

It Was “All” for “Nothing”: Sharp Phase Transitions for Noiseless Discrete Channels

Jonathan Niles-Weed and Ilias Zadik[✉]

Abstract—We establish a phase transition known as the “all-or-nothing” phenomenon for noiseless discrete channels. This class of models includes the Bernoulli group testing model and the planted Gaussian perceptron model. Previously, the existence of the all-or-nothing phenomenon for such models was only known in a limited range of parameters. Our work extends the results to all signals with arbitrary sublinear sparsity. Over the past several years, the all-or-nothing phenomenon has been established in various models as an outcome of two seemingly disjoint results: one positive result establishing the “all” half of all-or-nothing, and one impossibility result establishing the “nothing” half. Our main technique in the present work is to show that for noiseless discrete channels, the “all” half implies the “nothing” half, that is, a proof of “all” can be turned into a proof of “nothing.” Since the “all” half can often be proven by straightforward means—for instance, by the first-moment method—our equivalence gives a powerful and general approach towards establishing the existence of this phenomenon in other contexts.

Index Terms—All-or-nothing phenomenon, discrete channels, group testing, Gaussian perceptron.

I. INTRODUCTION

ASURPRISING feature of high-dimensional inference problems is the presence of *phase transitions*, where the behavior of estimator changes abruptly as the parameters of a problem vary. Often, these transitions help illuminate fundamental limitations of an optimal estimation procedure, by showing, for instance, that a certain inference task is impossible when the noise is too large or the number of samples too few. There is a large and growing literature on proving rigorously the presence of such transitions and on establishing their implications for learning and inference tasks in a variety of settings. (see, e.g. [15])

A particularly striking phase transition is known as the *all-or-nothing* phenomenon [10], [22], [26]. In problems evincing this phenomenon, there is a sharp break: below a critical number of samples, it is impossible to infer *almost any*

Manuscript received 4 March 2022; revised 31 August 2022; accepted 14 November 2022. Date of publication 11 January 2023; date of current version 14 July 2023. The work of Jonathan Niles-Weed was supported by the NSF under Grant DMS-2015291. The work of Ilias Zadik was supported in part by a CDS Moore-Sloan (postdoctoral) fellowship from NYU, in part by the Simons-NSF under Grant DMS-2031883, and in part by the Vannevar Bush Faculty Fellowship under Grant ONR-N00014-20-1-2826. (Corresponding author: Ilias Zadik.)

Jonathan Niles-Weed is with the Courant Institute of Mathematical Sciences and Center for Data Science, New York University, New York City, NY 10012 USA (e-mail: jnw@cims.nyu.edu).

Ilias Zadik is with the Department of Mathematics, Massachusetts Institute of Technology, Boston, MA 02142 USA (e-mail: ilias.zadik@gmail.com).

Communicated by S. Boucheron, Associate Editor for Machine Learning and Statistics.

Digital Object Identifier 10.1109/TIT.2022.3225802

information about a parameter of interest, but as soon as that critical point is reached, it is possible to infer the parameter almost perfectly. Such problems exhibit a sharp dichotomy, where either perfect inference is possible or nothing is.

In this work, we develop general tools for proving the all-or-nothing phenomenon for a class of models we call “noiseless discrete channels.” In such models, we fix a function g and observe identically distributed copies of a pair $(Y, X) \in \mathcal{Y} \times \mathbb{R}^L$ generated by

$$Y = g(X, \theta),$$

where X is a random draw from some known distribution on \mathbb{R}^L , and θ is an unknown parameter to be estimated. Under the assumption that $|\mathcal{Y}| < \infty$, we can view g as a discrete channel, parametrized by θ , which maps \mathbb{R}^L to \mathcal{Y} , and our goal is to ascertain how many samples (i.e., how many uses of this channel) we need to reliably recover θ .

We highlight two special cases of the above model which have seen recent attention:

- Group testing [1], [9]: $\theta \in \{0, 1\}^N$ indicates a subset of infected individuals in a population, and $X \in \{0, 1\}^N$ indicates a random subset chosen to be tested as a batch. We observe $g(X, \theta) = 1(\text{Support}(X) \cap \text{Support}(\theta) \neq \emptyset)$, where for a vector $v \in \mathbb{R}^N$, $\text{Support}(v) \subseteq [N]$ denotes the set of the non-zero coordinates of v . How many tests do we need to determine which individuals are infected?
- Planted Gaussian perceptron [27]: in this simple “teacher-student” setting, $\theta \in \{0, 1\}^N$ represents the weights of a “teacher” one-layer neural network, and we observe $g(X, \theta) = 1(\sum_{j=1}^N \theta_j x_j \geq 0)$, where the x_j are i.i.d. standard Gaussian random variables. How many samples do we need for a “student” to learn the teacher’s hidden weights?

Both models have recently been studied in the all-or-nothing framework [14], [25]. However, the range of parameters for which the all-or-nothing phenomenon has been rigorously established in either model is limited. [25] show that all-or-nothing holds for group testing in the extremely sparse regime when the number of infected individuals is $o(N^\varepsilon)$ for all $\varepsilon > 0$. Their proof is combinatorial and proceeds by the second-moment method. [14] give a heuristic derivation of the all-or-nothing phenomenon for the planted Gaussian perceptron based on the replica method from statistical physics, and establish that this phenomenon holds if $\|\theta\|_0 := |\{i \in [N] : \theta_i \neq 0\}|$ is both $\omega(N^{\frac{8}{9}})$ and $o(N)$.

A. Contribution

We give a simple criterion for the all-or-nothing phenomenon to hold in noiseless discrete channels. For such settings, we measure success in terms of the minimum mean squared error (MMSE) and the signal is assumed to lie on the Euclidean unit sphere. The “*all*” property corresponds to a vanishing MMSE, while the “*nothing*” property corresponds to MMSE being asymptotically equal to one, which is the mean squared error achieved by the trivial zero estimator. As a corollary of our result, we show that the all-or-nothing phenomenon holds for all relevant sparsity regimes in both the group testing and planted perceptron models, substantially generalizing prior work.

Our key technical contribution is to show that, under suitable conditions, proving the “*all*” condition *immediately* implies that the “*nothing*” condition holds as well. More specifically, we show that if the mean squared error vanishes for all $n \geq (1 + \epsilon)n^*$ for some critical n^* , then for $n \leq (1 - \epsilon)n^*$ no recovery is possible. In other words, for these models, “*all*” implies “*nothing*” in a suitable sense. Crucially, the “*all*” condition can often be proven directly, by simple means, as it suffices to establish that a specific estimator is successful, via for example a simple “union bound” or “first-moment” argument. On the other hand, the “*nothing*” lower bound requires proving the failure of any estimation method, and has typically been proven by using more subtle techniques, such as delicate second moment method arguments (see e.g. [22] for the regression setting and [25] for the Bernoulli group testing setting). Our “*all*” implies “*nothing*” result shows that this complication is unnecessary for a class of noiseless discrete channels.

We apply our techniques to both non-adaptive Bernoulli group testing and the planted Gaussian perceptron model. We report the following.

- For the *Bernoulli group testing model* (BGT), we focus on the case, common in the group testing literature, where there are k infected individuals, with $k = o(N)$. We model the infected individuals as a binary k -sparse vector on the unit sphere, and as mentioned above we measure success in terms of the MMSE. In the BGT setting each individual is assumed to participate in any given test in an i.i.d. fashion, and independently with everything, with probability $\frac{\nu}{k}$, for some $\nu = \nu_k$ satisfying $q = (1 - \frac{\nu}{k})^k$. Here $q \in (0, 1)$ is a fixed constant, again as customary in the literature of Bernoulli group testing [1]. We show as an application of our technique that the all-or-nothing phenomenon holds for the BGT design **for all** $k = o(N)$ **and for any** $q \leq \frac{1}{2}$ at the critical number of tests

$$n_q^* = k \log \frac{N}{k} / h(q),$$

where $h(q)$ denotes the (rescaled) binary entropy at q defined in (18). In words, with less than n_q samples the MMSE is not better than “random guess”, while with more than n_q samples it is almost zero. To the best of our knowledge this result was known before only in the case where $k = o(N^\varepsilon)$ for all $\varepsilon > 0$ and $q = \frac{1}{2}$ [25].

- For the *Gaussian perceptron model*, we focus on the case where θ is a binary k -sparse vector on the unit sphere, with $k = o(N)$. We study a more general class of noiseless Boolean models than the Gaussian perceptron, where $Y_i = 1(\langle X_i, \theta \rangle \in A)$ for some arbitrary Borel $A \subseteq \mathbb{R}$ with (standard) Gaussian mass equal to $\frac{1}{2}$. Equivalently we consider any Boolean function $f: \mathbb{R} \rightarrow \{-1, 1\}$ which is balanced under the standard Gaussian measure, i.e. $\mathbb{E}f(Z) = 0$, $Z \sim N(0, 1)$, and assume $Y_i = f(\langle X_i, \theta \rangle)$. Notice that the perceptron model corresponds to the case $A = [0, +\infty)$ and $f(t) = 21(t > 0) - 1$, but it includes other interesting models such as the symmetric binary perceptron $A = [-u, u]$ with u the median of $|Z|$, $Z \sim N(0, 1)$ which has recently been studied in the statistical physics literature [2]. We apply our technique in this setting to prove a generic result; **all such models** exhibit the all-or-nothing phenomenon at the same critical sample size

$$n^* = k \log_2 \frac{N}{k}.$$

To the best of our knowledge this sharp phase transition was known before only in the case where $A = [0, +\infty)$ and k is $\omega(N^{\frac{2}{3}})$ and $o(N)$ [14]

B. Comparison With Previous Work

1) *All-or-Nothing*: The all-or-nothing phenomenon has been investigated in a variety of models, and with different techniques [4], [5], [10], [14], [17], [21], [22], [25]. More specifically, the phenomenon was initially observed in the context of the maximum likelihood estimator for sparse regression in [10] and was later established in the context of MMSE for sparse regression [21], [22], sparse (tensor) PCA [4], [5], [17], Bernoulli group testing [25] and generalized linear models [14].

A common theme of these works is that all-or-nothing behavior can arise when the parameter to be recovered is sparse, with sparsity *sublinear* in the dimensions of the problem. Though it is expected that this phenomenon should arise for all sublinear scalings, technical difficulties often restrict the range of applicability of rigorous results. In the present work we circumvent this challenge by showing that a version of the “*all*” condition suffices to establish the all-or-nothing phenomenon for the whole sublinear regime. As mentioned above, usually the “*all*” result is easier to establish than the “*nothing*” result. Leveraging this, we are able to establish the all-or-nothing phase transitions throughout the sublinear sparsity regimes of both the Bernoulli group testing and Gaussian perceptron models, where only partial results have been established before [14], [25].

2) *“All” Implies “Nothing”*: As mentioned already, our key technical contribution is showing that the “*all*” result suffices to establish the all-or-nothing phenomenon. This potentially counterintuitive result relates to a technique used in information theory known as the *area theorem* [13], [16], [20]. A heuristic explanation of this connection in the regression context appears in [23, Section 1.1.]; however, despite this intuition, the authors of [22] do not proceed by this route.

To the best of our knowledge, our results are the first to rigorously prove that in certain sparse learning settings, the “all” result indeed implies the all-or-nothing sparse phase transition.

3) *Bernoulli Group Testing*: Now, we comment on our contribution for the BGT model, as compared to the BGT literature. In the Bernoulli group testing model, it is well-known that for all $k = o(N)$ and $q = (1 - \frac{\nu}{k})^k$, it is possible to obtain a vanishing MMSE (“all”) with access to $(1 + \epsilon)n_q^* = (1 + \epsilon)k \log \frac{N}{k} / h(q)$ tests [1], [25]. Furthermore, it is also known that if $q = 1/2$ with less than $(1 - \epsilon)n_{1/2}^*$ test it is impossible to achieve an “all” result [24, Theorem 3]. To the best of our knowledge, this (weak) negative result of “all” being impossible is *not known* when $q \neq \frac{1}{2}$ and one has access to fewer than $(1 - \epsilon)n_q^*$ tests, though some relevant discussion appears in [23]. Finally, as mentioned above, [25] do establish the strong negative “nothing” result that if $k = o(N^\delta)$ for all $\delta > 0$ and $q = \frac{1}{2}$ with less $(1 - \epsilon)n_{1/2}^*$ it is impossible to achieve a non-trivial MMSE [25].

In the present work, we show as a corollary of our methods that **for all** $k = o(N)$ and $q \leq \frac{1}{2}$, “nothing” holds when the number of tests is fewer than $(1 - \epsilon)n_q^*$, substantially improving the literature of impossibility results in Bernoulli group testing. While to the best of our knowledge, the appropriate “all” result needed for our argument to work is not known for any $q < \frac{1}{2}$ it has been established before when $q = \frac{1}{2}$ [see, e.g.] [Lemma 1.3.] [12]. Finally, it is worth pointing out that some form of non-trivial information can still be extracted from the Bernoulli group testing instance even in the “nothing” regime where the MMSE is trivial. For example, [25] showed that for some values of k it is possible even when $n < (1 - \epsilon)n_q^*$ to successfully hypothesis test between the Bernoulli group testing model and a “pure noise” model where the tests outcomes are random and independent from everything else (see also the more recent work [8] on the same topic).

4) *Gaussian Perceptron Model*: For the Gaussian perceptron model, to the best of our knowledge the most relevant result is in [14] where the authors prove the all-or-nothing phenomenon at $n^* = k \log_2 \frac{N}{k}$ samples when k is $\omega(N^{\frac{8}{9}})$ and $o(N)$. While they characterize the free energy of the model and therefore provide more precise results than we do, their results apply to a restricted sparsity regime. We do not precisely characterize the limiting free energy, but our much simpler argument shows that the all-or-nothing phenomenon holds **for all** sparsity levels $k = o(N)$.

II. MAIN RESULTS

A. General Framework: Noiseless Discrete Channels

1) *The Family of Models*: We define a sequence of observational models we study in this work, indexed by $N \in \mathbb{N}$. Assume that an unknown parameter, or “signal”, $\theta \in \mathbb{R}^N$ is drawn from some uniform prior $P_\Theta = (P_\Theta)_N$ supported on a discrete subset Θ of the unit sphere in \mathbb{R}^N . We set $|\Theta| = M = M_N$ and make the following “non-negativity” assumption on the overlap between two parameters that for any $\theta, \theta' \in \Theta$ it holds

$$\langle \theta, \theta' \rangle \geq 0.$$

For some distribution $\mathcal{D} = \mathcal{D}_N$ supported on \mathbb{R}^L , where $L = L_N$, we assume that for $n = n_N$ i.i.d. samples $X_i \sim \mathcal{D}_X, i = 1, 2, \dots, n$ we observe $(Y_i, X_i), i = 1, 2, \dots, n$ where

$$Y_i = g(X_i, \theta), i = 1, 2, \dots, n. \quad (1)$$

The function $g = g_N : \mathbb{R}^L \times \mathbb{R}^N \rightarrow \mathcal{Y}$ is referred to as the channel. We assume throughout that \mathcal{Y} is finite and of cardinality that remains constant as N grows, e.g., $\mathcal{Y} = \{0, 1\}$. We denote by Y^n the n -dimensional vector with entries $Y_i, i = 1, 2, \dots, n$ and X^n the $n \times L$ matrix with columns the vectors $X_i, i = 1, 2, \dots, n$. We write $P = P_N$ for the joint law of (Y^n, X^n, θ) .

We are given access to the pair (Y^n, X^n) , and our goal is to recover θ . We measure recovery with n samples in terms of the minimum mean squared error (MMSE),

$$\text{MMSE}_N(n) = \mathbb{E}\|\theta - \mathbb{E}[\theta|Y^n, X^n]\|^2. \quad (2)$$

2) *The All-or-Nothing Phenomenon*: We say that a sequence of models $((P_\Theta)_N, g_N, \mathcal{D}_N)$ satisfies the *all-or-nothing phenomenon* with critical sequence of sample sizes $n_c = (n_c)_N$ if

$$\lim_{N \rightarrow \infty} \text{MMSE}_N(\lfloor \beta n_c \rfloor) = \begin{cases} 1 & \text{if } \beta < 1 \\ 0 & \text{if } \beta > 1. \end{cases} \quad (3)$$

This condition expresses a very sharp phase transition: when $\beta > 1$, we can identify the signal nearly perfectly, but when $\beta < 1$, we can do no better than a trivial estimator which always outputs zero.

3) *Assumptions*: To establish our result we make throughout the following further assumptions on our models.

Recall that we have assumed that our prior P_Θ is the uniform distribution on some finite subset of cardinality $M = M_N$. We assume throughout that $M_N \rightarrow \infty$ as $N \rightarrow \infty$. We also make the following assumption, which requires that the distribution P_Θ is sufficiently spread out.

Assumption 1: For θ and θ' chosen independently from P_Θ we have

$$\lim_{\delta \rightarrow 0^+} \lim_{N \rightarrow \infty} \frac{\log(M P_\Theta^{\otimes 2}(\langle \theta, \theta' \rangle \geq 1 - \delta))}{\log M} = 0. \quad (4)$$

Moreover, we assume for θ and θ' chosen independently from P_Θ and any $\epsilon > 0$,

$$\lim_N P_\Theta^{\otimes 2}(\langle \theta', \theta \rangle \geq \epsilon) = 0. \quad (5)$$

Assumption (4) guarantees that the that for two independent draws from the prior θ, θ' , the asymptotic probability that θ' is very near to θ is dominated by the probability that $\theta = \theta'$. This condition is the same as the one employed by [17] in the analysis of the all-or-nothing phenomenon for Gaussian models.

Assumption (5) implies that independent samples from the prior are asymptotically uncorrelated with each other. This condition is natural in the context of the all-or-nothing phenomenon, since if Assumption (5) fails to hold, then it is possible to obtain an estimator with non-trivial correlation with the signal by simply drawing a fresh sample from the prior, independent of the observations.

Assumptions (4) and (5) are easy to verify in a variety of sparse models. For instance, they hold if the rate function

$$r(\rho) = -\lim_N \frac{1}{\log M} \log P_{\Theta}^{\otimes 2}(\langle \theta', \theta \rangle \geq \rho) \quad \rho \in [0, 1]$$

exists and is a strictly increasing continuous function on $[0, 1]$.

We make also assumptions on the probability a $\theta' \in \Theta \setminus \{\theta\}$ is able to fit the observations generated by the signal θ .

Assumption 2: We assume there exists a fixed function $R : [0, 1] \rightarrow [0, 1]$, independent of N , such that

$$P_N(g(X, \theta) = g(X, \theta')) = R(\langle \theta, \theta' \rangle) \quad \forall N \in \mathbb{N}, \theta, \theta' \in \Theta.$$

That is, that the probability that $g(X, \theta)$ and $g(X, \theta')$ agree is a function of $\langle \theta, \theta' \rangle$ alone. We assume that R is continuous at 0^+ and strictly increasing on $[0, 1]$.

B. Main Result: How “All” Implies “Nothing”

Notice that since our prior distribution is a uniform distribution over the finite parameter space Θ and our observation model is noiseless, the posterior distribution of θ given Y^n, X^n satisfies that for any θ' ,

$$\begin{aligned} P(\theta' | Y^n, X^n) &= \frac{P(\theta') P(Y^n | X^n, \theta')}{P(Y^n | X^n)} \\ &\propto P(Y^n | X^n, \theta') = \prod_{i=1}^n 1(Y_i = g(X_i, \theta')). \end{aligned}$$

In words, the posterior distribution is simply the uniform measure over the vectors $\theta' \in \Theta$ satisfying

$$Y_i = g(X_i, \theta'), i = 1, 2, \dots, n. \quad (6)$$

As an easy corollary, the distance of the posterior mean from the ground truth vector, or equivalently the $\text{MMSE}_N(n)$, can be naturally related to the behavior of the following “counting” random variables.

Definition 1: For any $N \in \mathbb{N}$ and $\delta \in [0, 2]$, let $Z_{N,\delta} = Z_{N,\delta}(Y^n, X^n)$ be the random variable which is equal to the number of solutions $\theta' \in \Theta$ of equations (6) where $\|\theta - \theta'\|^2 \geq \delta$.

Using the definition above, the following simple proposition holds.

Proposition 1: For θ' drawn from the posterior distribution of θ given Y^n, X^n it holds almost surely that

$$P(\|\theta - \theta'\|^2 \geq \delta | Y^n, X^n) = \frac{Z_{N,\delta}}{Z_{N,0}}. \quad (7)$$

Hence,

$$\text{MMSE}_N(n) = \frac{1}{2} \mathbb{E} \int_{\delta=0}^2 \frac{Z_{N,\delta}}{Z_{N,0}} d\delta. \quad (8)$$

Furthermore, the property that

$$\text{for all } \delta \in (0, 2], \lim_N \mathbb{E} \frac{Z_{N,\delta}}{Z_{N,0}} = 0 \quad (9)$$

is equivalent with the “all” property

$$\lim_N \text{MMSE}_N(n) = 0. \quad (10)$$

Finally the property

$$\text{for all } \delta \in (0, 2], \lim_N P(Z_{N,\delta} > 0) = 0, \quad (11)$$

implies the “all” condition (10).

Proposition 1 offers a clean combinatorial way of establishing the vanishing MMSE (“all”) in our context; one needs to prove the (relative) absence of solutions of (6) which are at a constant distance from θ , establishing for example (11). A clear benefit of such an approach is that one could possibly establish such a result by trying a (possibly conditional) union bound—or “first moment”—argument. We investigate further the power of establishing the “all” result in what follows.

We consider the following *critical sample size*,

$$n^* = (n^*)_N = \lfloor \frac{H(\theta)}{H(Y)} \rfloor, \quad (12)$$

where by $H(\cdot)$ we refer to the Shannon entropy of a discrete random variable and $Y = g(X, \theta)$ for a sample of $X \sim \mathcal{D}_N$ and $\theta \sim (P_{\Theta})_N$. The significance of the sample size n^* is highlighted in the following proposition which establishes that the “all” condition (10) can only hold if the number of samples is at least n^* .

Proposition 2: Suppose that Assumption 1 is true. If the “all” condition (10) holds for some sequence of sample sizes $n = n_N$, then

$$\liminf_N \frac{n}{n^*} \geq 1.$$

While we defer the proof of Proposition 2 to the Appendix VI, we highlight some aspects of it which will be important in what follows. The key identity behind the proof of the proposition is it always holds that

$$H(\theta) - H(\theta | Y^n, X^n) \quad (13)$$

$$= nH(Y) - D(P(Y^n, X^n) || Q(Y^n, X^n)) \quad (14)$$

$$\leq nH(Y), \quad (15)$$

where 1) D stands for the Kullback-Leibler (KL) divergence (see e.g. [20, Definition 2.1.]), 2) $P(Y^n, X^n)$ stands for the joint law of (Y^n, X^n) generated by the observation model (1) and 3) $Q(Y^n, X^n)$ stands for the law of a “null” model where the columns of X^n are i.i.d. samples drawn from \mathcal{D} and the entries of Y^n are drawn in an i.i.d. fashion from the distribution of $Y = g(X, \theta)$ but independently from X^n . As a result, the law of a single observation (X_i, Y_i) is the same under P and Q , but P and Q are distinct as joint distributions, as for example the latter has *no hidden signal*. The identity (14) follows from algebraic manipulations which can be found in Appendix II. The inequality in (15) is implied by the non-negativity of the KL divergence.

The proof of the proposition is based on the fact that the “all” condition (10) implies that the entropy of the posterior is of smaller order of magnitude than the entropy of the prior (see Proposition 4). This property allows us to conclude that the left hand side of (15) is $(1 - o(1))H(\theta)$ which concludes the proof.

Now we present the main technical result of the present work. We establish that if Proposition 2 is tight, that is if (10) can be proven to be true when $n \geq (1 + \epsilon)n^*$ for arbitrary

$\epsilon > 0$, then (10) is a sufficient to establish that the *the all-or-nothing phenomenon* holds at sample size n^* as well.

Theorem 1 (“All” Implies “Nothing”): Suppose that Assumptions 1, 2 are true. Assume that if $n \geq (1 + \epsilon)n^*$, for some arbitrary fixed $\epsilon > 0$, then the “all” condition (10) holds. Then if $n \leq (1 - \epsilon)n^*$ for arbitrary fixed $\epsilon \in (0, 1)$, it holds

$$\lim_N \text{MMSE}_N(n) = 1. \quad (16)$$

In particular, the all-or-nothing phenomenon (3) holds at critical samples sizes $n_c = n^*$.

We provide here some intuition behind such a potentially surprising implication. Notice that if (10) holds at sample sizes $(1 + \epsilon)n^*$ for arbitrary fixed $\epsilon > 0$, then from the sketch of the proof of Proposition 2 the inequality (15) needs to hold (approximately) with equality. In fact one can show that at $n = n^*$, it must necessarily hold that

$$\lim_N \frac{D(P(Y^{n^*}, X^{n^*}) || Q(Y^{n^*}, X^{n^*}))}{H(\theta)} = 0. \quad (17)$$

At an intuitive level, (17) seems already a significant step towards what we desire to prove. Indeed, (17) suggests that (Y^n, X^n) drawn from our model P are close in distribution to the samples (Y^n, X^n) drawn from the null model Q . This strongly suggests that outperforming the random guess in mean-squared error should be impossible.

While we think that this argument hints at the right direction for proving the “nothing” property, we do not know a complete proof along these lines. The reason is that one cannot conclude that P and Q are “sufficiently close”, e.g., in the total variation sense, to argue the above. The reason is that the KL distance in (17) vanishes only after rescaling by the factor $H(\theta) = \log M_N \rightarrow +\infty$. For this reason, (17) does not imply any nontrivial bound for the total variation distance between P, Q . Notably though, such an obstacle has already been tackled in the literature of “nothing” results in the context of sparse tensor PCA [6], [17]. In these cases the “nothing” result can be established by the use of the I-MMSE formula combined with weak detection lower bound such as (17). The I-MMSE formula is an identity for Gaussian channels between the derivative (with respect to the continuous *signal to noise ratio (SNR)*) of the corresponding KL divergence and the MMSE for this value of SNR. We are not aware of any such formula for the noiseless discrete models considered in this work. Nevertheless, inspired by the I-MMSE connection, we study the discrete derivative of the KL divergence for noiseless models. Specifically we prove a result of potentially independent interest, that a vanishing *discrete derivative* (with respect to *the sample size n*) of the sequence $D(P(Y^n, X^n) || Q(Y^n, X^n)) / H(\theta)$, $n \in \mathbb{N}$, at $n \leq n^*$, implies indeed a trivial MMSE at $n < n^*$. We conclude then the result by using classical real analysis results, to show that the convexity of the vanishing sequence for $n \leq n^*$, implies its the discrete derivative of the sequence is also vanishing for $n \leq n^*$.

C. The Case of Boolean Channels: A Simple Condition

One can naturally ask whether for various models of interest there exist a simple sufficient condition which can establish the positive result (11) at n^* (e.g. by a union bound argument) and therefore prove the all-or-nothing phenomenon. In this subsection, we provide such a simple sufficient condition for the subclass of *Boolean (or 1-bit) noiseless models*, i.e. when $\mathcal{Y} = \{0, 1\}$. Perhaps not surprisingly, our result follows from an appropriate “union-bound” or “first-moment” argument. In the next section we apply our condition to various such models.

Notice that in these Boolean binary settings the critical sample size simplifies to $n^* = \lfloor \frac{H(\theta)}{h(p)} \rfloor$ where $p = P(g(X, \theta) = 1)$ and h is the binary entropy

$$h(t) = -t \log t - (1-t) \log(1-t), t \in (0, 1), \quad (18)$$

where \log is, as always in this work, with base e .

To proceed, we need some additional definitions. The first is about the two possible outcomes of the channel, and it extends Assumption 2 to further properties on the probability of a fixed $\theta' \in \Theta \setminus \{\theta\}$ satisfying (6).

Assumption 3: There exist fixed functions $R_1 : [0, 1] \rightarrow [0, 1]$, $R_0 : [0, 1] \rightarrow [0, 1]$, independent of N , such that for all $N \in \mathbb{N}$ and $\theta, \theta' \in \Theta$ it holds

$$R_1(\langle \theta, \theta' \rangle) = P(g(X, \theta) = g(X, \theta') = 1), \quad (19)$$

$$R_0(\langle \theta, \theta' \rangle) = P(g(X, \theta) = g(X, \theta') = 0). \quad (20)$$

For all $\rho \in [0, 1]$, $R(\rho) = R_0(\rho) + R_1(\rho)$, where $R(\rho)$ is as in Assumption 2. We assume that both R_0 and R_1 are increasing on $[0, 1]$.

The second is about the distribution of the overlap between two independent copies of the prior distribution. Again we borrow the definition from [17].

Definition 2: Given a non-decreasing continuous function $r : [-1, 1] \rightarrow \mathbb{R}_{\geq 0}$, we say $\{P_\Theta\}$ admits an *overlap rate function* r , if for all $\rho \in [0, 1]$ it holds

$$\limsup_N \frac{1}{\log M_N} \log P_\Theta^{\otimes 2}[\langle \theta', \theta \rangle \geq \rho] \leq -r(\rho),$$

where θ and θ' are independent draws from P_Θ .

We state our result.

Corollary 1: Let $\mathcal{Y} = \{0, 1\}$ and let $p = P(g(X, \theta) = 1) \in (0, 1)$ be constant. Suppose that Assumptions 1, 2 and 3 are true. If $\{P_\Theta\}$ admits an overlap rate $r(\rho)$ satisfying

$$r(\rho) \geq \frac{1}{h(p)} \left(p \log \frac{R_1(\rho)}{p^2} + (1-p) \log \frac{R_0(\rho)}{(1-p)^2} \right), \quad \forall \rho \in [0, 1], \quad (21)$$

then the all-or-nothing phenomenon holds at $n^* = \lfloor \frac{\log M}{h(p)} \rfloor$.

III. APPLICATIONS

In this subsection we use our results, and specifically Corollary 1, to establish the all-or-nothing phenomenon for various *sparse Boolean* models of interest.

A. Application 1: Nonadaptive Bernoulli Group Testing

We start with nonadaptive Bernoulli group testing. In this context, we fix some parameter $k = k_N \in \mathbb{N}$ with $k \rightarrow +\infty$ and $k = o(N)$, which corresponds to k infected individuals out of a population of cardinality N . We also fix a constant $q \in (0, 1/2]$, which controls the size of each group which is getting tested. The goal is to perform nonadaptive tests for a disease on different “groups” or subsets of the N individuals at a time. The logic is that by doing so we may be able to use fewer tests, say, from testing each individual separately, and still recover the infected individuals. Notice that such a Bernoulli group testing model is characterized by the two parameters q and $k = k_N$.

1) The Model: We assume a uniform prior P_Θ , which following our notation we encode as the uniform measure on the k -sparse binary vectors on the sphere in N -dimensions, i.e.

$$\Theta = \{\theta \in \{0, \frac{1}{\sqrt{k}}\}^N : \|\theta\|_0 = k\},$$

where there is a natural correspondence between the identities of the infected individuals and the support of the vectors $\theta \in \Theta$. For each k , we define $\nu = \nu_k$ to be the unique positive number satisfying

$$\left(1 - \frac{\nu}{k}\right)^k = q.$$

The group of individuals being tested is modeled by the binary vector $X \in \{0, 1\}^N$, with $\mathcal{D} = \text{Bernoulli}(\frac{\nu}{k})^{\otimes N}$. In words, we choose whether an individual participates at any given test independently from everything and with probability ν/k , with the parameter ν chosen so that the probability that each group contains no infected individuals is exactly q .

We model the channel by the step function at $\frac{1}{\sqrt{k}}$, i.e. $Y = g(X, \theta) = 1(\langle X_i, \theta \rangle \geq \frac{1}{\sqrt{k}})$, which simply outputs the information of whether at least one of the individuals in the selected group is infected (which is equivalent to $\sqrt{k}\langle X_i, \theta \rangle = |\text{Support}(X_i) \cap \text{Support}(\theta)| \geq 1$) or not. The sample size n corresponds to the number of tests conducted.

Corollary 1 when applied to this context establishes the following result.

Theorem 2: Let $q \in (0, \frac{1}{2}]$ be a constant. Suppose $k = o(N)$ and ν_k satisfies $(1 - \frac{\nu_k}{k})^k = q$. Then the non-adaptive Bernoulli group testing model satisfies the all-or-nothing phenomenon at number of tests $n^* = \lfloor \frac{\log(\frac{N}{k})}{h(q)} \rfloor = (1 + o(1)) \frac{k \log \frac{N}{k}}{h(q)}$.

The $q = 1/2$, $k = N^{o(1)}$ case of Theorem 2 was proved by [25]. Our theorem extends their result to all sublinear sparsities and all $q \leq 1/2$. In particular, we cover the commonly used choices of $q = 1/e - o(1)$, with $\nu = 1$, and $q = 1/2 - o(1)$ with $\nu = \ln 2$ (see e.g. [1, Section 2]). When $q > 1/2$, condition (21) fails, and whether a result similar to Theorem 2 holds in this regime remains open.

B. Application 2: Sparse Gaussian Perceptron and Sparse Balanced Gaussian (SBG) Models

In this subsection, we turn our study to a family of what we call as *Sparse Balanced Gaussian (SBG) models*. Every such model can be characterized by some sparsity parameter

$k = k_N = o(N)$ and a fixed “balanced” Borel subset $A \subseteq \mathbb{R}$ with $P(Z \in A) = 1/2$ for $Z \sim N(0, 1)$.

1) The Model: We assume as above that the signal θ is sampled from the uniform measure on the k -sparse binary vectors on the sphere in N -dimensions, i.e. $\Theta = \{\theta \in \{0, \frac{1}{\sqrt{k}}\}^N : \|\theta\|_0 = k\}$. We assume that the distribution for $X \in \mathbb{R}^N$ is given by the standard Gaussian measure $\mathcal{D} = N(0, I_N)$. Finally the channel is given by the formula $Y = g(X, \theta) = 1(\langle X_i, \theta \rangle \in A)$.

We highlight two models of this class that have been studied in different contexts.

- The case $A = [0, +\infty)$ corresponds to the well-studied *Gaussian perceptron* model with a sparse planted signal $Y = 1(\langle X, \theta \rangle \geq 0)$. Variants of the Gaussian perceptron model have received enormous attention in learning theory and statistical physics (see e.g. [3], [27] and references therein). Recently the sparse version has been studied by [14].
- The case $A = [-u, u]$ for some u with u such that $P(|Z| \leq u) = \frac{1}{2}$, which corresponds to what is known as the *symmetric binary perceptron* model with a sparse planted signal $Y = 1(|\langle X, \theta \rangle| \leq u)$ [2].

We establish a general result that all SBG models exhibit the all-or-nothing phenomenon at the same critical sample size.

Theorem 3: Suppose $k = o(N)$ and $A \subseteq \mathbb{R}$ be an arbitrary fixed Borel subset with $P(Z \in A) = 1/2$ for $Z \sim N(0, 1)$. Then the Sparse Balanced Gaussian model defined by k and A exhibits the all-or-nothing phenomenon at $n^* = \lfloor \frac{\log(\frac{N}{k})}{\log 2} \rfloor = (1 + o(1))k \log_2 \frac{N}{k}$.

In the context of Gaussian perceptron it has recently been proven [14] that the all-or-nothing phenomenon holds for any $\omega(N^{\frac{8}{9}}) = k = o(N)$. Theorem 3 generalizes this result to any $k = o(N)$, even constant. To our knowledge, the existence of such a transition for the symmetric binary perceptron and other SBG models is new.

APPENDIX I AUXILIARY RESULTS AND IMPORTANT PRELIMINARY CONCEPTS

All the results in this auxiliary section holds under our framework and assumptions as described in Section II-A.

We start with an elementary lemma.

Lemma 1: The sequence of critical sample sizes n^* defined in (12) satisfies,

- (i) $\lim_N n^* = +\infty$
- (ii) $n^* = (1 + o(1)) \frac{H(\theta)}{H(Y)}$, as $N \rightarrow +\infty$.

Proof: Recall that in all our models we assume $H(\theta) = \log M_N \rightarrow +\infty$ as $N \rightarrow +\infty$. Furthermore, we assume that $Y = g(X, \theta)$ is a random variable supported on a subset of \mathcal{Y} where $|\mathcal{Y}| = O(1)$. Hence, $H(Y) \leq \log |\mathcal{Y}| = O(1)$. In particular in all our models it holds

$$\lim_N \frac{H(\theta)}{H(Y)} = +\infty.$$

This establishes the first property. The second property follows since for any real valued sequence $X_N, N \in \mathbb{N}$ with $\lim_N x_N = +\infty$, it holds $\lim_N \frac{x_N}{\lfloor x_N \rfloor} = 1$. \square

We continue with a crucial proposition for our main result. This proposition establishes a connection between the critical sample size n^* , the estimation error manifested in the form of the entropy of the posterior distribution.

To properly establish it we need some additional notation, and the definition of an appropriate “null” distribution on the observation (Y^n, X^n) .

Definition 3: We denote by $P_n = P_n(Y^n, X^n)$ the law of the observation under our model, i.e. for any measurable A , $P_n(A) = \mathbb{E}_{\theta \sim P_\Theta} P((Y^n, X^n, \theta) \in A)$. In words, P_n generates (Y^n, X^n) by first sampling θ from the prior, independently sampling X^n in an i.i.d. fashion from \mathcal{D} , and then generating Y^n by the conditional law $P(Y^n | \theta, X^n)$.

Notice that in the noiseless case studied in this work, the latter conditional law greatly simplifies to a dirac mass at $Y^n = (g(X_i, \theta))_{i=1}^n$.

Definition 4: Denote, as usual, by $Y = g(X, \theta)$ the random variable where $X \sim \mathcal{D}$ and $\theta \sim P_\Theta$ are independent. We define by $Q_n = Q_n(Y^n, X^n)$ the “null” distribution on n samples, where the observations are generated as follows. We sample X^n in an i.i.d. fashion from \mathcal{D} and then generate Y^n in an i.i.d. fashion from the law of Y , *independently* from X^n .

Notice first that the marginals of Y_i under Q_n are identical to the marginals of Y_i under P_n . Yet, the joint law of Q_n does not include any “signal” θ and Y^n are independent of X^n . Naturally, is not possible to estimate any signal θ with observations coming from the null model Q_n .

The following proposition holds.

Proposition 3: For (θ, Y^n, X^n) generated according to P_n , and Q_n defined in Definition 4 we have,

$$I(\theta; Y^n | [X^n]) = nH(Y) - D(P_n || Q_n). \quad (22)$$

and therefore

$$\left(1 - \frac{H(\theta | Y^n, X^n)}{H(\theta)}\right) + \frac{D(P_n || Q_n)}{H(\theta)} = (1 + o(1)) \frac{n}{n^*}. \quad (23)$$

Proof: Note that (23) follows from (22) directly from part (ii) of Lemma 1, along with the identity

$$\begin{aligned} I(\theta; Y^n | [X^n]) &= H(\theta | X^n) - H(\theta | Y^n, X^n) \\ &= H(\theta) - H(\theta | Y^n, X^n), \end{aligned}$$

where the second equality uses that X^n is independent of θ .

We now prove (23). We have

$$\begin{aligned} I(\theta; Y^n | [X^n]) &= \mathbb{E}_{X^n} \left[\mathbb{E}_{\theta, Y^n | X^n} \log \frac{P(Y^n | X^n, \theta)}{P(Y^n | X^n)} \right] \\ &= \mathbb{E}_{X^n} \left[\mathbb{E}_{\theta, Y^n | X^n} \log \frac{1}{P(Y^n | X^n)} \right] \\ &= \mathbb{E}_{X^n} \left[\mathbb{E}_{\theta, Y^n | X^n} \log \frac{1}{Q(Y^n)} \right] \\ &\quad - \mathbb{E}_{X^n} \left[\mathbb{E}_{\theta, Y^n | X^n} \log \frac{P(Y^n | X^n)}{Q(Y^n)} \right] \end{aligned}$$

The second term is $D(P_n || Q_n)$. For the first term, we have that $Q(Y^n) = \prod_{i=1}^n Q(Y_i)$ by the definition of Q . Since by

assumption also $Q(Y_i) = P(Y_i)$ for each i we conclude

$$\begin{aligned} \mathbb{E}_{X^n} \left[\mathbb{E}_{\theta, Y^n | X^n} \log \frac{1}{Q(Y^n)} \right] &= n \mathbb{E}_{X_1} \left[\mathbb{E}_{\theta, Y_1 | X_1} \log \frac{1}{P(Y_1)} \right] \\ &= nH(Y). \end{aligned}$$

The proof is complete. \square

Proposition 4: Suppose that (10) holds for some sequence of sample size $n = n_N$. Then we also have

$$\lim_{N \rightarrow +\infty} \frac{H(\theta | Y^n, X^n)}{H(\theta)} = 0. \quad (24)$$

Proof:

Recall that since the prior is uniform and the model is noiseless, the posterior is simply the uniform distribution over the solutions θ' of the system of equations (6). Hence, using the notation of Definition 1, $Z_{N,0}$ is the random variable which is equal to the number of such solutions. Hence, (24) is equivalent to

$$\lim_{N \rightarrow +\infty} \frac{\mathbb{E} \log Z_{N,0}}{\log M} = 0.$$

where we used that $H(\theta) = \log M$.

Now fix a $\delta \in (0, 2]$ and let $A_\delta := \{\theta'' \in \Theta : \|\theta'' - \theta\| \geq \delta\}$. Notice that almost surely

$$\frac{Z_{N,\delta}}{Z_{N,0}} Z_{N,0} = Z_{N,\delta} \leq |A_\delta|$$

and

$$(1 - \frac{Z_{N,\delta}}{Z_{N,0}}) Z_{N,0} = Z_{N,0} - Z_{N,\delta} \leq |\Theta \setminus A_\delta|.$$

Hence if we denote for simplicity $p_\delta := \frac{Z_{N,\delta}}{Z_{N,0}}$ we have,

$$\begin{aligned} \log Z_{N,0} &\leq p_\delta \log \left(\frac{|A_\delta|}{p_\delta} \right) + (1 - p_\delta) \log \left(\frac{|\Theta \setminus A_\delta|}{1 - p_\delta} \right) \\ &\leq p_\delta \log M + \log |\Theta \setminus A_\delta| + h(p_\delta), \end{aligned}$$

where h is the binary entropy. Since $h(p_\delta) \leq \ln 2$ and $\lim_N \log M_N = +\infty$, we conclude that for any $\delta \in (0, 2]$

$$\begin{aligned} &\frac{\mathbb{E} \log Z_{N,0}}{\log M} \\ &\leq \mathbb{E} p_\delta + \mathbb{E} \frac{\log |\Theta \setminus A_\delta|}{\log M} + \frac{\ln 2}{\log M} \\ &= \mathbb{E} \frac{Z_{N,\delta}}{Z_{N,0}} + \mathbb{E} \frac{\log(M P_{\theta' \sim P_\Theta}(\|\theta - \theta'\| < \delta))}{\log M} + \frac{\ln 2}{\log M} \\ &= \mathbb{E} \frac{Z_{N,\delta}}{Z_{N,0}} + \mathbb{E} \frac{\log(M P_{\theta' \sim P_\Theta}(\langle \theta, \theta' \rangle > 1 - \delta^2/2))}{\log M} \\ &\quad + \frac{\ln 2}{\log M} \\ &\leq \mathbb{E} \frac{Z_{N,\delta}}{Z_{N,0}} + \frac{\log(M P_\Theta^{\otimes 2}(\langle \theta, \theta' \rangle > 1 - \delta^2/2))}{\log M} \\ &\quad + \frac{\ln 2}{\log M}. \end{aligned}$$

where in the last inequality we used Jensen’s inequality and the fact that the logarithm is concave. Now we send first N to infinity and then δ to zero and show that the right hand side

of the last inequality converges to zero. The third term clearly vanishes. The second first term vanishes by the double limit by using Assumption 1. The first term vanishes by the first limit since (9). The proof follows. \square

We state and prove here a folklore result in the statistical physics literature called the “Nishimori” identity, which will be useful in what follows.

Lemma 2: It always holds that if θ' is a random variable drawn from the posterior distribution $P_{\theta|Y^n, X^n}$ that

$$\text{MMSE}_N(n) = 1 - \mathbb{E}\langle\theta, \theta'\rangle.$$

Proof: Bayes’ rule implies that the joint distribution of $\langle\theta', \theta\rangle$ is identical with the distribution of $\langle\theta', \theta''\rangle$ of two independent random variables drawn from the posterior distribution of θ given Y^n, X^n . Therefore,

$$\begin{aligned} \mathbb{E}\langle\mathbb{E}[\theta|Y^n, X^n], \theta\rangle &= \mathbb{E}\langle\theta', \theta\rangle \\ &= \mathbb{E}\langle\theta', \theta''\rangle = \mathbb{E}\|\mathbb{E}[\theta|Y^n, X^n]\|^2. \end{aligned}$$

The result follows since

$$\begin{aligned} \text{MMSE}_N(n) &= 1 + \mathbb{E}\|\mathbb{E}[\theta|Y^n, X^n]\|^2 \\ &\quad - 2\mathbb{E}\langle\mathbb{E}[\theta|Y^n, X^n], \theta\rangle. \end{aligned} \quad (26)$$

\square

APPENDIX II CONVEX ANALYSIS

A. Background

In this work, we use the following two results from convex analysis on the real line.

The first result concerns the left differentiability of a convex function on the interior of its domain.

Theorem 4: [11, Proposition I.4.1.1] For any interval $I \subset \mathbb{R}$, convex function $f : I \rightarrow \mathbb{R}$ and x in the interior of I , the left derivative of f exists at x .

The second result establishes that if a sequence of convex function defined on an open interval converges to a convex differentiable function, the pointwise convergence can be generalized to their (left) derivatives.

Theorem 5: [11, Proposition I.4.3.4] Fix an open interval $I \subset \mathbb{R}$, and consider a sequence $(f_n)_{n \in \mathbb{N}} : I \rightarrow \mathbb{R}$ of convex functions. Assume that f_n converges pointwise to a differentiable $f : I \rightarrow \mathbb{R}$. Then the left derivatives of f_n converge pointwise to the derivative of f , f' .

B. A Key Proposition

Towards employing certain analytic techniques we consider the following function defined on $\mathbb{R}_{>0}$, which simply linear interpolates between the values of the sequence $D(P_n||Q_n)/H(\theta)$, $n \in \mathbb{N}$. We establish that the analytic properties of this function express various fundamental statistical properties of the inference setting of interest. Here and throughout this section, P_n, Q_n are defined as in Definitions 3, 4.

Definition 5: Let $D_N : (0, +\infty) \rightarrow [0, +\infty)$, $N \in \mathbb{N}$ be the sequence of functions defined by

$$\begin{aligned} D_N(\beta) &:= (1 - \beta n^* + \lfloor \beta n^* \rfloor) \frac{D(P_{\lfloor \beta n^* \rfloor}||Q_{\lfloor \beta n^* \rfloor})}{H(\theta)} \\ &\quad + (\beta n^* - \lfloor \beta n^* \rfloor) \frac{D(P_{\lfloor \beta n^* \rfloor + 1}||Q_{\lfloor \beta n^* \rfloor + 1})}{H(\theta)}. \end{aligned} \quad (25)$$

Notice that the normalization of the argument of D_N is appropriately chosen such that $D_N(1) = D(P_{n^*}||Q_{n^*})/H(\theta)$.

Proposition 5: Consider the sequence of functions D_N , per Definition 5. Then under our framework and assumptions described in Section II-A the following hold.

- (a) For each N , D_N is a convex, increasing, nonnegative function.
- (b) For all fixed $\beta > 0$,

$$\limsup_N D_N(\beta) = \beta - 1 + \limsup_N \frac{H(\theta|X^{\lfloor \beta n^* \rfloor}, Y^{\lfloor \beta n^* \rfloor})}{H(\theta)}.$$

- (c) For all fixed $\beta > 0$ and for each N , the function D_N is left differentiable at β and the left derivative at β satisfies

$$(D_N)'_-(\beta) = 1 - \frac{H(Y_{\lfloor \beta n^* \rfloor} | Y^{\lceil \beta n^* \rceil - 1}, X^{\lceil \beta n^* \rceil})}{H(Y)} + o(1), \quad (26)$$

where the $o(1)$ term tends to zero as $N \rightarrow +\infty$.

Proof: We start with part (a).

Since D_N is a linear interpolation of the sequence $\frac{D(P_n||Q_n)}{H(\theta)}$, $n \in \mathbb{N}$ and $H(\theta) > 0$, it suffices to show the same properties for the sequence $D(P_n||Q_n)$, $n \in \mathbb{N}$. The nonnegativity is obvious. For a fixed $n \in \mathbb{N}$ we have using the identity (22) from Proposition 3 that

$$\begin{aligned} D(P_{n+1}||Q_{n+1}) - D(P_n||Q_n) \\ = H(Y) - I(\theta; Y^{n+1} | [X^{n+1}]) + I(\theta; Y^n | X^n) \end{aligned}$$

By the chain rule for mutual information, its definition and the independence of the X_i ’s we have

$$\begin{aligned} I(\theta; Y^{n+1} | X^{n+1}) \\ = I(\theta; Y_{n+1} | X^{n+1}, Y^n) + I(\theta; Y^n | X^{n+1}) \\ = I(\theta; Y_{n+1} | X^{n+1}, Y^n) + I(\theta; Y^n | X^n). \end{aligned}$$

Combining the above and using the definition of the mutual information and the fact that our channel $Y_i = g(X_i, \theta)$ is noiseless, we obtain

$$\begin{aligned} D(P_{n+1}||Q_{n+1}) - D(P_n||Q_n) \\ = H(Y) - I(\theta; Y_{n+1} | X^{n+1}, Y^n) \\ = H(Y) - H(Y_{n+1} | X^{n+1}, Y^n) \end{aligned}$$

$$\begin{aligned} &+ H(Y_{n+1} | X^{n+1}, Y^n, \theta) \\ &= H(Y) - H(Y_{n+1} | X^{n+1}, Y^n) \end{aligned} \quad (27)$$

$$= I(Y_{n+1}; X^{n+1}, Y^n) \quad (28)$$

Now the increasing property of the sequence follows from the fact that the mutual information is non-negative. For the convexity, it suffices to show that the right hand side of (28)

is nondecreasing. Indeed, notice that for each n from the fact that conditioning reduces entropy,

$$\begin{aligned} & I(Y_{n+1}; X^{n+1}, Y^n) \\ &= H(Y_{n+1}) - H(Y_{n+1}|X^{n+1}, Y^n) \\ &\geq H(Y_{n+1}) - H(Y_{n+1}|X_2, \dots, X_{n+1}, Y_2, \dots, Y_n) \\ &= I(Y_{n+1}; X_2, \dots, X_{n+1}, Y_2, \dots, Y_n) \\ &= I(Y_n; X_1, \dots, X_n, Y_1, \dots, Y_{n-1}) \\ &= I(Y_n; X^n, Y^{n-1}). \end{aligned}$$

This completes the proof of part (a).

For part (b) notice that from Proposition 3 we have for each fixed $\beta > 0$

$$\frac{D(P_{\lfloor \beta n^* \rfloor} \| Q_{\lfloor \beta n^* \rfloor})}{H(\theta)} \quad (29)$$

$$\begin{aligned} &= (1+o(1)) \frac{\lfloor \beta n^* \rfloor}{n^*} - 1 + \frac{H(\theta|X^{\lfloor \beta n^* \rfloor}, Y^{\lfloor \beta n^* \rfloor})}{H(\theta)} \\ &= \beta - 1 + \frac{H(\theta|X^{\lfloor \beta n^* \rfloor}, Y^{\lfloor \beta n^* \rfloor})}{H(\theta)} + o(1), \end{aligned} \quad (30)$$

since $n^* \rightarrow +\infty$ by Lemma 1. By (27),

$$\begin{aligned} & \frac{D(P_{\lfloor \beta n^* \rfloor + 1} \| Q_{\lfloor \beta n^* \rfloor + 1})}{H(\theta)} - \frac{D(P_{\lfloor \beta n^* \rfloor} \| Q_{\lfloor \beta n^* \rfloor})}{H(\theta)} \\ &\leq \frac{H(Y)}{H(\theta)} = o(1), \end{aligned}$$

since $H(Y) = O(1)$ and $H(\theta) \rightarrow \infty$.

Since $D_N(\beta)$ is a convex combination of $\frac{D(P_{\lfloor \beta n^* \rfloor} \| Q_{\lfloor \beta n^* \rfloor})}{H(\theta)}$ and $\frac{D(P_{\lfloor \beta n^* \rfloor + 1} \| Q_{\lfloor \beta n^* \rfloor + 1})}{H(\theta)}$, we conclude that

$$\limsup_N D_N(\beta) = \beta - 1 + \limsup_N \frac{H(\theta|X^{\lfloor \beta n^* \rfloor}, Y^{\lfloor \beta n^* \rfloor})}{H(\theta)},$$

as we wanted.

For part (c), recall that D_N is the piecewise linear interpolation of the convex sequence $\frac{D(P_n \| Q_n)}{H(\theta)}$. By [11, Proposition I.4.1.1], stated also in Theorem 4, D_N possesses a left derivative on the interior of its domain, so that D_N is left-differentiable at β for all $\beta > 0$. Moreover, this left derivative is simply the slope of the segment which connects $(\frac{\lfloor \beta n^* \rfloor - 1}{n^*}, \frac{D(P_{\lfloor \beta n^* \rfloor - 1} \| Q_{\lfloor \beta n^* \rfloor - 1})}{H(\theta)})$ and $(\frac{\lfloor \beta n^* \rfloor}{n^*}, \frac{D(P_{\lfloor \beta n^* \rfloor} \| Q_{\lfloor \beta n^* \rfloor})}{H(\theta)})$, which equals

$$\begin{aligned} & \frac{D(P_{\lfloor \beta n^* \rfloor} \| Q_{\lfloor \beta n^* \rfloor})}{H(\theta)} - \frac{D(P_{\lfloor \beta n^* \rfloor - 1} \| Q_{\lfloor \beta n^* \rfloor - 1})}{H(\theta)} \\ & \quad \frac{\lfloor \beta n^* \rfloor}{n^*} - \left(\frac{\lfloor \beta n^* \rfloor - 1}{n^*} \right) \\ &= n^* \frac{D(P_{\lfloor \beta n^* \rfloor} \| Q_{\lfloor \beta n^* \rfloor}) - D(P_{\lfloor \beta n^* \rfloor - 1} \| Q_{\lfloor \beta n^* \rfloor - 1})}{H(\theta)} \\ &= (1+o(1)) \frac{D(P_{\lfloor \beta n^* \rfloor} \| Q_{\lfloor \beta n^* \rfloor}) - D(P_{\lfloor \beta n^* \rfloor - 1} \| Q_{\lfloor \beta n^* \rfloor - 1})}{H(Y)}, \end{aligned}$$

where we use the second part of Lemma 1 for n^* and the $o(1)$ term tends to zero as n^* tends to infinity. Now using (27) we

conclude that the slope is

$$\begin{aligned} & (1+o(1)) \left(1 - \frac{H(Y_{\lceil \beta n^* \rceil})|Y^{\lceil \beta n^* \rceil - 1}, X^{\lceil \beta n^* \rceil})}{H(Y)} \right) \\ &= 1 - \frac{H(Y_{\lceil \beta n^* \rceil})|Y^{\lceil \beta n^* \rceil - 1}, X^{\lceil \beta n^* \rceil})}{H(Y)} + o(1). \end{aligned}$$

The proof is complete. \square

APPENDIX III

PROOF OF THEOREM 1: TURNING “ALL” INTO “NOTHING”

Recall that from Proposition 4, condition (10) implies that at $(1+\epsilon)n^*$ samples the entropy of the posterior distribution is of smaller order than the entropy of the prior. Our first result towards proving Theorem 1 establishes two implications of this property of the entropy of the posterior.

Lemma 3: Suppose that for all $\epsilon > 0$ and $n \geq (1+\epsilon)n^*$, (24) holds. Then we have,

(1) (KL closeness)

$$\lim_N \frac{D(P_{n^*} \| Q_{n^*})}{H(\theta)} = 0, \quad (31)$$

and

(2) (prediction “nothing”) for any fixed $\epsilon > 0$, if $n \leq (1-\epsilon)n^*$, then

$$\lim_N \frac{H(Y_{n+1}|Y^n, X^{n+1})}{H(Y)} = 1. \quad (32)$$

In words, the sublinear entropy of the posterior implies 1) a “KL-closeness” between the planted distribution P_{n^*} and the null distribution Q_{n^*} , and 2) that the entropy of the observation Y_{n+1} conditioned on knowing the past observations Y^n and X^{n+1} is (almost) equal to the unconditional entropy of $Y = Y_{n+1}$. While the first condition is, as already mentioned, hard to interpret (because of the $H(\theta)$ normalization), the second condition has rigorous implication of the recovery problem of interest, because of the following lemma. We emphasize to the reader that towards establishing this lemma the use of an assumption such as Assumption 2 is crucial.

Lemma 4: Suppose that (32) holds. Then for any fixed $\epsilon \in (0, 1)$, if $n = n_N \leq (1-\epsilon)n^*$, then

$$\lim_N \text{MMSE}_N(n) = 1,$$

i.e. (16) holds.

Notice that combining Proposition 4, the part (2) of Lemma 3 and Lemma 4, the proof of Theorem 1 follows in a straightforward manner. We proceed by establishing the two lemmas.

A. Proof of Lemma 3

Proof: We start with establishing (31). Notice that for any $\epsilon > 0$, combining Proposition 5 part (b) for $\beta = 1+\epsilon$ and the condition (24), we have

$$\limsup_N D_N(1+\epsilon) = \epsilon.$$

Using now that D_N is increasing from Proposition 5 part (a), we conclude

$$\limsup_N D_N(1) \leq \epsilon,$$

or as $\epsilon > 0$ was arbitrary and D_N is non-negative,

$$\lim_N D_N(1) = 0. \quad (33)$$

The identity (31) follows because for $\beta = 1$, $\beta n^* \in \mathbb{N}$ and therefore $D_N(1) = \frac{D(P_{n^*} \| Q_{n^*})}{H(\theta)}$.

We now show that (31) implies (32). Notice that using Proposition 5 part (a), D_N is a sequence of increasing, convex and non-negative functions which we restrict to be defined on the compact interval $[0, 1]$. Hence combining with (31) or the equivalent (33), we have

$$\lim_N \sup_{\beta \in [0, 1]} |D_N(\beta)| = \lim_N D_N(1) = 0.$$

Therefore D_N converges uniformly to the zero function.

Now to establish our result notice that since conditioning reduces entropy it suffices to consider the case where $n = \lceil \beta n^* \rceil - 1$ for some fixed $\beta \in (0, 1)$. Using standard analysis result [11, Proposition I.4.3.4], stated also in Theorem 5, since the functions D_N are convex and converge uniformly to 0 in the open interval $(0, 1)$ we can conclude the left derivatives of $D_N(\beta)$ converge to the derivative of the zero function as well, i.e.

$$\lim_N (D_N)'_-(\beta) = 0.$$

Using now (26) from Proposition 5 part (c) for this β we conclude the proof. \square

B. Proof of Lemma 4

Proof: Fix some $\epsilon \in (0, 1)$ and assume $n \leq (1 - \epsilon)n^*$. Denote the probability distribution \tilde{P}_{n+1} on (Y^{n+1}, X^{n+1}) where (Y^n, X^n) are drawn from P_n and Y_{n+1}, X_{n+1} are drawn independently from the marginals P_Y, \mathcal{D} respectively. Notice that \tilde{P} is carefully chosen so that

$$\begin{aligned} D(P_{n+1} \| \tilde{P}_{n+1}) \\ = \mathbb{E} \log \frac{P(Y_{n+1} | Y^n, X^{n+1})}{P(Y_{n+1})} \\ = H(Y) - H(Y_{n+1} | Y^{n-1}, X^{n+1}). \end{aligned}$$

Hence using (32),

$$D(P_{n+1} \| \tilde{P}_{n+1}) = o(H(Y)) = o(1),$$

where we used the assumption that $H(Y) \leq \log |\mathcal{Y}| = O(1)$. Using Pinsker’s inequality we conclude

$$\lim d_{\text{TV}}(P_{n+1}, \tilde{P}_{n+1}) = 0.$$

Now we denote by θ' a sample from the posterior distribution $P_n(\theta | Y^n, X^n)$. Using the total variation guarantee we have

$$\begin{aligned} P_{n+1} \{g(X_{n+1}, \theta') = Y_{n+1}\} \\ = \tilde{P}_{n+1} \{g(X_{n+1}, \theta') = Y_{n+1}\} + o(1). \end{aligned}$$

Under \tilde{P} , because of its definition, we can write Y_{n+1} as $g(X_{n+1}, \theta'')$, where $\theta'' \sim P_\theta$ is independent of everything else. Using Assumption 2 we conclude

$$\mathbb{E}R(\langle \theta, \theta' \rangle) = \mathbb{E}R(\langle \theta'', \theta' \rangle) + o(1).$$

Furthermore, using (5) and the fact that θ'' is independent from θ' we conclude that for any $\epsilon > 0$,

$$\limsup_N \mathbb{E}R(\langle \theta, \theta' \rangle) \leq R(\epsilon)$$

Hence by continuity of R at 0 we have

$$\limsup_N \mathbb{E}R(\langle \theta, \theta' \rangle) \leq R(0).$$

Recall that 0 is the unique minimizer of R on $[0, 1]$ which allows us to conclude

$$|\mathbb{E}R(\langle \theta, \theta' \rangle) - R(0)| = o(1).$$

and therefore by Markov’s inequality for any $\epsilon > 0$,

$$P(R(\langle \theta, \theta' \rangle) > R(0) + \epsilon) = o(1).$$

As R is strictly increasing we conclude that, for any $\epsilon > 0$,

$$P(\langle \theta, \theta' \rangle > \epsilon) = o(1).$$

Since the integrand is bounded from above we conclude that

$$\limsup_N \mathbb{E}\langle \theta, \theta' \rangle \leq 0.$$

Using now Lemma 2, we conclude

$$\liminf_N \text{MMSE}_N(n) \geq 1,$$

which concludes the proof. \square

APPENDIX IV

PROOF OF COROLLARY 1: ESTABLISHING THE “ALL”

Proof: We apply Theorem 1. We fix some $\epsilon > 0$ and want to show that for if $n \geq (1 + \epsilon)n^*$, (11) holds, which based on Proposition 1] implies the desired “all” condition (10). We also assume without loss of generality that $n \leq Cn^*$ for an absolute positive constant C , since the random variables $Z_{N,\delta}$ are decreasing in the stochastic order as functions of the samples size n . By assumption p is a fixed constant in $(0, 1)$ independent of N . Hence since $h(p)$ is a positive constant itself we conclude from the definition of n^* that for all $n = \Theta(n^*)$ it also holds

$$n = \Theta(\log M). \quad (34)$$

In particular, $n \rightarrow +\infty$ as $N \rightarrow +\infty$.

Consider n_1 the number of samples where $Y_i = 1$ and notice that n_1 is distributed as a Binomial distribution $\text{Bin}(n, p)$. We condition on the event that $\mathcal{F} = \{|n_1 - np| \leq \sqrt{n} \log \log n\}$. Standard large deviation theory on the Binomial distribution yields that since $p \in (0, 1)$ and $n \rightarrow +\infty$, it holds $\lim_N P(\mathcal{F}) = 1$.

Therefore, by Markov’s inequality it suffices to prove that for every $\delta > 0$,

$$\lim_N \mathbb{E}[Z_{N,\delta} \mathbf{1}(\mathcal{F})] = 0. \quad (35)$$

or equivalently by linearity of expectation and the independence of $Y_i, X_i, i = 1, 2, \dots, n$ given θ ,

$$\lim_N \mathbb{E} \left\{ 1(\mathcal{F}) \sum_{\theta': \|\theta' - \theta\|^2 \geq \delta} P \left(\bigcap_{i=1}^n \{Y_i = g(X_i, \theta')\} \middle| n_1 \right) \right\} = 0.$$

Now fix any θ' with $\|\theta' - \theta\|^2 \geq \delta$ or equivalently $\rho = \langle \theta, \theta' \rangle \leq 1 - \frac{\delta}{2}$ and some n_1 satisfying \mathcal{F} . Using the definitions of $R_i, i = 0, 1$ from Assumption 3 we have that

$$P \left(\bigcap_{i=1}^n \{Y_i = g(X_i, \theta')\} \middle| n_1 \right)$$

equals

$$\begin{aligned} & \binom{n}{n_1} P(g(X, \theta) = g(X, \theta') = 1)^{n_1} \\ & \times P(g(X, \theta) = g(X, \theta') = 0)^{n-n_1} \\ & = \binom{n}{n_1} R_1(\rho)^{n_1} R_0(\rho)^{n-n_1} \\ & = \exp \left(nh \left(\frac{n_1}{n} \right) + n_1 \log R_1(\rho) + (n - n_1) \log R_0(\rho) + o(n) \right) \end{aligned} \quad (36)$$

where h is defined in (18), and we used the standard application of Stirling's formula $\log \binom{n}{nx} = nh(x) + o(n)$ when x is bounded away from 0 and 1. The last expression equals to

$$\begin{aligned} & \exp \left(n \left(h \left(\frac{n_1}{n} \right) + \frac{n_1}{n} \log R_1(\rho) + \left(1 - \frac{n_1}{n} \right) \log R_0(\rho) \right) + o(n) \right) \\ & = \exp \left(n (h(p) + p \log R_1(\rho)) + (1 - p) \log R_0(\rho) + o(n) \right) \end{aligned} \quad (37)$$

$$= \exp \left(n \left(p \log \frac{R_1(\rho)}{p} + (1 - p) \log \frac{R_0(\rho)}{(1 - p)} \right) + o(n) \right), \quad (38)$$

and for (37) we used the continuity of h and that $n_1/n = p(1 + O(\log \log n / \sqrt{n})) = p(1 + o(1))$, since $n \rightarrow +\infty$. Importantly, since $p \in (0, 1)$ the $o(n)$ term in (38) can be taken to hold uniformly over the specific choices of n_1 satisfying \mathcal{F} .

Using (38) it suffices to establish for $G(\rho, p) = p \log \frac{R_1(\rho)}{p} + (1 - p) \log \frac{R_0(\rho)}{(1 - p)}$ that

$$\lim_N \mathbb{E}_\theta \sum_{\rho \in \mathcal{R}: \rho \leq 1 - \frac{\delta}{2}} |\{\theta' \in \Theta : \langle \theta, \theta' \rangle = \rho\}| e^{nG(\rho, p) + o(n)} = 0, \quad (39)$$

where \mathcal{R} denotes the support of the overlap distribution of two independent samples from the prior P_Θ . Now since the prior is uniform over Θ if ρ is drawn from the law of the inner product between two independent samples from the prior, (39) is equivalent with

$$\lim_N \mathbb{E}_\rho 1(\rho \leq 1 - \frac{\delta}{2}) M e^{nG(\rho, p) + o(n)} = 0. \quad (40)$$

Now notice that since $n^* = (1 + o(1)) \frac{\log M}{h(p)}$ by Proposition 1 we have

$$\begin{aligned} & M e^{nG(\rho, p) + o(n)} \\ & = \exp(\log M + n^* G(\rho, p) \\ & + (n - n^*) G(\rho, p) + o(n)) \\ & = \exp(n^* h(p) + n^* G(\rho, p) \\ & + (n - n^*) G(\rho, p) + o(n + \log M)) \\ & = \exp(n^* (p \log \frac{R_1(\rho)}{p} \\ & + (1 - p) \log \frac{R_0(\rho)}{(1 - p)^2}) + (n - n^*) G(\rho, p) + o(\log M)), \end{aligned} \quad (41)$$

where we used that n is of order $\log M$, by (34).

Now Assumption 3 implies that the functions $R_i, i = 0, 1$ are increasing in $[0, 1]$ and Assumption 2 that their sum is strictly increasing in $[0, 1]$. Furthermore, notice that at full correlation it holds $R_1(1) = p, R_0(1) = 1 - p$. Hence we conclude that for some $\delta' > 0$ the following holds; for any $\rho \leq 1 - \frac{\delta}{2}$,

$$\begin{aligned} & \min \{ \log \frac{R_1(\rho)}{p}, \log \frac{R_0(\rho)}{1-p} \} \\ & \leq \min \{ \log \frac{R_1(1)}{p}, \log \frac{R_0(1)}{1-p} \} - \delta' = -\delta' \end{aligned}$$

and

$$\begin{aligned} & \max \{ \log \frac{R_1(\rho)}{p}, \log \frac{R_0(\rho)}{1-p} \} \\ & \leq \max \{ \log \frac{R_1(1)}{p}, \log \frac{R_0(1)}{1-p} \} = 0. \end{aligned}$$

Hence, since $p \in (0, 1)$, from the definition of $G(\rho, p)$ we conclude that for $\delta'' = \delta' \min\{p, 1 - p\} > 0$ it holds that for all $\rho \leq 1 - \frac{\delta}{2}$, $G(\rho, p) \leq -\delta''$. Hence since $n \geq (1 + \epsilon)n^*$ and $n^* = \Theta(\log M)$ we conclude that for all $\rho \leq 1 - \frac{\delta}{2}$,

$$(n - n^*) G(\rho, p) \leq -\epsilon \delta'' n^* = -\Omega(\log M). \quad (42)$$

Combining (41) with (42), and then using

$$n^* = \frac{\log M}{h(p)} + o(\log M),$$

we conclude

$$\begin{aligned} & \mathbb{E}_\rho 1(\rho \leq 1 - \frac{\delta}{2}) M e^{nG(\rho, p) + o(n)} \\ & \leq e^{-\Omega(\log M)} \times \\ & \mathbb{E}_\rho \exp \left(n^* \left(p \log \frac{R_1(\rho)}{p^2} + (1 - p) \log \frac{R_0(\rho)}{(1 - p)^2} \right) \right) \\ & = e^{-\Omega(\log M)} \times \\ & \mathbb{E}_\rho \exp \left(\frac{\log M}{h(p)} \left(p \log \frac{R_1(\rho)}{p^2} + (1 - p) \log \frac{R_0(\rho)}{(1 - p)^2} \right) \right). \end{aligned}$$

Hence we are left with establishing

$$\limsup_N \frac{1}{\log M} \log \mathbb{E}_\rho \exp(W(\rho, p) \log M) = 0, \quad (43)$$

for $W(\rho, p) \triangleq \frac{1}{h(p)} \left(p \log \frac{R_1(\rho)}{p^2} + (1 - p) \log \frac{R_0(\rho)}{(1 - p)^2} \right)$.

To prove it, let us fix a positive integer k . We have

$$\begin{aligned} & \mathbb{E}_\rho \exp(W(\rho, p) \log M) \\ & \leq \sum_{\ell=0}^{k-1} P[\rho \geq \ell/k] \sup_{t \in [\ell/k, (\ell+1)/k]} \exp(W(t, p) \log M) \\ & \leq k \cdot \max_{0 \leq \ell \leq k} \sup_{t \in [\ell/k, (\ell+1)/k]} \exp(W(t, p) \log M + \log P[\rho \geq \ell/k]). \end{aligned}$$

Therefore by using the overlap rate function r ,

$$\begin{aligned} & \limsup_N \frac{1}{\log M} \log \mathbb{E}_\rho \exp(W(\rho, p) \log M) \\ & \leq \max_{0 \leq \ell \leq k} \sup_{t \in [\ell/k, (\ell+1)/k]} \left(W(t, p) - r\left(\frac{\ell}{k}\right) \right) \\ & \leq \sup_{t \in [0, 1]} (W(t, p) - r(t)) + \sup_{t, t' \in [0, 1]: |t-t'| \leq \frac{1}{k}} |r(t) - r(t')|. \end{aligned}$$

Sending $k \rightarrow +\infty$ using the uniform continuity of r (implied by e.g. the Heine-Cantor theorem) we conclude

$$\begin{aligned} & \limsup_N \frac{1}{\log M} \log \mathbb{E}_\rho \exp(W(\rho, p) \log M) \\ & \leq \sup_{t \in [0, 1]} (W(t, p) - r(t)). \end{aligned}$$

The assumption (21) completes the proof. \square

APPENDIX V APPLICATIONS: THE PROOFS

In this section we present the proofs for the three families of models we establish the all-or-nothing phenomenon using our technique. The proof concept remains the same across the different models; we apply Corollary 1 and check that all assumptions apply.

A. Proof of Theorem 2

Proof: We apply Corollary 1. Notice that for any fixed $\theta \in \Theta$ the random variable $\sqrt{k}\langle X_i, \theta \rangle$ follows a Binomial distribution $\text{Bin}(k, \frac{\nu}{k})$. Therefore

$$p = 1 - P(g(X, \theta) = 0) = 1 - \left(1 - \frac{\nu}{k}\right)^k = 1 - q. \quad (44)$$

Hence $h(p) = h(q)$ and the critical sample size is indeed $n^* = \lfloor \frac{\log(\frac{N}{k})}{h(q)} \rfloor$, and Stirling’s formula implies that since $k = o(N)$, $H(\theta) = \log \binom{N}{k} = (1 + o(1))k \log \frac{N}{k}$ and therefore it also holds $n^* = (1 + o(1))\frac{k \log \frac{N}{k}}{h(q)}$.

We now check the assumptions of the Corollary.

1) *Assumption 1:* We start with Assumption 1, which concerns properties of the prior. We use [18, Lemma 6] to conclude that the prior P_Θ admits the overlap rate function $r(t) = t$, per Definition 2. Now, notice that the first part of Assumption 1 is directly implied by the fact that $r(1 + \delta) = 1 + \delta > 1$ for any fixed $\delta > 0$. For the second part notice that since the law of the prior is permutation-invariant with respect to the N dimensions, for any fixed $\theta \in \Theta$ and θ' chosen from the prior, $\langle \theta, \theta' \rangle$ is equal in distribution to the law of $\langle \theta, \theta' \rangle$ where θ, θ' are two independent samples from

the prior. Hence using the overlap rate function $r(t) = t$ we have that for $M_N = \binom{N}{k}$ it holds that for any $\epsilon > 0$,

$$P(\langle \theta, \theta' \rangle > \epsilon) \leq \exp(-(\epsilon + o(1)) \log M_N) = o(1),$$

as desired.

2) *Assumptions 2, 3:* For Assumption 2 and Assumption 3 we directly compute by elementary combinatorics the functions $R_i(\rho)$, $i = 1, 2$ and $R(\rho)$. Recall that $\{g(X, \theta') = 1\}$ is the event that the supports of θ' and X have a non-empty intersection. Given a N -dimensional vector $v \in \mathbb{R}^n$, we denote its support by $S(v) := \{i \in [N] : v_i \neq 0\}$. First, fix some $\rho \in [0, 1]$ and we compute $R_\rho(1)$ by considering two arbitrary θ, θ' which share ρk indices in their support. Notice that for the argument to be non-vacuous we assume also that $\rho = \ell/k$ for some $\ell \in \{0, 1, 2, \dots, k\}$. Conditioning on whether $S(X)$ intersects $S(\theta) \cap S(\theta')$, we have

$$\begin{aligned} R_1(\rho) &= P(g(X, \theta) = g(X, \theta') = 1) \\ &= P(S(X) \cap S(\theta) \neq \emptyset, S(X) \cap S(\theta') \neq \emptyset) \\ &= \underbrace{1 - \left(1 - \frac{\nu}{k}\right)^\ell}_{\text{case } S(X) \cap S(\theta) \cap S(\theta') \neq \emptyset} \\ &\quad + \underbrace{\left(1 - \frac{\nu}{k}\right)^\ell \left(1 - \left(1 - \frac{\nu}{k}\right)^{k-\ell}\right)^2}_{\text{case } S(X) \cap S(\theta) \cap S(\theta') = \emptyset} \\ &= 1 - 2 \left(1 - \frac{\nu}{k}\right)^k + \left(1 - \frac{\nu}{k}\right)^{2k-\ell} \\ &= 1 - 2q + q^{2-\rho}. \end{aligned}$$

Likewise,

$$\begin{aligned} R_0(\rho) &= P(g(X, \theta) = g(X, \theta') = 0) \\ &= P(S(X) \cap S(\theta) = S(X) \cap S(\theta') = \emptyset) \\ &= \left(1 - \frac{\nu}{k}\right)^{2k-\ell} \\ &= q^{2-\rho} \end{aligned}$$

We obtain

$$R(\rho) = 1 - 2q + 2q^{2-\rho}.$$

It can be straightforwardly checked that all three functions are strictly increasing and continuous in $[0, 1]$.

3) *Condition (21):* Finally, we need to check the condition (21). First notice that since $r(t) = t$ we need to show that for all $\rho \in [0, 1]$,

$$\rho h(p) \geq \left(p \log \frac{R_1(\rho)}{p^2} + (1-p) \log \frac{R_0(\rho)}{(1-p)^2} \right),$$

or using the definition of h ,

$$0 \geq p \log \frac{R_1(\rho)}{p^{2-\rho}} + (1-p) \log \frac{R_0(\rho)}{(1-p)^{2-\rho}}. \quad (45)$$

Notice that for any ρ ,

$$\frac{R_0(\rho)}{(1-p)^{2-\rho}} = 1.$$

Therefore it suffices to show that for every $\rho \in [0, 1]$,

$$R_1(\rho) \leq p^{2-\rho}$$

or equivalently with respect to $q = 1 - p$,

$$1 - 2q + q^{2-\rho} \leq (1 - q)^{2-\rho}.$$

To prove the latter, recall that $q \leq \frac{1}{2}$ and consider the function $f(\rho) = (1 - q)^{2-\rho} - q^{2-\rho}$. Notice that $f(0) = f(1) = 1 - 2q$ and therefore it suffices to prove that f is concave in $[0, 1]$. The second derivative of f is

$$\begin{aligned} f''(\rho) &= \log(1 - q)^2(1 - q)^{2-\rho} - (\log q)^2 q^{2-\rho} \\ &= (1 - q)^{2-\rho} \left(\log(1 - q)^2 - (\log q)^2 \left(\frac{q}{1 - q} \right)^{2-\rho} \right) \\ &\leq (1 - q)^{2-\rho} \left(\log(1 - q)^2 - (\log q)^2 \left(\frac{q}{1 - q} \right)^2 \right), \end{aligned}$$

since $q \leq \frac{1}{2}$. Hence, it suffices to show $\log(1 - q)^2 \leq (\log q)^2 \left(\frac{q}{1 - q} \right)^2$ or $(1 - q) \log(1 - q) \geq q \log q$. To prove the latter consider the function $g(q) = (1 - q) \log(1 - q) - q \log q$, $q \in (0, \frac{1}{2}]$. Notice that $g(0^+) = g(1/2) = 0$, and that for each $q \in (0, \frac{1}{2})$ $g''(q) = \frac{2q-1}{q(1-q)} < 0$. Hence, $g(q)$ is concave on the interval $[0, 1/2]$, and $g(q) \geq \min\{g(0^+), g(1/2)\} = 0$ as we wanted. The proof is complete. \square

B. Proof of Theorem 3

Proof: We apply Corollary 1. Notice that since any fixed $\theta \in \Theta$ lies on the unit sphere in \mathbb{R}^N and $X_i \sim N(0, I_N)$, it holds that $\langle X_i, \theta \rangle \sim N(0, 1)$. Hence

$$p = P(g(X_i, \theta) = 1) = P(\langle X_i, \theta \rangle \in A) = \frac{1}{2}.$$

Hence indeed the critical sample size is $n^* = \lfloor \frac{\log(\frac{N}{k})}{h(\frac{1}{2})} \rfloor = (1 + o(1)) \lfloor k \log_2 \frac{N}{k} \rfloor$, where we have Stirling's formula and the assumption that $k = o(N)$.

We now check the assumptions of the Corollary.

1) *Assumption 1:* Assumption 1 concerns properties of the prior P_Θ and they are already established in the corresponding part of Theorem 2, since the prior is identical.

2) *Assumptions 2, 3:* For Assumption 2 and Assumption 3, we study the functions $R_i(\rho)$, $i = 0, 1$ and $R(\rho)$.

Now we compute the functions. Recall that $\{g(X, \theta) = 1\} = \{\langle X_i, \theta \rangle \in A\}$ and that for any $\theta, \theta' \in \Theta$ with $\langle \theta, \theta' \rangle = \rho$ the pair $\langle X_i, \theta \rangle, \langle X_i, \theta' \rangle$ is a bivariate pair of standard Gaussians with correlation ρ . Letting (Z, Z_ρ) be such a pair, we therefore have

$$\begin{aligned} R_1(\rho) &= P(g(X, \theta) = 1) = P(Z \in A, Z_\rho \in A), \\ R_0(\rho) &= P(g(X, \theta) = 0) = P(Z \notin A, Z_\rho \notin A), \\ R(\rho) &= P(g(X, \theta) = g(X, \theta')) \\ &= P(Z \in A, Z_\rho \in A) + P(Z \notin A, Z_\rho \notin A). \end{aligned}$$

Furthermore, because A is balanced we have for any $\rho \in [0, 1]$,

$$\begin{aligned} R(\rho) &= P(\{Z \in A, Z_\rho \in A\} \cup \{Z \notin A, Z_\rho \notin A\}) \\ &= P(Z \in A, Z_\rho \in A) + P(Z \notin A, Z_\rho \notin A) \\ &= P(Z \in A, Z_\rho \in A) + (1 - P(Z \in A)) \\ &\quad - P(Z_\rho \in A) + P(Z \in A, Z_\rho \in A)) \end{aligned}$$

$$\begin{aligned} &= 2P(Z \in A, Z_\rho \in A) \\ &= 2R_1(\rho) \\ &= 2R_0(\rho). \end{aligned}$$

The uniform limits are all strictly increasing with respect to $\rho \in [0, 1]$ and continuous at $0+$, by Lemma 5 applied to A and A^C .

3) *Condition (21):* Finally, we need to check the condition (21). First notice that similar to Theorem 2 the prior admits the overlap rate function $r(t) = t$ and therefore the condition is equivalent with (45). Notice that (45) simplifies since $p = 1/2$ in our case to

$$R_1(\rho) \leq 2^{\rho-2},$$

or

$$P(Z \in A, Z_\rho \in A) \leq 2^{\rho-2}. \quad (46)$$

By Borell's noise stability theorem [7], since $P(Z \in A) = 1/2 = P(Z \geq 0)$, we have

$$\begin{aligned} P(Z \in A, Z_\rho \in A) \\ \leq P(Z \geq 0, Z_\rho \geq 0) = \frac{1}{4} \left(1 + \frac{2}{\pi} \arcsin \rho \right), \end{aligned}$$

where the equality is by Sheppard's formula [24].

Hence it suffices to show that for all $\rho \in [0, 1]$ it holds $2^\rho \geq 1 + \frac{2}{\pi} \arcsin(\rho)$. We consider the function $g(\rho) = 2^\rho - 1 - \frac{2}{\pi} \arcsin(\rho)$, $\rho \in [0, 1]$. It suffices to show $g(\rho) \geq 0$ for all $\rho \in [0, 1]$.

Now notice $g(0) = g(1) = 0$ and $g(\frac{1}{2}) = \sqrt{2} - 4/3 > 0$. We claim that there is no root of g in $(0, 1)$ which implies the result by Bolzano's theorem. Arguing by contradiction, if there was a root then the derivative

$$g'(\rho) = 2^\rho \ln 2 - \frac{2}{\pi} \frac{1}{\sqrt{1 - \rho^2}}$$

would have two roots in $(0, 1)$ by Rolle's theorem. Rearranging, this is equivalent to the equation

$$\rho \ln 2 + \frac{1}{2} \ln(1 - \rho^2) = \ln\left(\frac{2}{\pi \ln 2}\right)$$

having two roots in $(0, 1)$. But the function on the left side is concave and is zero for $\rho = 0$, so it takes each negative value at most once. Since $\ln(\frac{2}{\pi \ln 2}) < 0$ we are done. \square

APPENDIX VI REMAINING PROOFS

Proof of Proposition 1: The equality (7) follows in a straightforward manner from the observation that the posterior of θ given Y^n, X^n is the uniform measure over the solutions θ' of equations (6), and the definition of $Z_{N,\delta}$.

For (8) notice that by Cauchy-Schwarz inequality that if θ' is drawn from the posterior $P_{\theta|Y^n, X^n}$,

$$\begin{aligned} \text{MMSE}_N(n) &= \mathbb{E}\|\theta - \mathbb{E}[\theta|Y^n, X^n]\|^2 \\ &= \frac{1}{2}\mathbb{E}\|\theta - \theta'\|^2 \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{2} \int_{\delta=0}^2 P(\|\theta - \theta'\|^2 \geq \delta) \\
&= \frac{1}{2} \int_{\delta=0}^2 \mathbb{E}P(\|\theta - \theta'\|^2 \geq \delta | Y^n, X^n) \\
&= \frac{1}{2} \mathbb{E} \int_{\delta=0}^2 P(\|\theta - \theta'\|^2 \geq \delta | Y^n, X^n),
\end{aligned}$$

where in the second line we have used Lemma 2 and where in the last line we are allowed to exchange the order of integration by Tonelli’s theorem as all integrands are non-negative. Notice that finally (7) allows us to conclude (8).

For the second part, fix some arbitrary $\epsilon \in (0, 2]$ and set $A = \{Z_{N,\epsilon}/Z_{N,0} > \epsilon\}$. From (11) we have $P(A) = o(1)$. Notice that for any ϵ' with $2 \geq \epsilon' \geq \epsilon$, it holds almost surely $Z_{N,\epsilon'}/Z_{N,0} \leq \epsilon$. Hence, we have using (8),

$$\begin{aligned}
&\text{MMSE}_N(n) \\
&\leq \mathbb{E} \int_{\delta=0}^2 \frac{Z_{N,\delta}}{Z_{N,0}} d\delta \\
&= \mathbb{E} \int_{\delta=0}^2 \frac{Z_{N,\delta}}{Z_{N,0}} \mathbf{1}(A^c) d\delta + \mathbb{E} \int_{\delta=0}^2 \frac{Z_{N,\delta}}{Z_{N,0}} \mathbf{1}(A) d\delta \\
&\leq \mathbb{E} \int_{\delta=0}^{\epsilon} \frac{Z_{N,\delta}}{Z_{N,0}} d\delta + 2\epsilon + P(A) \\
&\leq 3\epsilon + o(1).
\end{aligned}$$

Therefore,

$$\limsup_N \text{MMSE}_N(n) \leq 3\epsilon.$$

As $\epsilon \in (0, 2]$ was arbitrary we conclude (10).

The other direction follows in a straightforward manner since for any fixed $\epsilon > 0$,

$$\text{MMSE}_N(n) \geq \mathbb{E} \int_{\delta=0}^{\epsilon} \frac{Z_{N,\delta}}{Z_{N,0}} d\delta \geq \epsilon \mathbb{E} \frac{Z_{N,\epsilon}}{Z_{N,0}}.$$

□

Proof of Proposition 2: Using (23) from Proposition 3 we have that since the KL divergence is non-negative,

$$1 - \frac{H(\theta | Y^n, X^n)}{H(\theta)} \leq (1 + o(1)) \frac{n}{n^*} \quad (47)$$

Using now Proposition 4 we conclude that (24) holds. Combining (24) with (47) concludes the result. □

Lemma 5: Let Z and Z_ρ be a bivariate pair of standard Gaussians with correlation ρ . Then for any Borel set $A \subseteq \mathbb{R}$ such that $P(Z \in A) \in (0, 1)$, the function

$$\rho \mapsto P(Z \in A, Z_\rho \in A)$$

is strictly increasing on $[0, 1]$ and continuous on $[0, 1]$.

Proof: Write γ for the standard Gaussian measure on \mathbb{R} . We recall [see, e.g. [19], Proposition 11.37] that there exists an orthonormal basis $\{h_k\}_{k \geq 0}$ for $L_2(\gamma)$ such that, for any $f \in L_2(\gamma)$,

$$\mathbb{E}[f(Z)f(Z_\rho)] = \sum_{k \geq 0} \rho^k \hat{f}_k^2,$$

where the coefficients $\{\hat{f}_k\}_{k \geq 0}$ are defined by

$$f = \sum_{k \geq 0} \hat{f}_k h_k \text{ in } L_2(\gamma).$$

Moreover, $h_0 = 1$, so that if f is not γ -a.s. constant, there exists a $k > 0$ for which $\hat{f}_k \neq 0$. We obtain that, for any non-constant f , the function $\mathbb{E}[f(Z)f(Z_\rho)] = \sum_{k \geq 0} \rho^k \hat{f}_k^2$ is strictly increasing on $[0, 1]$. Furthermore, by Parseval’s identity $\sum_{k \geq 0} \hat{f}_k^2 = \mathbb{E}[f^2(Z)] < +\infty$. Hence then function $\mathbb{E}[f(Z)f(Z_\rho)] = \sum_{k \geq 0} \rho^k \hat{f}_k^2$ is also continuous on $[0, 1]$.

Applying this result to the non-constant function $f(x) = 1(x \in A) \in L_2(\gamma)$ yields the claim. □

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for the very helpful feedback.

REFERENCES

- [1] M. Aldridge, O. Johnson, and J. Scarlett, “Group testing: An information theory perspective,” *Found. Trend. Comms. Inf. Theory*, vol. 15, nos. 3–4, pp. 196–392, 2019.
- [2] B. Aubin, W. Perkins, and L. Zdeborová, “Storage capacity in symmetric binary perceptrons,” *J. Phys. A, Math. Theor.*, vol. 52, no. 29, Jul. 2019, Art. no. 294003.
- [3] J. Barbier, F. Krzakala, N. Macris, L. Miolane, and L. Zdeborová, “Optimal errors and phase transitions in high-dimensional generalized linear models,” *Proc. Nat. Acad. Sci. USA*, vol. 116, no. 12, pp. 5451–5460, Mar. 2019.
- [4] J. Barbier and N. Macris, “0–1 phase transitions in sparse spiked matrix estimation,” 2019, *arXiv:1911.05030*.
- [5] J. Barbier et al., “All-or-nothing statistical and computational phase transitions in sparse spiked matrix estimation,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 14915–14926.
- [6] J. Banks, C. Moore, R. Vershynin, N. Verzelen, and J. Xu, “Information-theoretic bounds and phase transitions in clustering, sparse PCA, and submatrix localization,” *IEEE Trans. Inf. Theory*, vol. 64, no. 7, pp. 4872–4994, Feb. 2018.
- [7] C. Borell, “Geometric bounds on the Ornstein-Uhlenbeck velocity process,” *Probab. Theory Related Fields*, vol. 70, no. 1, pp. 1–13, 1985.
- [8] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, A. S. Wein, and I. Zadik, “Statistical and computational phase transitions in group testing,” in *Proc. 35th Conf. Learn. Theory*, Jul. 2022, pp. 4764–4781.
- [9] R. Dorfman, “The detection of defective members of large populations,” *Ann. Math. Statist.*, vol. 14, no. 4, pp. 436–440, Dec. 1943.
- [10] D. Gamarnik and I. Zadik, “Sparse high-dimensional linear regression: Estimating squared error and a phase transition,” *Ann. Statist.*, vol. 50, no. 2, pp. 880–903, Apr. 2022.
- [11] J.-B. Hiriart-Urruty and C. Lemaréchal, “Convex analysis and minimization algorithms. I,” volume 305 of *Grundlehren der Mathematischen Wissenschaften*, in *Fundamental Principles of Mathematical Sciences*. Berlin, Germany: Springer-Verlag, 1993.
- [12] F. Iliopoulos and I. Zadik, “Group testing and local search: Is there a computational-statistical gap?” in *Proc. Conf. Learn. Theory*, 2021, pp. 2499–2551.
- [13] S. Kudekar, S. Kumar, M. Mondelli, H. D. Pfister, and R. L. Urbanke, “Reed–Müller codes achieve capacity on erasure channels,” *IEEE Trans. Inf. Theory*, vol. 63, no. 7, pp. 4298–4316, Jul. 2017.
- [14] C. Luneau et al., “Information theoretic limits of learning a sparse rule,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 9948–9959.
- [15] M. Mezard and A. Montanari, *Information, Physics, and Computation*. Oxford, U.K.: Oxford Univ. Press, 2009.
- [16] C. Méasson, A. Montanari, and R. L. Urbanke, “Maxwell construction: The hidden bridge between iterative and maximum *a posteriori* decoding,” *IEEE Trans. Inf. Theory*, vol. 54, no. 12, pp. 5277–5307, Dec. 2008.
- [17] J. Niles-Weed et al., “The all-or-nothing phenomenon in sparse tensor PCA,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, pp. 17674–17684.
- [18] R. O’Donnell, *Analysis of Boolean Functions*. New York, NY, USA: Cambridge Univ. Press, 2014.
- [19] Y. Polyanskiy and Y. Wu, “Lecture notes on information theory,” *Inf. Theory From Coding Learn.*, Tech. Rep., 2015. [Online]. Available: <https://people.lids.mit.edu/yp/homepage/data/itbook-export.pdf>

- [20] G. Reeves and H. D. Pfister, "The replica-symmetric prediction for compressed sensing with Gaussian matrices is exact," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2016, pp. 665–669.
- [21] G. Reeves, J. Xu, and I. Zadik, "All-or-nothing phenomena: From single-letter to high dimensions," in *Proc. IEEE 8th Int. Workshop Comput. Adv. Multi-Sensor Adapt. Process. (CAMSAP)*, Dec. 2019, pp. 654–658.
- [22] G. Reeves, J. Xu, and I. Zadik, "The all-or-nothing phenomenon in sparse linear regression," in *Proc. 32nd Conf. Learn. Theory*, Phoenix, AZ, USA, Jun. 2019, pp. 2652–2663.
- [23] J. Scarlett and V. Cevher, "Phase transitions in group testing," in *Proc. 27th Annu. ACM-SIAM Symp. Discrete Algorithms*, Jan. 2016, pp. 40–53.
- [24] W. F. Sheppard, "On the application of the theory of error to cases of normal distribution and normal correlation," *Phil. Trans. Roy. Soc. London A, Containing Papers Math. Phys. Character*, vol. 192, pp. 101–167, Dec. 1899.
- [25] L. V. Truong, M. Aldridge, and J. Scarlett, "On the all-or-nothing behavior of Bernoulli group testing," *IEEE J. Sel. Areas Inf. Theory*, vol. 1, no. 3, pp. 669–680, Nov. 2020.
- [26] I. Zadik, "Computational and statistical challenges in high dimensional statistical models," Ph.D. dissertation, Massachusetts Inst. Technol., Cambridge, MA, USA, 2019.
- [27] L. Zdeborová and F. Krzakala, "Statistical physics of inference: Thresholds and algorithms," *Adv. Phys.*, vol. 65, no. 5, pp. 453–552, 2016.

Jonathan Niles-Weed received the Ph.D. degree in mathematics and statistics from MIT in 2019. From 2019 to 2020, he was a Post-Doctoral Member at the Institute for Advanced Study, Princeton, NJ, USA. Since 2020, he has been an Assistant Professor of mathematics and data science at New York University. His research interests include statistics, probability, and the mathematics of data science.

Ilias Zadik received the Ph.D. degree in operations research from MIT in 2019. For two years, he was a CDS Moore-Sloan Post-Doctoral Fellow with the Center for Data Science, New York University. He is currently a Post-Doctoral Associate with the Department of Mathematics, MIT. His research interests include the interface of high-dimensional statistics, the theory of machine learning, and applied probability.