + m^+(a) \vee m^-(a)}}.$$

*Proof.* The proof is a straightforward adaptation of [59] and thus we omit it.

*Remark* 10. Although Theorem 5 assumes a random (Poisson distributed) number of samples, the results carry over to the deterministic case with no modification, due to the sub-exponential concentration of the Poisson distribution.

We are ready to prove our likelihood-free hypothesis testing lower bound using Theorem 5.

**Proposition 14.** For a constant c > 0 independent of  $\varepsilon$  and k,

$$c\{m \ge 1/\varepsilon^2, n^2m \ge k^2/\varepsilon^4, m \le n\} \supseteq \mathcal{R}_{\mathsf{LF}}(\varepsilon, \mathsf{TV}, \mathcal{P}_{\mathsf{D}}) \cap \mathbb{N}^2_{m \le n}\}$$

where  $\mathbb{N}_{m\leq n}^2 = \{(n,m)\in\mathbb{N}^2:m\leq n\}.$ 

*Proof.* We focus on the regime  $n \le k$ , as otherwise the result is subsumed by Proposition 12. Suppose that  $\varepsilon \in (0, 1/2), \eta = 0.01$  (say) and  $n/\eta \le k/2$ . Define  $\gamma = n/\eta$  and let p, q be pmfs on [k] with weight  $(1 - \varepsilon)/\gamma$  on  $[\gamma]$  and k/4 light elements with weight  $4\varepsilon/k$  on [k/2, 3k/4] and [3k/4, k] respectively. To apply Valiant's wishful

thinking theorem, we take  $p^+ = (p, q, p)$  and  $p^- = (p, q, q)$  with corresponding hypotheses  $H^{\pm}$ . The (n, n, m)-based moments of  $p^{\pm}$  are given by

$$\frac{1}{n^{a+b}m^c}m^+(a,b,c) = \begin{cases} k & \text{if } a+c=0, b=0\\ \left(\frac{1-\varepsilon}{\alpha}\right)^{a+b+c}\alpha + \left(\frac{4\varepsilon}{k}\right)^{a+b+c}\frac{k}{4} & \text{if } a+c=0 \text{ xor } b=0\\ \left(\frac{1-\varepsilon}{\alpha}\right)^{a+b+c}\alpha & \text{if } a+c\geq 1, b\geq 1, \end{cases}$$
$$\frac{1}{n^{a+b}m^c}m^-(a,b,c) = \begin{cases} k & \text{if } a=0, b+c=0\\ \left(\frac{1-\varepsilon}{\alpha}\right)^{a+b+c}\alpha + \left(\frac{4\varepsilon}{k}\right)^{a+b+c}\frac{k}{4} & \text{if } a=0 \text{ xor } b+c=0\\ \left(\frac{1-\varepsilon}{\alpha}\right)^{a+b+c}\alpha & \text{if } a\geq 1, b+c\geq 1. \end{cases}$$

By the wishful thinking theorem we know that

$$\mathsf{TV}(f^+, f^-) \le 0.061 + 27.41 \sum_{a, b, c \in \mathbb{N}} \frac{|m^+(a, b, c) - m^-(a, b, c)|}{\sqrt{1 + \max(m^+, m^-)}}.$$

Let us consider the possible values of  $|m^+(a, b, c) - m^-(a, b, c)|$ . It is certainly zero if  $a \wedge b \ge 1$  or a = b = c = 0. Suppose that a = 0 so that necessarily  $b + c \ge 1$ . Then

$$\frac{1}{n^b m^c} |m^+(0,b,c) - m^-(0,b,c)| = \left(\frac{4\varepsilon}{k}\right)^{b+c} \frac{k}{4} \mathbb{1}(b \wedge c \ge 1)$$

Using the symmetry between a and b and that  $1 + m^+ \vee m^- \ge n^b m^c ((1 - \varepsilon)/\gamma)^{b+c} \gamma$  (for  $m^+ \ne m^-$ ), we can bound the infinite sum above as

$$\lesssim \sum_{b,c \ge 1} \frac{n^b m^c k^{1-(b+c)} \varepsilon^{b+c}}{\sqrt{n^b m^c \gamma^{1-(b+c)} (1-\varepsilon)^{b+c}}}$$
$$\lesssim \sum_{b,c \ge 1} n^{b/2} m^{c/2} \left(\frac{\sqrt{\gamma}}{k}\right)^{b+c-1} \varepsilon^{b+c}$$

Plugging in  $\gamma = n/\eta \asymp n$ , and using  $m \le n \le k$ , we obtain

$$\begin{split} \mathsf{TV}(f^+, f^-) &- 0.061 \lesssim \sum_{b,c \geq 1} n^{b+\frac{c}{2}-\frac{1}{2}} m^{c/2} \frac{1}{k^{b+c-1}} \varepsilon^{b+c} \\ &= \frac{n\sqrt{m}\varepsilon^2}{k} \sum_{b,c \geq 0} \left(\frac{n}{k}\right)^{b+\frac{c}{2}} \left(\frac{m}{k}\right)^{\frac{c}{2}} \varepsilon^{b+c} \\ &\leq \frac{n\sqrt{m}\varepsilon^2}{k} \sum_{b,c \geq 0} \varepsilon^{b+c} \\ &\lesssim \frac{n\sqrt{m}\varepsilon^2}{k}, \end{split}$$

where we use that  $\varepsilon < 1/2$ . Thus, likelihood-free hypothesis testing is impossible for  $m \le n$  unless  $n^2m \gtrsim k^2/\varepsilon^4$ .

## APPENDIX C Proof of Theorem 4

## A. Upper bound

We deduce the upper bound by applying the corresponding result for  $\mathcal{P}_D$  as a black-box procedure.

**Theorem 6** ([56]). For a constant independent of  $\varepsilon$  and k,

$$n_{\mathsf{GoF}}(\varepsilon,\mathsf{H},\mathcal{P}_{\mathsf{D}}) \asymp \sqrt{k}/\varepsilon^2.$$

Write  $\mathcal{G}_{\ell}$  for the regular grid of size  $\ell^d$  on  $[0, 1]^d$  and let  $P_{\ell}$  denote the  $L^2$ -projector onto the space of functions piecewise constant on the cells of  $\mathcal{G}_{\ell}$ . For convenience let us recall Proposition 3.

**Proposition 15.** Let  $f, g \in \mathcal{P}_{\mathsf{H}}(\beta, d, C_{\mathsf{H}})$  with  $\beta \in (0, 1]$  and suppose that  $\mathsf{H}(f, g) \gtrsim \varepsilon$ . Then

$$\mathsf{H}(f,g) \lesssim \mathsf{H}(P_{\kappa}f,P_{\kappa}g)$$

for  $\kappa \simeq \varepsilon^{-2/\beta}$  where the constants depend only on  $\beta, d, C_{\mathsf{H}}$ .

With the above approximation result, the proof of Theorem 4 is straightforward.

Proof of Theorem 4. Suppose we are testing goodness-of-fit to  $f_0 \in \mathcal{P}_{\mathsf{H}}$  based on an i.i.d. sample  $X_1, \ldots, X_n$  from  $f \in \mathcal{P}_{\mathsf{H}}$ . Take  $\kappa \asymp \varepsilon^{-2/\beta}$  and bin the observations on  $\mathcal{G}_{\kappa}$ , denoting the pmf of the resulting distribution as  $p_f$ . Then, under the alternative hypothesis that  $\mathsf{H}(f, f_0) \ge \varepsilon$ , by Proposition 3

$$\varepsilon \lesssim \mathsf{H}(P_{\kappa}f_0, P_{\kappa}f) = \mathsf{H}(p_{f_0}, p_f).$$

In particular, applying the algorithm achieving the upper bound in Theorem 6 to the binned observations, we see that  $n \gtrsim \sqrt{\kappa^d}/\varepsilon^2 = \varepsilon^{-(2\beta+d)/\beta}$  samples suffice.

## B. Lower bound

The proof is extremely similar to the TV case, except we put the perturbations at density level  $\varepsilon^2$  instead of 1. *Proof.* Let  $\phi : [0,1] \to [0,1]$  be a smooth function such that  $\phi(x) = 0$  for  $x \le 1/3$  and  $\phi(x) = 1$  for  $x \ge 2/3$ . Let  $h : \mathbb{R}^d \to \mathbb{R}$  also be smooth with  $\int h = 0$  and  $\int h^2 = 1$  and support in  $[0,1]^d$ . Take  $\varepsilon \in (0,1)$  and let

$$f_0(x) = \varepsilon^2 + \frac{\phi(x_1)}{\|\phi\|_1} (1 - \varepsilon^2),$$

which is a density on  $[0,1]^d$ . For a large integer  $\kappa$  and  $j \in [\kappa/3] \times [\kappa]^{d-1}$  let

$$h_j(x) = \kappa^{d/2} h(\kappa x - j + 1)$$

for  $x \in [0,1]^d$ . Then  $h_j$  is supported on  $[(j-1)/\kappa, j/\kappa] \subseteq [0,1/3] \times [0,1]^{d-1}$  and  $\int h_j^2 = 1$ . For  $\eta \in \{\pm 1\}^{[\kappa/3] \times [\kappa]^{d-1}}$  and  $\rho > 0$  let

$$f_{\eta}(x) = f_0 + \rho \sum_{j \in [\kappa/3] \times [\kappa]^{d-1}} \eta_j h_j(x)$$

Then  $f_{\eta}$  is positive provided that  $\varepsilon^2 \ge \rho \kappa^{d/2} |h|_{\infty} \asymp \rho \kappa^{d/2}$ . Further,  $||f_{\eta}||_{\mathcal{C}^{\beta}}$  is of constant order provided  $\rho \kappa^{d/2+\beta} \lesssim 1$ . Under these assumptions  $f_{\eta} \in \mathcal{P}_{\mathsf{H}}$ . Note that the Hellinger distance between  $f_{\eta}$  and  $f_0$  is

$$\begin{aligned} H^{2}(f_{0},f_{\eta}) &= \sum_{j \in [\kappa/3] \times [\kappa]^{d-1}} \int_{[\frac{j-1}{\kappa},\frac{j}{\kappa}]} \left(\sqrt{f_{0}(x)} - \sqrt{f_{\eta}(x)}\right)^{2} \mathrm{d}x \\ &= \sum_{j \in [\kappa/3] \times [\kappa]^{d-1}} \int_{[\frac{j-1}{\kappa},\frac{j}{\kappa}]} \frac{\rho^{2}h_{j}^{2}(x)}{(\sqrt{f_{0}(x)} + \sqrt{f_{\eta}(x)})^{2}} \mathrm{d}x \\ &\geq \sum_{j \in [\kappa/3] \times [\kappa]^{d-1}} \int_{[\frac{j-1}{\kappa},\frac{j}{\kappa}]} \frac{\rho^{2}h_{j}^{2}(x)}{4\varepsilon^{2}} \mathrm{d}x \\ &\gtrsim \frac{\rho^{2}\kappa^{d}}{\varepsilon^{2}}. \end{aligned}$$

Suppose we draw  $\eta$  uniformly at random. Via Ingster's trick we compute

$$\chi^{2}(\mathbb{E}_{\eta}f_{\eta}^{\otimes n}||f_{0}^{\otimes n}) + 1 = \int \mathbb{E}_{\eta\eta'} \prod_{i=1}^{n} \frac{f_{\eta}(x_{i})f_{\eta'}(x_{i})}{f_{0}(x_{i})} \mathrm{d}x_{1} \dots \mathrm{d}x_{n}$$
$$= \mathbb{E}_{\eta\eta'} \left( \int \frac{f_{\eta}(x)f_{\eta'}(x)}{f_{0}(x)} \mathrm{d}x \right)^{n}.$$

Looking at the integral term on the inside we get

$$\begin{split} \int \frac{f_{\eta}(x)f_{\eta'}(x)}{f_{0}(x)} \mathrm{d}x &= \int \frac{\left(f_{0}(x) + \rho \sum_{j \in [\kappa/3] \times [\kappa]^{d-1}} \eta_{j} h_{j}(x)\right) \left(f_{0}(x) + \rho \sum_{j \in [\kappa/3] \times [\kappa]^{d-1}} \eta'_{j} h_{j}(x)\right)}{f_{0}(x)} \mathrm{d}x \\ &= 1 + \rho \sum_{j} (\eta_{j} + \eta'_{j}) \int h_{j}(x) \mathrm{d}x + \rho^{2} \sum_{j} \eta_{j} \eta'_{j} \int \frac{h_{j}(x)^{2}}{f_{0}(x)} \mathrm{d}x \\ &= 1 + \frac{\rho^{2}}{\varepsilon^{2}} \sum_{j} \eta_{j} \eta'_{j} \int h_{j}(x)^{2} \mathrm{d}x \\ &= 1 + \frac{\rho^{2}}{\varepsilon^{2}} \langle \eta, \eta' \rangle, \end{split}$$

where we've used that  $h_j$  and  $h_{j'}$  have disjoint support unless j = j',  $\int h_j = 0$ ,  $\int h_j^2 = 1$ , and that  $f_0(x) = \varepsilon^2$  for all x with  $x_1 \le 1/3$ . Plugging in, using the inequalities  $1 + x \le \exp(x)$  and  $\cosh(x) \le \exp(x^2)$  we obtain

$$\begin{split} \chi^2(\mathbb{E}_{\eta} f_{\eta}^{\otimes n} \| f_0^{\otimes n}) + 1 &\leq \mathbb{E}_{\eta\eta'} (1 + \frac{\rho^2}{\varepsilon^2} \langle \eta, \eta' \rangle)^n \\ &\leq \mathbb{E}_{\eta\eta'} \exp(\frac{\rho^2 n}{\varepsilon^2} \langle \eta, \eta' \rangle) \\ &= \cosh(\frac{\rho^2 n}{\varepsilon^2})^{\kappa^d/3} \\ &\leq \exp(\frac{\rho^4 n^2 \kappa^d}{3\varepsilon^4}). \end{split}$$

Choosing  $\kappa = \varepsilon^{-2/\beta}$  and  $\rho = \varepsilon^{(2\beta+d)/\beta}$  we see that goodness-of-fit testing of  $f_0$  is impossible unless

$$n \gtrsim \frac{\varepsilon^2}{\rho^2 \kappa^{d/2}} = \varepsilon^{-\frac{2\beta+d}{\beta}}.$$

## APPENDIX D AUXILIARY TECHNICAL RESULTS

A. Proof of Lemma 1

*Proof.* We prove the upper bound first. Let  $\mathbb{P}_0, \mathbb{P}_1 \in \mathcal{P}$  be arbitrary. Then by Lemma 6,

$$\begin{split} \inf_{\psi} \max_{i=0,1} \mathbb{P}_{i}^{\otimes m}(\psi \neq i) &\leq \inf_{\psi} \left( \mathbb{P}_{0}^{\otimes m}(\psi = 1) + \mathbb{P}_{1}^{\otimes m}(\psi = 0) \right) \\ &= 1 - \mathsf{TV}(\mathbb{P}_{0}^{\otimes m}, \mathbb{P}_{1}^{\otimes m}) \\ &\leq 1 - \frac{1}{2} \mathsf{H}^{2}(\mathbb{P}_{0}^{\otimes m}, \mathbb{P}_{1}^{\otimes m}) \triangleq (\dagger). \end{split}$$

By tensorization of the Hellinger affinity, we have

$$\mathsf{H}^{2}(\mathbb{P}_{0}^{\otimes m}, \mathbb{P}_{1}^{\otimes m}) = 2 - 2\left(1 - \frac{1}{2}\mathsf{H}^{2}(\mathbb{P}_{0}, \mathbb{P}_{1})\right)^{m}.$$
 (D.1)

Plugging in, along with  $1 + x \le \exp(x)$  gives

$$(\dagger) \le \exp(-\frac{m}{2}\mathsf{H}^2\left(\mathbb{P}_0^{\otimes m}, \mathbb{P}_1^{\otimes m})\right).$$

Taking  $m > 2\log(3)/\mathsf{H}^2(\mathbb{P}_0,\mathbb{P}_1)$  shows the existence of a successful test. Let us turn to the lower bound. Using Lemma 6 we have

$$\begin{split} \inf_{\psi} \max_{i=0,1} \mathbb{P}_{i}^{\otimes m}(\psi \neq i) \geq \frac{1}{2} \left( 1 - \mathsf{TV}(\mathbb{P}_{0}^{\otimes m}, \mathbb{P}_{1}^{\otimes m}) \right) \\ \geq \frac{1}{2} \left( 1 - \mathsf{H}(\mathbb{P}_{0}^{\otimes m}, \mathbb{P}_{1}^{\otimes m}) \right). \end{split}$$

Note that it is enough to restrict the maximization in Lemma 1 to  $\mathbb{P}_0$ ,  $\mathbb{P}_1 \in \mathcal{P}$  with  $\mathsf{H}^2(\mathbb{P}_0, \mathbb{P}_1) < 1$ . Now, by (D.1) and the inequalities  $1 - x \ge e^{-2x}$  for  $x \in [0, 1/2]$  and  $e^{-x} \ge 1 - x$  for  $x \ge 0$ , we obtain

$$\begin{aligned} \mathsf{H}^{2}(\mathbb{P}_{0}^{\otimes m},\mathbb{P}_{1}^{\otimes m}) &= 2 - 2\left(1 - \frac{1}{2}\mathsf{H}^{2}(\mathbb{P}_{0},\mathbb{P}_{1})\right)^{*} \\ &\leq 2 - 2\exp(-m\mathsf{H}^{2}(\mathbb{P}_{0},\mathbb{P}_{1})) \\ &\leq 2m\mathsf{H}^{2}(\mathbb{P}_{0},\mathbb{P}_{1}). \end{aligned}$$

Taking  $m = 1/(18 \mathsf{H}^2(\mathbb{P}_0, \mathbb{P}_1))$  concludes the proof via Lemma 5.

## B. Proof of Proposition 4

For arbitrary  $f \in L^2(\mu)$  write  $f_i = \langle f \phi_i \rangle$  and  $f_{ii'} = \langle f \phi_i \phi_{i'} \rangle$ , assuming that the quantities involved are well-defined. We record some useful properties of  $P_r$  that we will use throughout our proofs.

**Lemma 7.**  $P_r$  is self-adjoint and has operator norm

$$||P_r|| \triangleq \sup_{f \in L^2(\mu) : ||f||_2 \le 1} ||P_r(f)||_2 \le 1.$$

Suppose that  $f, g, h, t \in L^2(\mu)$  and that each quantity below is finite. Then

$$\sum_{ii'} f_i g_{i'} h_{ii'} = \langle h P_r(f) P_r(g) \rangle,$$
  
$$\sum_{ii'} f_i g_i h_{i'} t_{i'} = \langle f P_r(g) \rangle \langle h P_r(t) \rangle$$
  
$$\sum_{ii'} f_{ii'} g_{ii'} = \sum_i \langle f \phi_i P_r(g \phi_i) \rangle,$$

where the summation is over  $i, i' \in [r]$ .

*Proof.* Let  $P_r^{\perp}$  denote the orthogonal projection onto the orthogonal complement of span( $\{\phi_1, \ldots, \phi_r\}$ ). Then for any  $f, g \in L^2(\mu)$  we have

$$\langle fP_r(g)\rangle = \langle (P_r(f) + P_r^{\perp}(f))P_r(g)\rangle = \langle P_r(f)P_r(g)\rangle = \langle P_r(f)g\rangle,$$

where the last equality is by symmetry. We also have

$$\|P_r(f)\|_2^2 \le \|P_r(f)\|_2^2 + \|P_r^{\perp}(f)\|_2^2 = \|P_r(f) + P_r^{\perp}(f)\|^2 = \|f\|_2^2$$

Let  $f, g, h, t \in L^2(\mu)$ . Then

$$\sum_{ii'} f_i g_{i'} h_{ii'} = \sum_i f_i \sum_{i'} g_{i'} h_{ii'} = \sum_i f_i \langle hP_r g, \phi_i \rangle = \langle P_r(f) hP_r(g) \rangle$$
$$\sum_{ii'} f_i g_i h_{i'} t_{i'} = (\sum_i f_i g_i) (\sum_{i'} h_{i'} t_{i'}) = \langle fP_r(g) \rangle \langle hP_r(t) \rangle$$
$$\sum_{ii'} f_{ii'} g_{ii'} = \sum_i \langle f\phi_i \sum_{i'} \langle g\phi_i \phi_{i'} \rangle \phi_{i'} \rangle = \sum_i \langle f\phi_i P_r(g\phi_i) \rangle.$$

*Proof of Proposition 4.* Let us label the different terms of the statistic  $T_{LF}^{-d}$ :

$$T_{\mathsf{LF}}^{-\mathsf{d}} = \sum_{i=1}^{r} \left\{ \frac{2}{n^2} \sum_{j < j'}^{n} \phi_i(X_j) \phi_i(X_{j'}) - \frac{2}{n^2} \sum_{j < j'}^{n} \phi_i(Y_j) \phi_i(Y_{j'}) \right. \\ \left. - \frac{2}{nm} \sum_{j=1}^{n} \sum_{u=1}^{m} \phi_i(X_j) \phi_i(Z_u) + \frac{2}{nm} \sum_{j=1}^{n} \sum_{u=1}^{m} \phi_i(Y_j) \phi_i(Z_u) \right\} \\ \left. = \frac{2}{n^2} \mathsf{I} - \frac{2}{n^2} \mathsf{II} - \frac{2}{nm} \mathsf{III} + \frac{2}{nm} \mathsf{IV}. \right\}$$

Recall that  $X, Y, Z \sim f^{\otimes n}, g^{\otimes n}, h^{\otimes m}$  respectively. A straightforward computation yields

$$\mathbb{E}T_{\mathsf{LF}} = \|P_r(f-h)\|_2^2 - \|P_r(g-h)\|_2^2 - \frac{1}{n} \left(\|P_r[f]\|_2^2 - \|P_r[g]\|_2^2\right).$$

We decompose the variance as

$$\operatorname{var}(T_{\mathsf{LF}}) = \frac{4}{n^4} \operatorname{var}(\mathsf{I}) + \frac{4}{n^4} \operatorname{var}(\mathsf{II}) + \frac{4}{n^2 m^2} \operatorname{var}(\mathsf{III}) + \frac{4}{n^2 m^2} \operatorname{var}(\mathsf{IV}) - \frac{8}{n^3 m} \operatorname{Cov}(\mathsf{I}, \mathsf{III}) - \frac{8}{n^3 m} \operatorname{Cov}(\mathsf{II}, \mathsf{IV}) - \frac{8}{n^2 m^2} \operatorname{Cov}(\mathsf{III}, \mathsf{IV}),$$

where we used independence of the pairs (I, II), (I, IV), (II, III). Expanding the variances we obtain

$$\begin{aligned} \operatorname{var}(\mathsf{I}) &= \sum_{ii'} \left( \binom{n}{2} (f_{ii'}^2 - f_i^2 f_{i'}^2) + (\binom{n}{2}^2 - \binom{n}{2} - \binom{4}{2} \binom{n}{4}) (f_i f_{i'} f_{ii'} - f_i^2 f_{i'}^2) \right) \\ \operatorname{var}(\mathsf{II}) &= \sum_{ii'} \left( \binom{n}{2} (g_{ii'}^2 - g_i^2 g_{i'}^2) + (\binom{n}{2}^2 - \binom{n}{2} - \binom{4}{2} \binom{n}{4}) (g_i g_{i'} g_{ii'} - g_i^2 g_{i'}^2) \right) \\ \operatorname{var}(\mathsf{III}) &= \sum_{ii'} \left( nm(f_{ii'} h_{ii'} - f_i f_{i'} h_i h_{i'}) + nm(m-1)(f_{ii'} h_i h_{i'} - f_i f_{i'} h_i h_{i'}) + mn(n-1)(f_i f_{i'} h_{ii'} - f_i f_{i'} h_i h_{i'}) \right) \\ \operatorname{var}(\mathsf{IV}) &= \sum_{ii'} \left( nm(h_{ii'} g_{ii'} - h_i h_{i'} g_i g_{i'}) + mn(n-1)(h_{ii'} g_i g_{i'} - h_i h_{i'} g_i g_{i'}) + mn(m-1)(g_{ii'} h_i h_{i'} - h_i h_{i'} g_i g_{i'}) \right). \end{aligned}$$

For the covariance terms we obtain

$$Cov(I, III) = \sum_{ii'} 2m \binom{n}{2} (f_{ii'}f_ih_{i'} - f_i^2 f_{i'}h_{i'})$$
$$Cov(II, IV) = \sum_{ii'} 2m \binom{n}{2} (g_{ii'}g_ih_{i'} - g_i^2 g_{i'}h_{i'})$$
$$Cov(III, IV) = \sum_{ii'} mn^2 (h_{ii'}f_ig_{i'} - f_ig_{i'}h_ih_{i'}).$$

We can now start collecting the terms, applying the calculation rules from Lemma 7 repeatedly. Note that  $\binom{n}{2}^2 - \binom{n}{2} - \binom{4}{2}\binom{n}{4} = n^3 - 3n^2 + 2n$ , and by inspection we can conclude that  $1/n, 1/m, 1/m, 1/n^2$  and  $1/n^3$  are the only terms with nonzero coefficients. We look at each of them one-by-one:

$$Coef(\frac{1}{n}) = \sum_{ii'} \left( \underbrace{4(f_i f_{ii'} f_{ii'} - f_i^2 f_{i'}^2)}_{var(I)} + \underbrace{4(g_i g_{i'} g_{ii'} - g_i^2 g_{i'}^2)}_{var(II)} + \underbrace{4(h_i h_{i'} f_{ii'} - f_i f_{i'} h_i h_{i'})}_{var(III)} + \underbrace{4(g_{ii'} h_{ih'} - h_i h_{i'} g_i g_{i'})}_{var(IV)} - \underbrace{8(f_{ii'} f_i h_{i'} - f_i^2 f_{i'} h_{i'})}_{Cov(I,III)} - \underbrace{8(g_{ii'} g_i h_{i'} - g_i^2 g_{i'} h_{i'})}_{Cov(II,IV)} \right)$$
  
$$\leq 4A_{ffh} + 4A_{ggh},$$

recalling the definition  $A_{uvt} = \langle u [P_r(v-t)]^2 \rangle$  for  $u, v, t \in L^2(\mu)$ . Similarly, we get

$$\operatorname{Coef}(\frac{1}{m}) = \sum_{ii'} \left( \underbrace{4(h_{ii'}f_if_{i'} - f_if_{i'}h_ih_{i'})}_{\operatorname{var}(III)} + \underbrace{4(h_{ii'}g_ig_{i'} - h_ih_{i'}g_ig_{i'})}_{\operatorname{var}(IV)} - \underbrace{8(h_{ii'}f_ig_{i'} - f_ih_ih_{i'}g_{i'})}_{\operatorname{Cov}(III,IV)} \right) \\ \leq 4A_{hfg}.$$

For the lower order terms we obtain

$$\operatorname{Coef}(\frac{1}{nm}) = \sum_{ii'} \left( \underbrace{4(f_{ii'}h_{ii'} - f_if_{i'}h_ih_{i'}) - 4(f_{ii'}h_ih_{i'} - f_if_{i'}h_ih_{i'}) - 4(f_if_{i'}h_{ii'} - f_if_{i'}h_ih_{i'})}_{\operatorname{var}(III)} + \underbrace{4(h_{ii'}g_{ii'} - h_ih_{i'}g_ig_{i'}) - 4(h_{ii'}g_ig_{i'} - h_ih_{i'}g_ig_{i'}) - 4(g_{ii'}h_ih_{i'} - h_ih_{i'}g_ig_{i'})}_{\operatorname{var}(IV)} \right)_{\operatorname{var}(IV)} \leq 4\langle fP_r(h)\rangle^2 + 4\langle gP_r(h)\rangle^2 + 4B_{fh} + 4B_{gh} \\ \lesssim |B_{fh}| + |B_{gh}| + ||f + g + h||_2^4$$

where we recall the definition  $B_{uv} = \sum_i \langle u\phi_i P_r(v\phi_i) \rangle$  for  $u, v \in L^2(\mu)$  and apply the Cauchy-Schwarz inequality. Next, we look at the coefficient of  $1/n^2$  and find

$$\begin{aligned} \operatorname{Coef}(\frac{1}{n^2}) &= \sum_{ii'} \left( \underbrace{2(f_{ii'}^2 - f_i^2 f_{i'}^2) - 12(f_{ii'} f_i f_{i'} - f_i^2 f_{i'}^2)}_{\operatorname{var}(\mathsf{I})} + \underbrace{2(g_{ii'}^2 - g_i^2 g_{i'}^2) - 12(g_{ii'} g_i g_{i'} - g_i^2 g_{i'}^2)}_{\operatorname{var}(\mathsf{II})} \right. \\ &+ \underbrace{8(f_{ii'} f_i h_{i'} - f_i^2 f_{i'} h_{i'})}_{\operatorname{Cov}(\mathsf{I},\mathsf{III})} + \underbrace{8(g_{ii'} g_i h_{i'} - g_i^2 g_{i'} h_{i'})}_{\operatorname{Cov}(\mathsf{I},\mathsf{IV})} \right)}_{\operatorname{Cov}(\mathsf{I},\mathsf{IV})} \\ &\lesssim |B_{ff}| + |B_{gg}| + ||f + g + h||_2^4 + ||f + g + h||_2^3. \end{aligned}$$

Finally, we look at the coefficient of  $1/n^3$ :

$$\operatorname{Coef}\left(\frac{1}{n^{3}}\right) = \sum_{ii'} \left(\underbrace{-2(f_{ii'}^{2} - f_{i}^{2}f_{i'}^{2}) + 8(f_{ii'}f_{i}f_{i'} - f_{i}^{2}f_{i'}^{2})}_{\operatorname{Cov}(\mathsf{I},\mathsf{III})} \underbrace{-2(g_{ii'}^{2} - g_{i}^{2}g_{i'}^{2}) + 8(g_{ii'}g_{i}g_{i'} - g_{i}^{2}g_{i'}^{2})}_{\operatorname{Cov}(\mathsf{I},\mathsf{III})}\right)$$

$$\lesssim |B_{ff}| + |B_{gg}| + ||f + g||_{2}^{3}.$$

## C. Lemma 8

**Lemma 8.** Suppose that a, b, c > 0 and  $N = (N_1, \ldots, N_k) \sim \text{Mult}(n, (\frac{1}{k}, \ldots, \frac{1}{k}))$ . Then

$$\mathbb{E}_N \prod_{j \in [k]} (a + b(1+c)^{N_j}) \le (a + be^{cn/k})^k.$$

*Proof.* Expanding via the binomial formula and using the fact that sums of  $N_j$ 's are binomial random variables, we get

$$\mathbb{E}_N \prod_{j \in k} (a+b(1+c)^{N_j}) = \mathbb{E} \sum_{\ell=0}^k \binom{k}{\ell} b^\ell (1+c)^{\operatorname{Bin}(n,\ell/k)} a^{k-\ell}$$
$$= \sum_{\ell=0}^k \binom{k}{\ell} b^\ell (1+\frac{c\ell}{k})^n a^{k-\ell}$$
$$\leq (a+be^{cn/k})^k,$$

where we used  $1 + x \leq e^x$  for all  $x \in \mathbb{R}$ .

## D. Proof of Proposition 3

Let us write  $a_+ \triangleq a \lor 0$  for both functions and real numbers. We start with some known results of approximation theory.

Definition 7. For  $f:[0,1]^d \to \mathbb{R}$  define the modulus of continuity as

$$\omega(\delta; f) = \sup_{\|x-y\|_2 \le \delta} |f(x) - f(y)|$$

**Lemma 9.** For any real-valued function f and  $\delta \ge 0$ ,

$$\omega(\delta;\sqrt{f_+}) \leq \omega(\delta;f)^{1/2}$$

*Proof.* Follows from the inequality  $|\sqrt{a_+} - \sqrt{b_+}|^2 \le |a - b|$  valid for all  $a, b \in \mathbb{R}$ . Lemma 10. Let  $f : [0,1]^d \to \mathbb{R}$  be  $\beta$ -smooth for  $\beta \in (0,1]$ . Then

$$\omega(\delta; f) \le c \, \delta^\beta$$

for a constant c depending only on  $||f||_{C^{\beta}}$ .

Proof. Follows by the definition of Hölder continuity.

**Lemma 11** ([66, Theorem 4]). For any continuous function  $f : [0,1]^d \to \mathbb{R}$  the best polynomial approximation  $p_n$  of degree n satisfies

$$||p_n - f||_{\infty} \le c \omega \left(\frac{d^{3/2}}{n}; f\right)$$

for a universal constant c > 0.

**Definition 8.** Given a function  $f:[0,1]^d \to \mathbb{R}$ ,  $\ell \ge 1$  and  $j \in [\ell]^d$ , let  $\pi_{j,\ell}f:[0,1]^d \to \mathbb{R}$  denote the function

$$\pi_{j,\ell}f(x) \triangleq f\left(\frac{x+j-1}{\ell}\right).$$

In other words,  $\pi_{j,\ell}f$  is equal to f zoomed in on the j'th bin of the regular grid  $\mathcal{G}_{\ell}$ .

Recall that here  $P_{\ell}$  denotes the  $L^2$  projector onto the space of functions piecewise constant on the bins of  $\mathcal{G}_{\ell}$ . We are ready for the proof of Proposition 3.

*Proof.* Let  $\kappa \ge r \ge 1$  whose values we specify later. We treat the parameters  $\beta, d, \|f\|_{\mathcal{C}^{\beta}}, \|g\|_{\mathcal{C}^{\beta}}$  as constants in our analysis. Let  $u_f : [0, 1]^d \to \mathbb{R}$  denote the (piecewise polynomial) function that is equal to the best polynomial approximation of  $\sqrt{f}$  on each bin of  $\mathcal{G}_{\kappa/r}$  with maximum degree  $\alpha$ . By lemmas 9 and 10 for any  $\ell \ge 1$  and  $j \in [\ell]^d$ 

$$\omega(\delta; \pi_{j,\ell}\sqrt{f}) \le \omega(\delta/\ell; \sqrt{f}) \lesssim (\delta/\ell)^{\beta/2}, \tag{D.2}$$

so that by Lemma 11

$$|u_f - \sqrt{f}|_{\infty} = \sup_{j \in [\kappa/r]^d} |\pi_{j,\kappa/r}(u_f - \sqrt{f})|_{\infty}$$
  
$$\lesssim \sup_{j \in [\kappa/r]^d} \omega(d^{3/2}/\alpha; \pi_{j,\kappa/r}\sqrt{f})$$
  
$$\lesssim (\alpha \kappa/r)^{-\beta/2}.$$

Regarding r as a constant independent of  $\kappa$ ,  $\alpha$  can be chosen large enough independently of  $\kappa$  such that  $|u_f - \sqrt{f}|_{\infty} \le c_1 \kappa^{-\beta/2}$  for  $c_1$  arbitrarily small. Define  $u_q$  analogously to  $u_f$ . We have the inequalities

We can write

$$||u_f - u_g||_2^2 = \frac{1}{(\kappa/r)^d} \sum_{j \in [\kappa/r]^d} ||\pi_{j,\kappa/r}(u_f - u_g)||_2^2$$

Now, by [42, Lemma 7.4] we can take r large enough (depending only on  $\beta$ , d,  $||f||_{C^{\beta}}$ ,  $||g||_{C^{\beta}}$ ) such that

$$\|\pi_{j,\kappa/r}(u_f - u_g)\|_2 \le c_2 \|P_r \pi_{j,\kappa/r}(u_f - u_g)\|_2$$

where the implied constant depends on the same parameters as r. Thus, we get

$$\begin{aligned} \mathsf{H}^{2}(f,g) &\leq 8c_{1}^{2}\kappa^{-\beta} + \frac{2c_{2}^{2}}{(\kappa/r)^{d}} \sum_{j \in [\kappa/r]^{d}} \|P_{r}\pi_{j,\kappa/r}(u_{f} - u_{g})\|_{2}^{2} \\ &\leq 8c_{1}^{2}\kappa^{-\beta} + \frac{6c_{2}^{2}}{(\kappa/r)^{d}} \sum_{j \in [\kappa/r]^{d}} \left( \|P_{r}\pi_{j,\kappa/r}u_{f} - \sqrt{P_{r}\pi_{j,\kappa/r}f}\|_{2}^{2} + \|P_{r}\pi_{j,\kappa/r}u_{g} - \sqrt{P_{r}\pi_{j,\kappa/r}g}\|_{2}^{2} \right) \\ &+ 6c_{2}^{2}\mathsf{H}^{2}(P_{\kappa}f, P_{\kappa}f), \end{aligned}$$

where  $c_1, c_2$  depend only on the unimportant parameters, and  $c_1$  can be taken arbitrarily small compared to  $c_2$ . We also used the fact that  $P_r \pi_{j,\kappa/r} = \pi_{j,\kappa/r} P_{\kappa}$ . Looking at the terms separately, we have

$$\begin{aligned} \|P_{r}\pi_{j,\kappa/r}u_{f} - \sqrt{P_{r}\pi_{j,\kappa/r}f}\|_{2} &\leq \|P_{r}\pi_{j,\kappa/r}u_{f} - P_{r}\sqrt{\pi_{j,\kappa/r}f}\|_{2} + \|P_{r}\sqrt{\pi_{j,\kappa/r}f} - \sqrt{P_{r}\pi_{j,\kappa/r}f}\|_{2} \\ &\leq c\kappa^{-\beta/2} + \|P_{r}\sqrt{\pi_{j,\kappa/r}f} - \sqrt{P_{r}\pi_{j,\kappa/r}f}\|_{2}, \end{aligned}$$

since  $P_r$  is a contraction by Lemma 7. We can decompose the second term as

$$\|P_r\sqrt{\pi_{j,\kappa/r}f} - \sqrt{P_r\pi_{j,\kappa/r}f}\|_2^2 = \sum_{\ell \in [r]^d} \int_{\left[\frac{\ell-1}{r},\frac{\ell}{r}\right]} \left( r^d \int_{\left[\frac{\ell-1}{r},\frac{\ell}{r}\right]} \sqrt{\pi_{j,\kappa/r}f(x)} \mathrm{d}x - \sqrt{r^d \int_{\left[\frac{\ell-1}{r},\frac{\ell}{r}\right]} \pi_{j,\kappa/r}f(x) \mathrm{d}x} \right)^2 = (\dagger).$$

For  $x \in [(\ell - 1)/r, \ell/r]$  we always have

$$|\pi_{j,\kappa/r}f(x) - \pi_{j,\kappa/r}f(\ell/r)| \le \omega(\frac{\sqrt{d}}{r};\pi_{j,\kappa/r}f) \lesssim \left(\frac{\sqrt{d}/r}{\kappa/r}\right)^{\beta} \lesssim \kappa^{-\beta}.$$

Using the inequality  $\sqrt{a+b} - \sqrt{(a-b)_+} \le 2\sqrt{b}$  valid for all  $a, b \ge 0$ , we can bound (†) by  $\kappa^{-\beta}$  up to constant and the result follows.

#### ACKNOWLEDGMENT

We thank Julien Chhor for pointing out the interpretation (IV.2). PG was supported in part by the NSF award IIS-1838071. YP was supported in part by the NSF under grant No CCF-2131115 and by the MIT-IBM Watson AI Lab.

#### REFERENCES

- [1] S. Chatrchyan, V. Khachatryan, A. M. Sirunyan, A. Tumasyan, W. Adam, E. Aguilo, T. Bergauer, M. Dragicevic, J. Erö, C. Fabjan *et al.*, "Observation of a new boson at a mass of 125 gev with the cms experiment at the lhc," *Physics Letters B*, vol. 716, no. 1, pp. 30–61, 2012.
- [2] C. Adam-Bourdarios, G. Cowan, C. Germain, I. Guyon, B. Kégl, and D. Rousseau, "The higgs boson machine learning challenge," in NIPS 2014 workshop on high-energy physics and machine learning. PMLR, 2015, pp. 19–55.
- [3] S. Agostinelli, J. Allison, K. a. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand et al., "Geant4—a simulation toolkit," *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250–303, 2003.
- [4] S. Frixione, P. Nason, and C. Oleari, "Matching nlo qcd computations with parton shower simulations: the powheg method," *Journal of High Energy Physics*, vol. 2007, no. 11, p. 070, 2007.
- [5] G. Corcella, I. G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M. H. Seymour, and B. R. Webber, "Herwig 6: an event generator for hadron emission reactions with interfering gluons (including supersymmetric processes)," *Journal of High Energy Physics*, vol. 2001, no. 01, p. 010, 2001.
- [6] T. Sjöstrand, S. Mrenna, and P. Skands, "Pythia 6.4 physics and manual," Journal of High Energy Physics, vol. 2006, no. 05, p. 026, 2006.
- [7] J. Alwall, P. Demin, S. De Visscher, R. Frederix, M. Herquet, F. Maltoni, T. Plehn, D. L. Rainwater, and T. Stelzer, "Madgraph/madevent v4: the new web generation," *Journal of High Energy Physics*, vol. 2007, no. 09, p. 028, 2007.
- [8] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, "Asymptotic formulae for likelihood-based tests of new physics," *The European Physical Journal C*, vol. 71, no. 2, pp. 1–19, 2011.
- [9] L. Lista, Statistical methods for data analysis in particle physics. Springer, 2017, vol. 941.
- [10] L. Devroye and G. Lugosi, Combinatorial methods in density estimation. Springer Science & Business Media, 2001.
- [11] Y. I. Ingster, "Minimax testing of nonparametric hypotheses on a distribution density in the l\_p metrics," *Theory of Probability & Its Applications*, vol. 31, no. 2, pp. 333–337, 1987.
- [12] O. Goldreich and D. Ron, "On testing expansion in bounded-degree graphs," 2000.
- [13] B. G. Kelly, T. Tularak, A. B. Wagner, and P. Viswanath, "Universal hypothesis testing in the learning-limited regime," in 2010 IEEE International Symposium on Information Theory. IEEE, 2010, pp. 1478–1482.
- [14] J. Ziv, "On classification with empirically observed statistics and universal data compression," *IEEE Trans. Inf. Theory*, vol. 34, pp. 278–286, 1988.
- [15] M. Gutman, "Asymptotically optimal classification for multiple tests with empirically observed statistics," *IEEE Transactions on Information Theory*, vol. 35, no. 2, pp. 401–408, 1989.
- [16] L. Zhou, V. Y. Tan, and M. Motani, "Second-order asymptotically optimal statistical classification," Information and Inference: A Journal of the IMA, vol. 9, no. 1, pp. 81–111, 2020.
- [17] H.-W. Hsu and I.-H. Wang, "On binary statistical classification from mismatched empirically observed statistics," in 2020 IEEE International Symposium on Information Theory (ISIT). IEEE, 2020, pp. 2533–2538.
- [18] H. He, L. Zhou, and V. Y. Tan, "Distributed detection with empirically observed statistics," *IEEE Transactions on Information Theory*, vol. 66, no. 7, pp. 4349–4367, 2020.
- [19] M. Haghifam, V. Y. Tan, and A. Khisti, "Sequential classification with empirically observed statistics," *IEEE Transactions on Information Theory*, vol. 67, no. 5, pp. 3095–3113, 2021.

- [20] P. Boroumand et al., "Universal neyman-pearson classification with a known hypothesis," arXiv preprint arXiv:2206.11700, 2022.
- [21] B. G. Kelly, A. B. Wagner, T. Tularak, and P. Viswanath, "Classification of homogeneous data with large alphabets," IEEE transactions on information theory, vol. 59, no. 2, pp. 782-795, 2012.
- [22] D. Huang and S. Meyn, "Classification with high-dimensional sparse samples," in 2012 IEEE International Symposium on Information Theory Proceedings. IEEE, 2012, pp. 2586-2590.
- B. Bhattacharya and G. Valiant, "Testing closeness with unequal sized samples," Advances in Neural Information Processing Systems, [23] vol. 28, 2015.
- [24] I. Diakonikolas and D. M. Kane, "A new approach for testing properties of discrete distributions," in 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS). IEEE, 2016, pp. 685-694.
- [25] L. Devroye, L. Gyorfi, and G. Lugosi, "A note on robust hypothesis testing," IEEE Transactions on Information Theory, vol. 48, no. 7, pp. 2111-2114, 2002.
- [26] J. Friedman, "On multivariate goodness-of-fit and two-sample testing," Citeseer, Tech. Rep., 2004.
  [27] D. Lopez-Paz and M. Oquab, "Revisiting classifier two-sample tests," *arXiv preprint arXiv:1610.06545*, 2016.
- [28] M. U. Gutmann, R. Dutta, S. Kaski, and J. Corander, "Likelihood-free inference via classification," Statistics and Computing, vol. 28, no. 2, pp. 411-425, 2018.
- [29] F. Liu, W. Xu, J. Lu, G. Zhang, A. Gretton, and D. J. Sutherland, "Learning deep kernels for non-parametric two-sample tests," in International conference on machine learning. PMLR, 2020, pp. 6316-6326.
- [30] I. Kim, A. Ramdas, A. Singh, and L. Wasserman, "Classification accuracy as a proxy for two-sample testing," The Annals of Statistics, vol. 49, no. 1, pp. 411-434, 2021.
- [31] S. Hediger, L. Michel, and J. Näf, "On the use of random forest for two-sample testing," Computational Statistics & Data Analysis, vol. 170, p. 107435, 2022.
- [32] P. R. Gerber, Y. Han, and Y. Polyanskiy, "Minimax optimal testing by classification," in The Thirty Sixth Annual Conference on Learning Theory. PMLR, 2023, pp. 5395–5432.
- [33] Y. I. Ingster, "On the minimax nonparametric detection of signals in white gaussian noise," Problemy Peredachi Informatsii, vol. 18, no. 2, pp. 61-73, 1982.
- [34] O. Goldreich, Introduction to property testing. Cambridge University Press, 2017.
- [35] N. Dalmasso, R. Izbicki, and A. Lee, "Confidence sets and hypothesis testing in a likelihood-free inference setting," in International Conference on Machine Learning. PMLR, 2020, pp. 2323-2334.
- [36] P. J. Huber, "A robust version of the probability ratio test," The Annals of Mathematical Statistics, pp. 1753–1758, 1965.
- [37] P. J. Huber and V. Strassen, "Minimax tests and the neyman-pearson lemma for capacities," The Annals of Statistics, pp. 251-263, 1973.
- [38] Y. G. Yatracos, "Rates of convergence of minimum distance estimators and kolmogorov's entropy," The Annals of Statistics, vol. 13, no. 2, pp. 768-774, 1985.
- [39] L. Birgé, Sur un théoreme de minimax et son application aux tests. Univ. de Paris-Sud, Dép. de Mathématique, 1979.
- [40] L. Birgé et al., "Robust tests for model selection," From probability to statistics and back: high-dimensional models and processes-A Festschrift in honor of Jon A. Wellner, pp. 47-64, 2013.
- [41] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White, "Testing closeness of discrete distributions," Journal of the ACM (JACM), vol. 60, no. 1, pp. 1-25, 2013.
- [42] E. Arias-Castro, B. Pelletier, and V. Saligrama, "Remember the curse of dimensionality: The case of goodness-of-fit testing in arbitrary dimension," Journal of Nonparametric Statistics, vol. 30, no. 2, pp. 448-471, 2018.
- [43] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White, "Testing that distributions are close," in Proceedings 41st Annual Symposium on Foundations of Computer Science. IEEE, 2000, pp. 259-269.
- [44] S.-O. Chan, I. Diakonikolas, P. Valiant, and G. Valiant, "Optimal algorithms for testing closeness of discrete distributions," in Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms. SIAM, 2014, pp. 1193–1203.
- [45] G. Valiant and P. Valiant, "An automatic inequality prover and instance optimal identity testing," SIAM Journal on Computing, vol. 46, no. 1, pp. 429-455, 2017.
- [46] T. Li and M. Yuan, "On the optimality of gaussian kernel based nonparametric tests against smooth alternatives," arXiv preprint arXiv:1909.03302, 2019.
- [47] I. A. Ibragimov and R. Z. Khas'minskii, "On the estimation of an infinite-dimensional parameter in gaussian white noise," in Doklady Akademii Nauk, vol. 236, no. 5. Russian Academy of Sciences, 1977, pp. 1053-1055.
- [48] A. B. Tsybakov, Introduction to Nonparametric Estimation, 1st ed. Springer Publishing Company, Incorporated, 2008.
- [49] I. M. Johnstone, "Gaussian estimation: Sequence and wavelet models," 2019.
- [50] E. Giné and R. Nickl, Mathematical foundations of infinite-dimensional statistical models. Cambridge university press, 2021.
- [51] C. L. Canonne, "A short note on learning discrete distributions," arXiv preprint arXiv:2002.11457, 2020.
- [52] B. Nazer, O. Ordentlich, and Y. Polyanskiy, "Information-distilling quantizers," in 2017 IEEE International Symposium on Information Theory (ISIT). IEEE, 2017, pp. 96-100.
- [53] A. Pensia, V. Jog, and P.-L. Loh, "Communication-constrained hypothesis testing: Optimality, robustness, and reverse data processing inequalities," arXiv preprint arXiv:2206.02765, 2022.
- [54] L. Devroye, A. Mehrabian, and T. Reddad, "The total variation distance between high-dimensional gaussians," arXiv preprint arXiv:1810.08693, 2018.
- [55] G. C. Kamath, "Modern challenges in distribution testing," Ph.D. dissertation, Massachusetts Institute of Technology, 2018.
- [56] C. Daskalakis, G. C. Kamath, and J. Wright, "Which distribution distances are sublinearly testable?" in Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM, 2018, pp. 2747–2764.
- [57] L. Birgé, "On estimating a density using hellinger distance and some other strange facts," Probability theory and related fields, vol. 71, no. 2, pp. 271-291, 1986.

- [58] L. Paninski, "A coincidence-based test for uniformity given very sparsely sampled discrete data," *IEEE Transactions on Information Theory*, vol. 54, no. 10, pp. 4750–4755, 2008.
- [59] P. Valiant, "Testing symmetric properties of distributions," *SIAM Journal on Computing*, vol. 40, no. 6, pp. 1927–1968, 2011.
- [60] S. Balakrishnan and L. Wasserman, "Hypothesis testing for densities and high-dimensional multinomials: Sharp local minimax rates," The
- Annals of Statistics, vol. 47, no. 4, pp. 1893-1927, 2019.
- [61] J. Chhor and A. Carpentier, "Sharp local minimax rates for goodness-of-fit testing in large random graphs, multivariate poisson families and multinomials," *arXiv preprint arXiv:2012.13766*, 2020.
- [62] —, "Goodness-of-fit testing for hölder-continuous densities: Sharp local minimax rates," arXiv preprint arXiv:2109.04346, 2021.
- [63] J. Lam-Weil, A. Carpentier, and B. K. Sriperumbudur, "Local minimax rates for closeness testing of discrete distributions," *Bernoulli*, vol. 28, no. 2, pp. 1179–1197, 2022.
- [64] J. Acharya, H. Das, A. Jafarpour, A. Orlitsky, S. Pan, and A. Suresh, "Competitive classification and closeness testing," in *Conference on Learning Theory*. JMLR Workshop and Conference Proceedings, 2012, pp. 22–1.
- [65] Y. I. Ingster and I. A. Suslina, Nonparametric goodness-of-fit testing under Gaussian models. Springer Science & Business Media, 2003, vol. 169.
- [66] D. Newman and H. Shapiro, "Jackson's theorem in higher dimensions," in On Approximation Theory/Über Approximationstheorie. Springer, 1964, pp. 208–219.