

Optimal Estimation of Change in a Population of Parameters

Ramya Korlakai Vinayak, Weihao Kong, Sham M. Kakade

Paul. G. Allen School of Computer Science & Engineering, University of Washington
{ramya, whkong, sham}@cs.washington.edu

Abstract

Paired estimation of change in parameters of interest over a population plays a central role in several application domains including those in the social sciences, epidemiology, medicine and biology. In these domains, the size of the population under study is often very large, however, the number of observations available per individual in the population is very small (*sparse observations*) which makes the problem challenging. Consider the setting with N independent individuals, each with unknown parameters (p_i, q_i) drawn from some unknown distribution on $[0, 1]^2$. We observe $X_i \sim \text{Bin}(t, p_i)$ before an event and $Y_i \sim \text{Bin}(t, q_i)$ after the event. Provided these paired observations, $\{(X_i, Y_i)\}_{i=1}^N$, our goal is to accurately estimate the *distribution of the change in parameters*, $\delta_i := q_i - p_i$, over the population and properties of interest like the ℓ_1 -*magnitude of the change* with sparse observations ($t \ll N$). We provide *information theoretic lower bounds* on the error in estimating the distribution of change and the ℓ_1 -magnitude of change. Furthermore, we show that the following two step procedure achieves the optimal error bounds: first, estimate the full joint distribution of the paired parameters using the maximum likelihood estimator (MLE) and then estimate the distribution of change and the ℓ_1 -magnitude of change using the joint MLE. Notably, and perhaps surprisingly, these error bounds are of the same order as the minimax optimal error bounds for learning the *full* joint distribution itself (in Wasserstein-1 distance); in other words, estimating the magnitude of the change of parameters over the population is, in a minimax sense, as difficult as estimating the full joint distribution itself.

1 Introduction

The problem of estimating change in the parameters of interest over a population plays a key role in social sciences, epidemiology, medicine, biology, and sports analytics [Wil45, WW82, MP06, RLP⁺11, Bro08]. Consider an example where we want to assess the impact of a policy of administering free flu vaccines on the health of a population. Suppose for a large random sample of the population we have observations of whether or not an individual caught the flu in each year for 5 years before and 5 years after the policy is introduced. While it is hard to accurately estimate the change in the probability of contracting flu per individual with only 5 observations before and after, assessing the *change over the population* is crucial for the public health policy makers. Learning the distribution and the magnitude of change is extremely useful for downstream analysis of testing and estimating properties of the distribution of change, e.g., was there a change? How much did the parameters change? Was the overall change positive?

One approach to estimating the change is to first learn the distributions before and after the event of interest and then estimate the change between these learned distributions. However, this approach is not sufficient in a *paired observation setting*, where observations are available

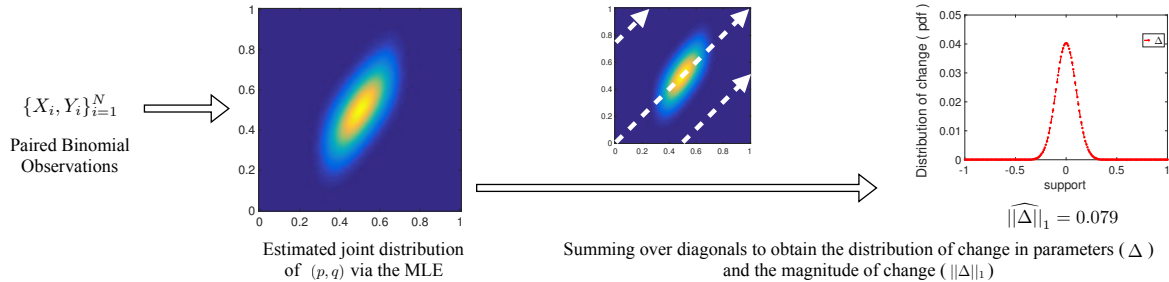


Figure 1: Illustration of optimal estimators for the distribution of change and the ℓ_1 -magnitude of change from paired Binomial observations obtained by using the estimated joint distribution of the parameters (p, q) .

before and after for each individual. For example, let half the population have parameter 0.25 and the other half 0.75 before an event. Suppose after the event the parameter increases by 0.5 for those with parameter 0.25 (before) and decreases by 0.5 for those with parameter 0.75 (before). Clearly, the distribution of change is supported on $\{-0.5, 0.5\}$ with half the mass on each point. However, the distribution of parameters before looks the same as that of after. In contrast, *paired estimation of the distribution of change* can overcome this to recover the distribution of change. Classically, this problem is studied using *paired difference test*, where the goal is to test whether the population mean or median changed after the event (e.g. introduction of a policy). *Paired t-test* is typically applied when the change is normally distributed. For other cases, e.g., Binomial observations as in our setting, *sign test* and *Wilcoxon's signed-rank test* [Wil45, LR06] are often applied. However, these tests can only answer the question of whether the population *mean* or *median* changed, but not how the changes distribute. Furthermore, when a classical paired test only detects change in the mean or the median, it does not capture the *effect size*, a measure of magnitude of the change [SF12].

Our Contributions: In this work, we consider the problem of estimating the distribution and the l_1 -magnitude of change in the paired Binomial observation setting and address the following question: *What is the sample complexity for estimating the distribution and magnitude of change?*

- We provide *information theoretic lower bounds* on the error in estimating the distribution of change in Wasserstein-1 distance (or the earth mover's distance) and the absolute error in estimating the ℓ_1 -magnitude of the change with sparse observations.
- We also show that estimators derived by using the maximum likelihood estimator (MLE) of the joint distribution of the paired observations achieve the minimax optimal error bounds in the sparse observation regime (Figure 1).

Perhaps surprisingly, our analysis shows that the minimax error bounds on estimating the distribution and the magnitude of the change, which is a functional of the distribution of change, is of the same order as the minimax error bound on estimating the full joint distribution of the paired observations over $[0, 1]^2$. That is, in a minimax sense, estimating the distribution of change in the parameters over the population and its ℓ_1 -magnitude (a functional of the distribution of change) is as difficult as estimating the *full* joint distribution of the parameters itself. We also evaluate the empirical performance of our estimators through extensive simulations and demonstrate the use of ℓ_1 -magnitude of change for hypothesis testing.

1.1 Problem Setup and Summary of Results

We can model individuals with parameters $\in [0, 1]$ in the Binomial setting as coins with biases equal to the parameters. Let R^* on $[0, 1]^2$ denote the true joint distribution of biases before and after the event. Let N be the number of coins. For each coin i , the before and after bias tuple (p_i, q_i) is drawn independently from R^* , i.e., $(p_i, q_i) \sim R^*$. We get to observe t tosses per coin before and t tosses after an event. Let $\{(X_i, Y_i)\}_{i=1}^N$ be the set of observations with $X_i \sim \text{Bin}(t, p_i)$ and $Y_i \sim \text{Bin}(t, q_i)$. Δ^* supported on $[-1, 1]$, denotes the distribution of change in the biases after the event compared to the biases before (defined formally later, see Equation (1)). $\|\Delta^*\|_1 := \int_{z=-1}^1 |z| d\Delta^*(z)$ denotes the l_1 -magnitude of the change. Our goal is to estimate Δ^* and its l_1 -magnitude $\|\Delta^*\|_1$ given the paired Binomial observations $\{(X_i, Y_i)\}_{i=1}^N$. We establish information theoretic lower bounds on how well we can estimate them over the population when $t \ll N$ (*sparse observations*). Furthermore, we construct estimators for the distribution and the magnitude of change in the biases over the population, denoted by $\widehat{\Delta}$ and $\|\widehat{\Delta}\|_1$ respectively, which achieve the information theoretic lower bounds and therefore are minimax optimal in the sparse regime. These estimators are simple and easily to compute, derived using the MLE of the joint distribution of (p, q) . Figure 1 shows an illustration of integrating over the diagonal and sub-diagonals of the joint distribution over $[0, 1]^2$ to obtain the distribution of change supported on $[-1, 1]$. These estimators are defined formally later in Definitions (2.1) and (2.2). We bound the error in the estimation of the distribution of change (under Wasserstein-1 distance, Definition 2.3) and the absolute error in estimating the l_1 -magnitude of change and in conjunction with the lower bounds show that,

Theorem 1.1. (Informal)

- When $t < c \log N$: $W_1(\Delta^*, \widehat{\Delta}) = \Theta_\alpha\left(\frac{1}{t}\right)$; $\left| \|\Delta^*\|_1 - \|\widehat{\Delta}\|_1 \right| = \Theta_\alpha\left(\frac{1}{t}\right)$.¹
- When $t \in [\Omega(\log N), \mathcal{O}(N^{1/8-\epsilon})]$: $W_1(\Delta^*, \widehat{\Delta}) = \Theta_\alpha\left(\frac{1}{\sqrt{t \log N}}\right)$; $\left| \|\Delta^*\|_1 - \|\widehat{\Delta}\|_1 \right| = \Theta_\alpha\left(\frac{1}{\sqrt{t \log N}}\right)$.

Note that $t < c \log N$ implies that the size of the population, $N \gtrsim \exp(t)$. When $t \in [\Omega(\log N), \mathcal{O}(N^{1/8-\epsilon})]$, then $\text{poly}(t) \lesssim N \lesssim \exp(t)$.

1.2 Related Works

Estimating the distribution of population of parameters: The MLE for estimating the unknown distribution of parameters in general nonparametric mixture model setting has been studied extensively in the literature [LC75, Tur76, Sim76, Lai78, Cre79, Lin83a, Lin83b, Böh89, LK92]. There have been several works that focus on the Binomial model [Lin83a, Lin83b, Woo99]. Most of the prior literature focused on understanding the geometry of the MLE, identifiability and uniqueness of the solutions, and optimality criteria of the MLE. Recently, Vinayak et. al [VKVK19] studied the Wasserstein error rate of the MLE in the one dimensional setting, and showed that the MLE achieves error $\mathcal{O}\left(\frac{1}{t}\right)$ in with $t = \mathcal{O}(\log N)$ coin tosses and $\mathcal{O}\left(\frac{1}{\sqrt{t \log N}}\right)$ in $t = \Omega(\log N)$ regime, which are minimax optimal in both cases. Li et. al [LRSS15] proposed a linear programming based estimator that achieves an error bound of $\mathcal{O}\left(\frac{1}{t}\right)$ when $t = \mathcal{O}(\log N)$ and a suboptimal error bound of $\mathcal{O}\left(\frac{1}{\sqrt{t}}\right)$ otherwise. Tian et. al [TKV17] proposed a method of moments estimator to estimate the unknown distribution of biases in the regime where the number of coin tosses is $t = \mathcal{O}(\log N)$. In the multivariate setting where the biases are drawn from a distribution supported on $[0, 1]^d$ with constant d , their estimator achieves Wasserstein error $\mathcal{O}\left(\frac{1}{t}\right)$ which is minimax optimal. The main weakness of this method of moments approach is that it fails to obtain the optimal rate when $t > c \log N$.

¹ $\Theta_\alpha(\cdot)$ hides $\log(1/\alpha)$ in the bound for it to hold with probability at least $1 - 2\alpha$.

We will comment on using the joint distribution learned using moment matching estimator in Remark 2.2.

Symmetric properties of discrete distributions: Estimating the ℓ_1 -magnitude of change is a special case of functional estimation problems where the goal is to estimate a certain functional of the parameters. Several recent works [OSVZ04, AOP09, ADM⁺10, VV11, VV13, JVHW15, WY15, VV16, WY16, OSW16, ADOS17, JHW18, HJW18] have focused on estimating symmetric functionals of discrete distributions, such as entropy and support size. Many of these works have observed that some functionals can be accurately estimated even in the setting where the number of samples is not sufficient to reliably learn the distribution. For example, Valiant and Valiant [VV11] showed that accurate estimation of the entropy of a discrete distribution with support size n only requires $O(\frac{n}{\log n})$ samples while learning the distribution itself requires at least $\Omega(n)$ samples.

PML framework: A series of works [AOP09, ADM⁺10, ADOS17] examined the *profile* or *pattern* maximum likelihood (PML) as a unifying framework for estimating symmetric properties of a discrete distribution. A limitation of this framework is that it only applies to scalar symmetric properties of discrete distributions. To the best of our knowledge, it is not known how to apply these approaches to estimating asymmetric properties, for examples ℓ_1 distance between two discrete distributions which is related to the setting we considered. Another limitation is that the near-optimal sample complexity of using the PML for estimating a function estimation relies on the existence of another stable estimator of the function, and the PML itself does not give the minimax rate. In contrast, we establish the optimality of using derived estimators for the distribution as well as ℓ_1 magnitude of change by providing matching information theoretic lower bounds for these estimation problems. Furthermore, the setting in our work is different from that of discrete distribution setting (as emphasized in the following paragraph).

Estimating ℓ_1 -distance: Jiao et. al [JHW18] considers the problem of estimating the ℓ_1 -distance between two discrete distributions. There are two major differences between this work and the results in [JHW18]. First, Jiao et. al [JHW18] considered the discrete distribution setting where the parameters sum to 1 and the observations follow Poisson distribution. In contrast, in the setting considered in this work, the observations are Binomial and the parameters need not sum to 1. Second, on the algorithmic side, the estimator in Jiao et. al [JHW18] is based on a multivariate piece-wise polynomial function that is specifically designed to approximate the function $|y - x|$. In contrast, our estimator for the ℓ_1 -magnitude of change is constructed by using the MLE solution of the joint distribution, which can be generalized to estimate other functionals.

2 Main Results

In this section we present our main results. We start with some notations and definitions that help us in presenting the formal results. Recall that we get N paired binomial observations, $\{(X_i, Y_i)\}_{i=1}^N$, where $X_i \sim \text{Bin}(t, p_i)$ and $Y_i \sim \text{Bin}(t, q_i)$ are the number of heads seen before and after by coin i . The binomial parameters (p_i, q_i) are drawn i.i.d. from the joint distribution R^* . Let $n_{u,v} := \sum_{i=1}^N \mathbf{1}\{(X_i, Y_i) = (u, v)\}$, denote the number of coins that sees tuple (u, v) , that is, turned up u heads before and v heads after out of t tosses. Let $h_{u,v} = n_{u,v}/N$ denote the fraction of coins that show (u, v) observation tuple. The matrix $H \in \mathbb{R}_{\geq 0}^{(t+1) \times (t+1)}$, with $H_{u,v} = h_{u,v}$ is the *fingerprint* matrix for the joint observation.

Definition 2.1. (Distribution of Change) Define random variable $\delta = q - p$ where $(p, q) \sim R^*$. Let Δ^* be the distribution of δ , i.e., the true distribution of change in the biases over $[-1, 1]$. Then Δ^* can be defined using R^* as follows, for $w \in [-1, 1]$,

$$\Delta^*(w) := \int_{z=-1}^w \int_{x=\max\{0, -z\}}^{\min\{1, 1-z\}} r^*(x, x+z) dx dz =: \int_{z=-1}^w \delta^*(z) dz, \quad (1)$$

where r^* and δ^* denotes the probability density functions (pdf) of the distributions R^* and Δ^* respectively.

Definition 2.2. (Magnitude of Change) The l_1 -magnitude of the change in the biases, denoted by $\|\Delta^*\|_1$, is given by the following equation,

$$\|\Delta^*\|_1 = \int_{x=0}^1 \int_{y=0}^1 |y-x| r^*(x, y) dx dy = \int_{z=-1}^1 |z| \delta^*(z) dz. \quad (2)$$

Goal: Given the paired observations $\{(X_i, Y_i)\}_{i=1}^N$, our goal is to estimate the distribution Δ^* and its l_1 -magnitude $\|\Delta^*\|_1$.

The accuracy of the estimator of the distribution of change is measured in *Wasserstein-1 distance* or the *Earth Mover's Distance* (EMD) defined below:

Definition 2.3. (Wasserstein-1 distance, dual definition [KR58]) The Wasserstein-1 distance or the Earth Mover's Distance (EMD) between two distributions P and Q supported on \mathcal{M} is defined as,

$$W_1(P, Q) = \sup_{f \in \text{Lip}(1)} \int_{\mathcal{M}} f(x) (dP(x) - dQ(x)), \quad (3)$$

where $\text{Lip}(1)$ denotes the set of all Lipschitz-1 functions over \mathcal{M} .

2.1 Information Theoretic Lower Bounds

The following theorem provides information theoretic lower bounds on the error in estimating the l_1 -magnitude of change and the distribution of change under the Wasserstein-1 distance.

Theorem 2.1. Let R be a distribution over $[0, 1]^2$. Let $\mathbf{S} := \{(X_i, Y_i)\}_{i=1}^N$ be random variables with $X_i \sim \text{Binomial}(t, p_i)$, $Y_i \sim \text{Binomial}(t, q_i)$ where (p_i, q_i) is drawn independently from R . Define random variable $\delta = q - p$ where $(p, q) \sim R$, and let Δ be the distribution of δ .

1. Define $\|\Delta_R\|_1 = \mathbb{E}_{(p,q) \sim R}[|q-p|]$. Let ζ be an estimator that maps \mathbf{S} to a real value $\zeta(\mathbf{S})$. Then for every t, N s.t. $t \leq \frac{N^{2(\epsilon^4-1)}}{36}$,

$$\inf_{\zeta} \sup_R \mathbb{E}[|\|\Delta_R\|_1 - \zeta(\mathbf{S})|] = \Omega\left(\frac{1}{t}\right) \vee \Omega\left(\frac{1}{\sqrt{t \log N}}\right).$$

2. Let g be an estimator that maps \mathbf{S} to a probability distribution $g(\mathbf{S})$ supported on $[-1, 1]$. Then for every t, N s.t. $t \leq \frac{N^{2(\epsilon^4-1)}}{36}$,

$$\inf_g \sup_R \mathbb{E}[W_1(\Delta, g(\mathbf{S}))] = \Omega\left(\frac{1}{t}\right) \vee \Omega\left(\frac{1}{\sqrt{t \log N}}\right),^2$$

² $a \vee b := \max\{a, b\}$.

In the setting in Theorem 2.1, the following proposition provides the information theoretic lower bound on the Wasserstein-1 error in estimating the joint distribution of (p, q) .

Corollary 2.1. (Lower bound on estimating the joint distribution.) Let ρ be an estimator that maps \mathbf{S} to a probability distribution $\rho(\mathbf{S})$ supported on $[0, 1]^2$. Then for every t, N s.t. $t \leq \frac{N^2(e^4-1)}{36}$,

$$\inf_{\rho} \sup_R \mathbb{E}[W_1(R, \rho(\mathbf{S}))] = \Omega\left(\frac{1}{t}\right) \vee \Omega\left(\frac{1}{\sqrt{t \log N}}\right).$$

2.2 Estimators for the Distribution and the Magnitude of Change

Our estimators for the distribution and the magnitude of change are derived from the MLE of the joint distribution of the parameters (p, q) given the observations. Given the paired Binomial observations $\{(X_i, Y_i)\}_{i=1}^N$, the MLE of the joint distribution of (p, q) denoted by \hat{R}_{mle} is as follows,

$$\begin{aligned} \hat{R}_{\text{mle}} &\in \arg \max_{S \in \mathcal{S}} \frac{1}{N} \log \Pr(\{(X_i, Y_i)\}_{i=1}^N | S), \\ &= \arg \max_{S \in \mathcal{S}} \frac{1}{N} \sum_{i=1}^N \log \Pr((X_i, Y_i) | S), \\ &= \arg \max_{S \in \mathcal{S}} \frac{1}{N} \sum_{u=0}^t \sum_{v=0}^t n_{u,v} \log \Pr((u, v) | S), \end{aligned} \quad (4)$$

where \mathcal{S} is the set of all distributions on $[0, 1]^2$. Note that $\Pr((u, v) | S)$ can be written as,

$$\int_0^1 \int_0^1 \binom{t}{u} x^u (1-x)^{t-u} \binom{t}{v} y^v (1-y)^{t-v} dS(x, y) =: \mathbb{E}_S[h_{u,v}],$$

where $\mathbb{E}_S[h_{u,v}]$ is the expected fingerprint, i.e., the fraction of the population that sees the tuple (u, v) under the distribution S . Therefore, the MLE of the joint distribution of (p, q) is written in terms of the expected fingerprints as follows,

$$\hat{R}_{\text{mle}} \in \arg \max_{S \in \mathcal{S}} \sum_{u=0}^t \sum_{v=0}^t h_{u,v} \log \mathbb{E}_S[h_{u,v}]. \quad (5)$$

We note that the objective function is concave in the distribution S and it is strictly concave in the expected fingerprints. Therefore, while there can be multiple distributions maximizing the objective function (5), the expected fingerprints under them will be the same.

Estimator for the distribution of change Δ^* over the population is denoted by $\hat{\Delta}$ and is derived using the joint MLE is as follows, for $w \in [-1, 1]$,

$$\hat{\Delta}(w) := \int_{z=-1}^w \int_{x=\max\{0, -z\}}^{\min\{1, 1-z\}} \hat{r}_{\text{mle}}(x, x+z) dx dz, \quad (6)$$

where \hat{r}_{mle} is the probability density function of \hat{R}_{mle} .

Estimator for the magnitude of change $\|\Delta^*\|_1$ over the population, denoted by $\|\hat{\Delta}\|_1$, is derived from the joint MLE as follows:

$$\|\hat{\Delta}\|_1 := \int_{x=0}^1 \int_{y=0}^1 |y-x| \hat{r}_{\text{mle}}(x, y) dx dy = \int_{z=-1}^1 |z| \hat{\delta}(z) dz, \quad (7)$$

where $\hat{\delta}$ is the probability density function of $\hat{\Delta}$.

With the setting in Section 1.1, the following theorem provides upper bounds on the accuracy of the estimators above:

Theorem 2.2. (Estimation of the Distribution of Change) There exists $\epsilon > 0$ such that with probability at least $1 - 2\alpha$,

1. The Wasserstein-1 error incurred by the estimator in Equation (6) can be upper bounded as follows,

$$W_1(\Delta^*, \hat{\Delta}) = \begin{cases} \mathcal{O}_\alpha\left(\frac{1}{t}\right), & \text{when } t = \mathcal{O}(\log N), \\ \mathcal{O}_\alpha\left(\frac{1}{\sqrt{t \log N}}\right), & \text{when } t \in [\Omega(\log N), \mathcal{O}(N^{1/8-\epsilon})]. \end{cases}$$

2. The error incurred by the estimator in Equation (7) can be upper bounded as,

$$\left| \|\Delta^*\|_1 - \|\hat{\Delta}\|_1 \right| = \begin{cases} \mathcal{O}_\alpha\left(\frac{1}{t}\right), & \text{when } t = \mathcal{O}(\log N), \\ \mathcal{O}_\alpha\left(\frac{1}{\sqrt{t \log N}}\right), & \text{when } t \in [\Omega(\log N), \mathcal{O}(N^{1/8-\epsilon})]. \end{cases}$$

Along with the lower bound in Theorem 2.1, the above upper bounds show that our estimators for the distribution of change and the l_1 -magnitude of change given in Equation (6) and Equation (7) respectively are optimal in the sparse observation regime.

With the setting in Section 1.1, the following proposition bounds the estimation error for the joint MLE.

Proposition 2.1. (MLE for the Joint Distribution) There exists $\epsilon > 0$ such that with probability at least $1 - 2\alpha$,

$$W_1(R^*, \hat{R}_{\text{mle}}) = \begin{cases} \mathcal{O}_\alpha\left(\frac{1}{t}\right), & \text{when } t = \mathcal{O}(\log N), \\ \mathcal{O}_\alpha\left(\frac{1}{\sqrt{t \log N}}\right), & \text{when } t \in [\Omega(\log N), \mathcal{O}(N^{1/8-\epsilon})]. \end{cases}$$

Along with the lower bound in Proposition 2.1, the above upper bound shows that the MLE is minimax optimal for estimating the joint distribution in the sparse observation regime. Furthermore, we note that the upper bound on the error in estimating the joint distribution of the parameters (p, q) is of the same order as the lower bounds on error in estimating the distribution and magnitude of change. Even though the l_1 -magnitude of change is a scalar value, estimating it is no easier than estimating the whole joint distribution itself.

Remark 2.1. The standard **naive plug-in empirical estimator** which uses the empirical histogram of the observations, $\left\{\left(\frac{X_i}{t}, \frac{Y_i}{t}\right)\right\}_{i=1}^N$, would be sub-optimal in estimating the quantities of interest in the sparse observation regime since it incurs an error of $\Theta\left(\frac{1}{\sqrt{t}}\right)$.

Remark 2.2. Moment based estimation: Our estimators for the distribution and l_1 -magnitude of change in Equations (6) and (7) can be easily modified by replacing r_{mle} with the estimated joint distribution of (p, q) using moment matching [TKV17] to obtain corresponding estimators derived by moment matching. These derived estimators would achieve the minimax optimal error rate of $\Theta\left(\frac{1}{t}\right)$ when $t < c \log N$ (i.e., when $N \gtrsim \exp(t)$). However, they fail to achieve optimal error rates when $t > c \log N$ due to higher variance in larger moments.

Remark 2.3. Estimating change with only observation of change: Binomial distribution is not transnational invariant, therefore, observing only the change, $Z = Y - X$, is not enough to estimate the distribution of change in parameters, $\delta = q - p$. Unbiased estimate of the average change in the parameters, $E[\delta]$, can be obtained using only the difference,

$\frac{1}{N} \sum_{i=1}^N \frac{Z_i}{t}$. However, unbiased estimates of the higher moments would need the joint moments. For example, an unbiased estimate of the second moment $E[\delta^2]$ would be

$$\widehat{E}[\delta^2] := \frac{1}{N} \sum_{i=1}^N \left(\frac{\binom{X_i}{2}}{\binom{t}{2}} + \frac{\binom{Y_i}{2}}{\binom{t}{2}} - 2 \frac{X_i Y_i}{t^2} \right),$$

which involves all the estimated joint second moments of (p, q) . In contrast, a naive empirical estimate using just the change Z gives rise to a biased estimate, $\frac{1}{N} \sum_{i=1}^N \left(\frac{Y_i - X_i}{t} \right)^2 = \frac{1}{N} \sum_{i=1}^N \delta_i^2 + \frac{1}{N} \sum_{i=1}^N \frac{q_i(1-q_i) + p_i(1-p_i)}{t}$. For small t , the second term is not zero unless all (p_i, q_i) are 0 or 1.

3 Proof Sketches for Main Results

In this section we provide a brief sketch of the proofs the main results.

Proof sketches for lower bounds: To show the information theoretic lower bound on the error in estimating the ℓ_1 -magnitude of change, we use the following proposition.

Proposition 3.1. For any integer $s > 0$, \exists a pair of distributions P and Q supported on $[a, b]$ where $0 < a < b$, such that their first s moments match, and

$$\left| \int_a^b \left| x - \frac{a+b}{2} \right| P(x) dx - \int_a^b \left| x - \frac{a+b}{2} \right| Q(x) dx \right| = \Omega \left(\frac{b-a}{s} \right).$$

Using Proposition 3.1 we construct two distributions P^* and Q^* supported on $[0, 1]$ such that their first t moments match and,

$$\int_0^1 \left| x - \frac{1}{2} \right| P(x) dx - \int_0^1 \left| x - \frac{1}{2} \right| Q(x) dx \geq \Omega \left(\frac{1}{t} \right).$$

Let $U = \delta_{1/2}$ be a distribution supported only at $1/2$. Define 2-dimensional distribution $P = U \times P^*, Q = U \times Q^*$. We first show that the difference in the ℓ_1 -magnitude of change of these two joint distributions is lower bounded by $1/t$. We then argue that if the tuple (p, q) are drawn from P or Q is information theoretically not distinguishable from the binomial observations (X, Y) . To show the lower bound for the case when $t > c \log N$, we construct distributions similar to above supported on a small interval around $1/2$ whose first $\log N$ moments match while the ℓ_1 -magnitudes of change differ by $\frac{1}{\sqrt{t \log N}}$. The lower bounds on the error in estimating the ℓ_1 -magnitude of change imply the lower bounds for estimating the distribution of change in the parameters as well.

Proof Sketches for the Upper Bounds: Using the definition of Δ^* (Equation (1)) and changing variables from $(x, z := y - x)$ to (x, y) , the Wasserstein-1 distance between the true distribution of change Δ^* and the estimated change distribution $\widehat{\Delta}$ can be written as follows,

$$W_1(\Delta^*, \widehat{\Delta}) = \sup_{f \in \text{Lip}(1)} \int_0^1 \int_0^1 f(y-x) \left(dR^*(x, y) - d\widehat{R}_{\text{mle}}(x, y) \right). \quad (8)$$

The estimation error for the magnitude of change is given by:

$$\left| \|\Delta^*\|_1 - \|\widehat{\Delta}\|_1 \right| = \left| \int_0^1 \int_0^1 |y-x| (dR^*(x, y) - d\widehat{R}_{\text{mle}}(x, y)) \right|. \quad (9)$$

By the definition the Wasserstein-1 distance between the true joint distribution R^* and \widehat{R}_{mle} is as follows,

$$W_1(R^*, \widehat{R}_{\text{mle}}) = \sup_{f \in \text{Lip}(1)} \int_0^1 \int_0^1 f(x, y) (dR^*(x, y) - d\widehat{R}_{\text{mle}}(x, y)). \quad (10)$$

Noting that $|y - x|$ is a Lipschitz-1 function, all the estimator error expressions in Equations (8), (9) and (10) involve Lipschitz-1 functions. It follows from the definition that,

$$\left| \|\Delta^*\|_1 - \|\widehat{\Delta}\|_1 \right| \leq W_1(\Delta^*, \widehat{\Delta}) \leq W_1(R^*, \widehat{R}_{\text{mle}}).$$

The following bound holds for any $f \in \text{Lip}(1)$ supported on $[0, 1]^2$ and any polynomial approximation \widehat{f} of f on $[0, 1]^2$.

$$\begin{aligned} & \left| \int_{x=0}^1 \int_{y=0}^1 f(x, y) (dR^*(x, y) - d\widehat{R}_{\text{mle}}(x, y)) \right| \\ &= \left| \int_0^1 \int_0^1 (f(x, y) - \widehat{f}(x, y)) (dR^* - d\widehat{R}_{\text{mle}}) + \int_0^1 \int_0^1 \widehat{f}(x, y) (dR^* - d\widehat{R}_{\text{mle}}) \right|, \\ &\leq 2\|f - \widehat{f}\|_\infty + \left| \mathbb{E}_{R^*}[\widehat{f}] - \mathbb{E}_{\widehat{R}_{\text{mle}}}[\widehat{f}] \right|, \end{aligned} \quad (11)$$

where $\|f - \widehat{f}\|_\infty := \max_{(x, y) \in [0, 1]^2} |f(x, y) - \widehat{f}(x, y)|$ is the uniform polynomial approximation error.

Let \widehat{f}_t be the 2-d Bernstein polynomial of degree $2t$ defined by,

$$\widehat{f}_t(x, y) = \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} B_u^t(x) B_v^t(y), \quad (12)$$

where $B_u^t(x) := \binom{t}{u} x^u (1-x)^{t-u}$ is the u -th Bernstein polynomial of degree t . Using this in Equation (11), the EMD between the true joint distribution of (p, q) , R^* and the MLE estimator \widehat{R}_{mle} can be bounded as,

$$W_1(R^*, \widehat{R}_{\text{mle}}) \leq \sup_{f \in \text{Lip}(1)} \left\{ 2\|f - \widehat{f}\|_\infty + \left| \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} \left(\mathbb{E}_{R^*}[h_{uv}] - \mathbb{E}_{\widehat{R}_{\text{mle}}}[h_{uv}] \right) \right| \right\}, \quad (13)$$

This upper bound also holds for Equations (8) and (9). The second term in the bound above Equation (13) can be bounded by the following expression,

$$\left| \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} \left(\mathbb{E}_{R^*}[h_{uv}] - \mathbb{E}_{\widehat{R}_{\text{mle}}}[h_{uv}] \right) \right| \stackrel{(a)}{\leq} \mathcal{O} \left(\max_{u, v} |\kappa_{uv}| \sqrt{\frac{(t+1)^2}{2N} \log \frac{4N}{(t+1)^2} + \frac{1}{N} \log \frac{3e}{\alpha}} \right),$$

where the inequality (a) uses triangle inequality followed by concentration of fingerprints and bound on the error in 2-d fingerprint due to MLE solution. Note that the term $\max_{(u, v)} |\kappa_{uv}|$ appears in the bound above. Therefore, it is crucial to bound the uniform approximation error $\|f - \widehat{f}\|_\infty$ while keeping the coefficients of the approximation (12) from being too large. Polynomial approximation of Lipschitz-1 functions on $[0, 1]^2$ using Bernstein polynomials play a crucial role in obtaining a tight bound that match the information theoretic lower bounds. We extend the analysis of [VKVK19] to obtain these bounds. The details are available in the appendix.

4 Empirical Performance

Solving the MLE: Recall the objective function for the joint MLE in Equation 5. The objective function is concave in the distribution S and strictly concave in the expected fingerprints. Furthermore, the set \mathcal{S} of all distributions over $[0, 1]^2$ is convex. Thus the MLE of the joint distribution in this setting is an optimization problem of maximizing a concave function over a convex set. We solve this optimization problem on a 100×100 uniform grid on $[0, 1]^2$. The objective function of the joint MLE on a $m \times m$ grid on $[0, 1]^2$ is as follows,

$$F(S) := \sum_{u=0}^t \sum_{v=0}^t h_{u,v}^{\text{obs}} \log \left(\sum_{i=0}^m \sum_{j=0}^m \binom{t}{u} \left(\frac{i}{m}\right)^u \left(1 - \frac{i}{m}\right)^{t-u} \binom{t}{v} \left(\frac{j}{m}\right)^v \left(1 - \frac{j}{m}\right)^{t-v} S_{ij} \right).$$

So, the MLE objective is to maximize $F(S)$ for S on the simplex, that is, $S_{ij} \geq 0$ and $\sum_i \sum_j S_{ij} = 1$. Let $\text{vec}(S)$ denote vectorized version of the matrix S by stacking the columns and $\nabla_S F$ denote the gradient of F at S . We use exponential gradient descent algorithm (or equivalently, mirror descent with KL divergence) to optimize this objective as described below:

1. Input: Observed fingerprints $\{h_{u,v}^{\text{obs}}\}_{u,v=0}^t$, learning rate η , grid size m , tolerance ε , maximum iterations: *MaxIters*.
2. Start with $S(0)$ being Uniform distribution on $[0, 1]^2$, i.e., $[S(0)]_{ij} = \frac{1}{(m+1)^2}$.
3. For $\text{iter} = 0, 1, 2, \dots, \text{MaxIters}$
 - Update: $[S(\text{iter} + 1)]_{ij} = [S(\text{iter})]_{ij} \exp \{ \eta [\nabla_{S(\text{iter})} F]_{ij} \}$, for $i, j = 0, \dots, m$, where
 - Terminate if $\|\text{vec}(S(\text{iter} + 1)) - \text{vec}(S(\text{iter}))\|_2 < \varepsilon$.

We run the above algorithm with a learning rate 0.5, stopping criteria of $\|\text{vec}(s_{\text{iter}}) - \text{vec}(s_{\text{iter}-1})\|_2 < 5 \times 10^{-7}$ and maximum iterations set to 2×10^4 . Estimates $\hat{\Delta}$ and $\|\hat{\Delta}\|_1$ are obtained using the discrete joint MLE solution. All the experiments were run on Matlab[®]R2017b.

Estimation of the distribution of change: Figure 2 shows the empirical performance of the MLE for estimating joint distribution of (p, q) given the paired before and after observations (second row) and the derived estimate of the distribution of change (third row) for a variety of true joint distributions (first row). The number of coins is $N = 10^6$ and the number of tosses before and after is $t = 5$ for all the examples in the figure.

Estimation of the magnitude of change Our estimation gets better as the number of observations per coin increases as the theory suggests. Figure 3 shows the error in estimating the distribution and the l_1 -magnitude of change for varying number of tosses $t \in \{2, 3, 4, 5, 6, 8, 10, 12, 15, 20, 25, 30, 50, 100\}$ for a variety of change distributions. The number of coins is $N = 10^6$. Each experiment is run 10 times. For all the figures presented, the distribution of biases before was truncated Gaussian on $[0.3, 0.7]$ with $\mu = 0.5$ and $\sigma = 0.1$. For comparison we have plotted estimators using the joint distribution estimated using the MLE (in red), the moment matching estimator (in green; also implemented using exponential gradient descent) and the naive empirical estimator (in blue). We observed the same trend for other distributions for biases before, e.g., 3-spikes of equal mass at $\{0.3, 0.5, 0.7\}$ and for the uniform distribution over $[0.3, 0.7]$.

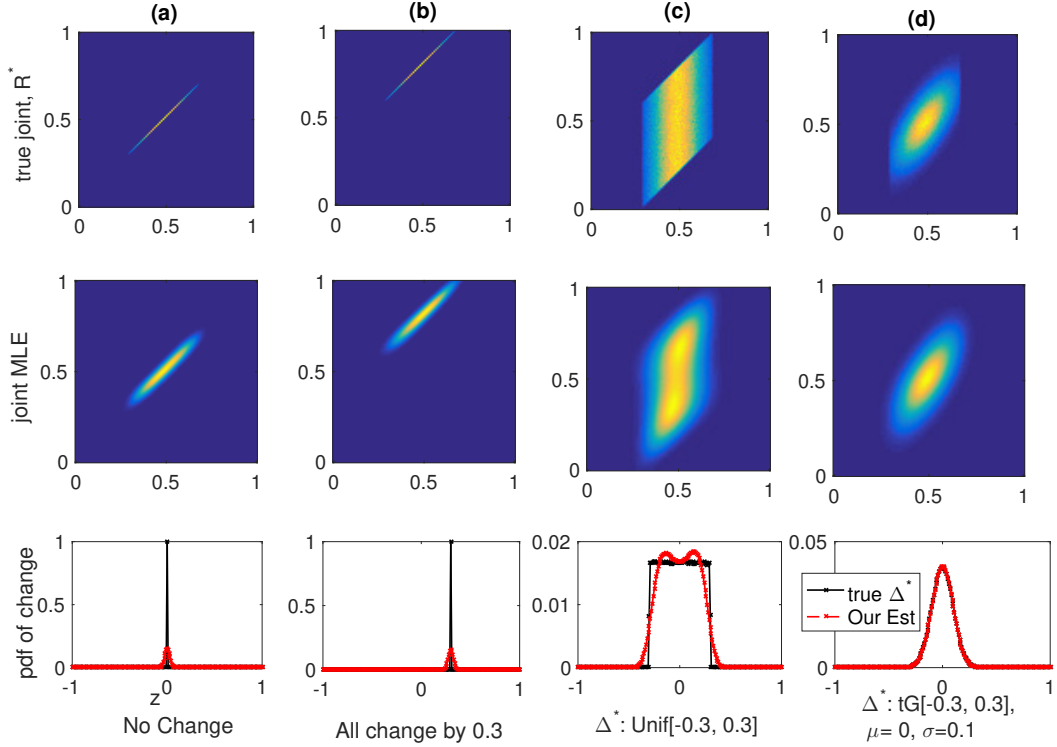


Figure 2: Empirical performance of the MLE in estimating the joint distribution (second row) and our estimator for distribution of change derived from the joint MLE (third row) for various true distributions (first row): (a) no change from before to after, (b) all the biases are increased by 0.3, (c) trapezoid joint distribution where the change distribution is uniform on $[-0.3, 0.3]$ and (d) truncated Gaussian (tG) joint, with change distribution also a tG supported on $[-0.3, 0.3]$ with mean $\mu = 0$ and standard deviation $\sigma = 0.1$. $N = 1e6$, $t = 5$.

Paired test using the l_1 -magnitude of change: l_1 -magnitude of change in the biases is a natural test statistic for nonparametric paired hypothesis testing. The null hypothesis is that there was no change in the biases from before to after. That is, under the null, the change distribution is a spike at 0 and consequently the l_1 -magnitude of change is 0. Commonly used nonparametric paired difference tests like the sign test and the Wilcoxon's signed rank test are designed to test for change in mean or median. If the change in the biases in our setting

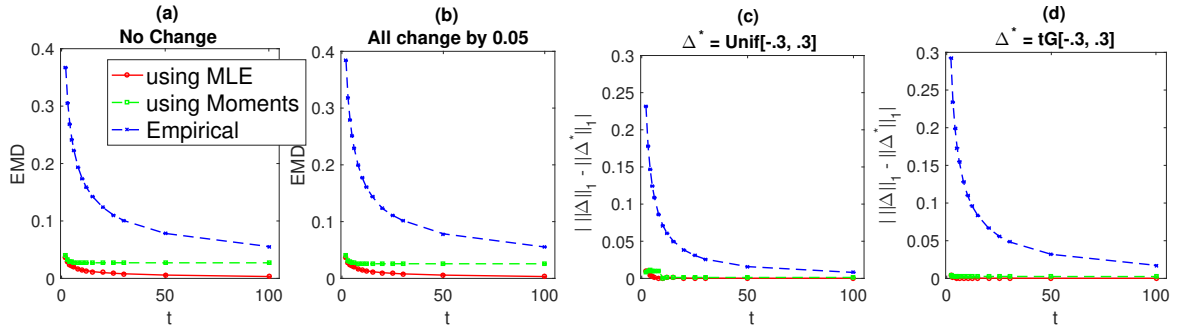


Figure 3: Error in estimating the distribution of change (measured in EMD) for various estimators as a function of varying number of tosses t under (a) no change, (b) all biases increased by 0.05, and l_1 -magnitude of change for (c) Δ^* : $\text{Unif}[-0.3, 0.3]$, (d) Δ^* : $\text{tG}[-0.3, 0.3]$ with $\mu = 0$, $\sigma = 0.1$. $N = 1e6$. 10 experiments.




Figure 4: Estimated magnitude of change under the null hypothesis for $t = \{5, 10, 15, 20\}$ and for various before distributions: 1 spike at 0.5 (1sp), 3 spikes (3sp), uniform (unif) and truncated Gaussian (tG). 100 Experiments, $N=1e6$. Box edges show 25-th and 75-th percentile, whisker edges are $\pm 2.7\text{std}$ [Mat].

is symmetric, e.g., uniform or truncated Gaussian around 0, then both the sign test and the Wilcoxon’s signed rank test fail to reject the null hypothesis. In contrast, using $\|\widehat{\Delta}\|_1$ as a test statistic we can successfully reject null hypothesis even when the change is symmetric around 0. Furthermore, while the sign test and the Wilcoxon’s signed rank test can detect systematic shifts in change for large N , they do not provide a sense of how much change occurred. Often quantities like correlation, rank correlation and other metrics are estimated to infer the *effect size*. In contrast, using the l_1 -magnitude of change for testing provides a transparent and natural quantitative measure of the effect size. Figure 4 shows the $\|\widehat{\Delta}\|_1$ for various joint distributions that lead to the null hypothesis (no change) for $t = \{5, 10, 15, 20\}$ and $N = 10^6$. For a given t and N , as long as the the l_1 -magnitude of change is higher than the critical value (that can be picked via simulations by setting desired level of false positives admissible), hypothesis test using $\|\widehat{\Delta}\|_1$ can reject the null hypothesis.

Acknowledgements

We thank our colleagues Jennifer Brennan, Maru Cabrera, Lalit Jain, John Thickstun and Jesse Thomason for the helpful discussions and comments. We also thank anonymous reviewers for their feedback and suggestions.

References

- [ADM⁺10] Jayadev Acharya, Hirakendu Das, Hosein Mohimani, Alon Orlitsky, and Shengjun Pan. Exact calculation of pattern probabilities. In *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*, pages 1498–1502. IEEE, 2010.

- [ADOS17] Jayadev Acharya, Hirakendu Das, Alon Orlitsky, and Ananda Theertha Suresh. A unified maximum likelihood approach for estimating symmetric properties of discrete distributions. In *International Conference on Machine Learning*, pages 11–21, 2017.
- [AOP09] Jayadev Acharya, Alon Orlitsky, and Shengjun Pan. Recent results on pattern maximum likelihood. In *Networking and Information Theory, 2009. ITW 2009. IEEE Information Theory Workshop on*, pages 251–255. IEEE, 2009.
- [BBL02] Thomas Bagby, Len Bos, and Norman Levenberg. Multivariate simultaneous approximation. *Constructive approximation*, 18(4):569–577, 2002.
- [Böh89] Dankmar Böhning. Likelihood inference for mixtures: geometrical and other constructions of monotone step-length algorithms. *Biometrika*, 76(2):375–383, 1989.
- [Bro08] Lawrence D Brown. In-season prediction of batting averages: A field test of empirical bayes and bayes methodologies. *The Annals of Applied Statistics*, pages 113–152, 2008.
- [Cre79] Noel Cressie. A quick and easy empirical bayes estimate of true scores. *Sankhyā: The Indian Journal of Statistics, Series B*, pages 101–108, 1979.
- [CT12] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [HJW18] Yanjun Han, Jiantao Jiao, and Tsachy Weissman. Local moment matching: A unified methodology for symmetric functional estimation and distribution estimation under wasserstein distance. *arXiv preprint arXiv:1802.08405*, 2018.
- [JHW18] Jiantao Jiao, Yanjun Han, and Tsachy Weissman. Minimax estimation of the l_1 distance. *IEEE Transactions on Information Theory*, 2018.
- [JVHW15] Jiantao Jiao, Kartik Venkat, Yanjun Han, and Tsachy Weissman. Minimax estimation of functionals of discrete distributions. *IEEE Transactions on Information Theory*, 61(5):2835–2885, 2015.
- [KR58] L. V. Kantorovich and G. S. Rubinstein. On a space of completely additive functions. *Vestnik Leningrad. Univ*, 13(7):52–59, 1958.
- [Lai78] Nan Laird. Nonparametric maximum likelihood estimation of a mixing distribution. *Journal of the American Statistical Association*, 73(364):805–811, 1978.
- [LC75] Frederic M Lord and Noel Cressie. An empirical bayes procedure for finding an interval estimate. *Sankhyā: The Indian Journal of Statistics, Series B*, pages 1–9, 1975.
- [Lin83a] Bruce G Lindsay. The geometry of mixture likelihoods: a general theory. *The annals of statistics*, pages 86–94, 1983.
- [Lin83b] Bruce G Lindsay. The geometry of mixture likelihoods, part ii: the exponential family. *The Annals of Statistics*, 11(3):783–792, 1983.
- [LK92] Mary L Lesperance and John D Kalbfleisch. An algorithm for computing the nonparametric mle of a mixing distribution. *Journal of the American Statistical Association*, 87(417):120–126, 1992.

- [LR06] Erich L Lehmann and Joseph P Romano. *Testing statistical hypotheses*. Springer Science & Business Media, 2006.
- [LRSS15] Jian Li, Yuval Rabani, Leonard J Schulman, and Chaitanya Swamy. Learning arbitrary statistical mixtures of discrete distributions. In *Proceedings of the forty-seventh annual ACM symposium on Theory of Computing*, pages 743–752. ACM, 2015.
- [Mat] Matlab boxplot. <https://www.mathworks.com/help/stats/boxplot.html>. Accessed: 2019-05-22.
- [MJT⁺18] Jay Mardia, Jiantao Jiao, Ervin Tánčzos, Robert D Nowak, and Tsachy Weissman. Concentration inequalities for the empirical distribution. *arXiv preprint arXiv:1809.06522*, 2018.
- [MP06] Chaya S Moskowitz and Margaret S Pepe. Comparing the predictive values of diagnostic tests: sample size and analysis for paired study designs. *Clinical trials*, 3(3):272–279, 2006.
- [OSVZ04] Alon Orlitsky, Narayana P Santhanam, Krishnamurthy Viswanathan, and Junan Zhang. On modeling profiles instead of values. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pages 426–435. AUAI Press, 2004.
- [OSW16] Alon Orlitsky, Ananda Theertha Suresh, and Yihong Wu. Optimal prediction of the number of unseen species. *Proceedings of the National Academy of Sciences*, 113(47):13283–13288, 2016.
- [Rab03] Abedallah Rababah. Transformation of chebyshev–bernstein polynomial basis. *Computational Methods in Applied Mathematics Comput. Methods Appl. Math.*, 3(4):608–622, 2003.
- [RLP⁺11] Bob Roozenbeek, Hester F Lingsma, Pablo Perel, Phil Edwards, Ian Roberts, Gordon D Murray, Andrew IR Maas, Ewout W Steyerberg, et al. The added value of ordinal analysis in clinical trials: an example in traumatic brain injury. *Critical Care*, 15(3):R127, 2011.
- [SF12] Gail M Sullivan and Richard Feinn. Using effect size—or why the p value is not enough. *Journal of graduate medical education*, 4(3):279–282, 2012.
- [Sim76] Leopold Simar. Maximum likelihood estimation of a compound poisson process. *The Annals of Statistics*, pages 1200–1209, 1976.
- [TKV17] Kevin Tian, Weihao Kong, and Gregory Valiant. Learning populations of parameters. In *Advances in Neural Information Processing Systems*, pages 5778–5787, 2017.
- [Tur76] Bruce W Turnbull. The empirical distribution function with arbitrarily grouped, censored and truncated data. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 290–295, 1976.
- [VKVK19] Ramya Korlakai Vinayak, Weihao Kong, Gregory Valiant, and Sham M Kakade. Maximum likelihood estimation for learning populations of parameters. In *International Conference on Machine Learning*, 2019.

- [VV11] Gregory Valiant and Paul Valiant. Estimating the unseen: an $n/\log(n)$ -sample estimator for entropy and support size, shown optimal via new clts. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 685–694. ACM, 2011.
- [VV13] Paul Valiant and Gregory Valiant. Estimating the unseen: improved estimators for entropy and other properties. In *Advances in Neural Information Processing Systems*, pages 2157–2165, 2013.
- [VV16] Gregory Valiant and Paul Valiant. Instance optimal learning of discrete distributions. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 142–155. ACM, 2016.
- [Wil45] Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945.
- [Woo99] G. R. Wood. Binomial mixtures: geometric estimation of the mixing distribution. *The Annals of Statistics*, 27(5):1706–1721, 1999.
- [WW82] Sholom Wacholder and Clarice Ring Weinberg. Paired versus two-sample design for a clinical trial of treatments with dichotomous outcome: Power considerations. *Biometrics*, pages 801–812, 1982.
- [WY15] Yihong Wu and Pengkun Yang. Chebyshev polynomials, moment matching, and optimal estimation of the unseen. *arXiv preprint arXiv:1504.01227*, 2015.
- [WY16] Yihong Wu and Pengkun Yang. Minimax rates of entropy estimation on large alphabets via best polynomial approximation. *IEEE Transactions on Information Theory*, 62(6):3702–3720, 2016.

A Proofs of Minimax Lower Bounds

We will first prove the following lemma that will be used in the proof of information theoretic lower bound on the error in estimating the ℓ_1 -magnitude of the change.

Proposition A.1. For any positive integer s , there exists a pair of distributions P and Q supported on $[a, b]$ where $0 < a < b$ such that P and Q have identical first s moments, and $|\int_a^b |x - \frac{a+b}{2}|P(x)dx - \int_a^b |x - \frac{a+b}{2}|Q(x)dx| = \Omega(\frac{b-a}{s})$.

Proof. Our proof leverages Lemma 29 from [JHW18], which is restated as follows,

Proposition A.2 (Lemma 29 of [JHW18]). For any given even integer $s > 0$, there exist two probability measures P^* and Q^* on $[-1, 1]$ that satisfy the following conditions:

1. $\int_{-1}^1 P^*(x)x^l dx = \int_{-1}^1 Q^*(x)x^l dx$, for $l = 0, 1, 2, \dots, s$
2. $\int_{-1}^1 P^*(x)|x|dx - \int_{-1}^1 Q^*(x)|x|dx = 2E_s[|x|; [-1, 1]]$

where $E_s[|t|; [-1, 1]]$ is the distance in the uniform norm on $[-1, 1]$ from the absolute value function $|t|$ to the space poly_s . It is known that $E_s[|x|; [-1, 1]] = \Omega(1/s)$.

Given the pair of distributions P^*, Q^* constructed by Proposition A.2, we define distribution $P(x) = \frac{2}{b-a}P^*\left(\frac{x-(a+b)/2}{(b-a)/2}\right)$ and $Q(x) = \frac{2}{b-a}Q^*\left(\frac{x-(a+b)/2}{(b-a)/2}\right)$. Since the transformation is linear, the first s moments of P, Q are identical. Further, $\int_a^b P(x)|x - (a+b)/2|dx - \int_a^b Q(x)|x - (a+b)/2|dx = \Omega\left(\frac{b-a}{s}\right)$. \square

Minimax lower bound for estimating ℓ_1 -magnitude of change:

Theorem A.1. Let R be a distribution over $[0, 1] \times [0, 1]$. Let $\mathbf{S} := \{(X_i, Y_i)\}_{i=1}^N$ be random variables with $X_i \sim \text{Binomial}(t, p_i)$, $Y_i \sim \text{Binomial}(t, q_i)$ where (p_i, q_i) is drawn independently from R . Define $\delta(R) = \mathbb{E}_{(p,q) \sim R}[|q - p|]$. Let ζ be an estimator that maps \mathbf{S} to a real value $\zeta(\mathbf{S})$. For every t, N s.t. $t \leq \frac{N^2(e^4-1)}{36}$, the following lower bound holds:

$$\inf_{\zeta} \sup_R \mathbb{E}[|\delta(R) - \zeta(\mathbf{S})|] = \Omega\left(\frac{1}{\sqrt{t \log N}}\right). \quad (14)$$

Proof. Let $U = \delta_{1/2}$ be a distribution supported only at $1/2$. We first apply Proposition A.1 to construct a pair of distributions P^* and Q^* supported on $\left[\frac{1}{2} - \sqrt{\frac{\log N}{t}}, \frac{1}{2} + \sqrt{\frac{\log N}{t}}\right]$ such that their first $L := e^4 \log N$ moments match, and

$$\int_{\frac{1}{2} - \sqrt{\frac{\log N}{t}}}^{\frac{1}{2} + \sqrt{\frac{\log N}{t}}} |x - \frac{1}{2}| P(x) dx - \int_{\frac{1}{2} - \sqrt{\frac{\log N}{t}}}^{\frac{1}{2} + \sqrt{\frac{\log N}{t}}} |x - \frac{1}{2}| Q(x) dx \geq \Omega\left(\frac{1}{\sqrt{t \log N}}\right).$$

Let $p \sim P^*$ and $q \sim Q^*$. Let $X \sim \text{Binomial}(t, p)$ and $X' \sim \text{Binomial}(t, q)$. When use the following proposition from [VKVK19],

Proposition A.3. (Proposition 4.4 of [VKVK19]) Let P and Q be two distributions, supported on $\left[\frac{1}{2} - \sqrt{\frac{\log N}{t}}, \frac{1}{2} + \sqrt{\frac{\log N}{t}}\right]$, whose first $L := e^4 \log N$ moments match. Let $p \sim P$ and $q \sim Q$. Let $X \sim \text{Binomial}(t, p)$ and $Y \sim \text{Binomial}(t, q)$. Then the total variation distance between X and Y satisfies,

$$\text{TV}(X, Y) \leq \frac{2\sqrt{t}}{Ne^4}.$$

Using Proposition 4.4 of [VKVK19], the total variation distance between X and X' satisfies,

$$\text{TV}(X, X') \leq \frac{2\sqrt{t}}{Ne^4}.$$

Define 2-dimensional distribution $P = U \times P^*, Q = U \times Q^*$. Let $(p_1, p_2) \sim P$ and $(q_1, q_2) \sim Q$. Let $X \sim \text{Binomial}(t, p_1), Y \sim \text{Binomial}(t, p_2)$ and $X' \sim \text{Binomial}(t, q_1), Y' \sim \text{Binomial}(t, q_2)$. Notice that $\delta(P) - \delta(Q) \geq \Omega\left(\frac{1}{\sqrt{t \log N}}\right)$. Since the distribution of X and X' are identical, we have

$$\text{TV}((X, Y), (X', Y')) = \text{TV}(Y, Y') \leq \frac{2\sqrt{t}}{Ne^4}.$$

Let $\mathbf{S} := \{(X_i, Y_i)\}_{i=1}^N$ be random variables with $X_i \sim \text{Binomial}(t, p_i)$, $Y_i \sim \text{Binomial}(t, q_i)$ where (p_i, q_i) is drawn independently from P , and denote the joint distribution of \mathbf{S} as P_N . Accordingly, let $\mathbf{S}' := \{(X_i, Y_i)\}_{i=1}^N$ be random variables with $X_i \sim \text{Binomial}(t, p_i)$, $Y_i \sim \text{Binomial}(t, q_i)$ where (p_i, q_i) is drawn independently from Q , and denote the joint distribution of \mathbf{S}' as Q_N . By the property of the product distribution, for $t \leq \frac{N^2(e^4-1)}{36}$, $\text{TV}(P_N, Q_N) \leq \frac{2\sqrt{t}}{Ne^4-1} \leq \frac{1}{3}$, which implies that the minimax error is at least $\Omega\left(\frac{1}{\sqrt{t \log N}}\right)$.

The lower bound for the case when $t = \mathcal{O}(\log N)$ can be constructed very similarly. We leverage Proposition 3.1 and construct distributions similar to above supported on $[0, 1]$ whose first t moments match while the ℓ_1 -magnitudes of change differ by $\frac{1}{t}$. \square

Minimax lower bounds for estimating the distribution of change and the joint distribution are implied by the lower bounds on the estimation of ℓ_1 -magnitude of change.

B Proofs of Upper Bounds

Using the definition of Δ^* (Equation (1)) and changing variables from (x, z) to (x, y) (Figure 1), for any function f , the expectation of f under Δ^* can be written as,

$$\mathbb{E}_{\Delta^*}[f] = \int_{z=-1}^1 f(z) d\Delta^*(z) = \int_{z=-1}^1 \int_{x=\max\{0,-z\}}^{\min\{1,1-z\}} f(z) dR^*(x, x+z) \quad (15)$$

$$= \int_0^1 \int_0^1 f(y-x) dR^*(x, y). \quad (16)$$

Similarly, from the definition of the estimator of distribution of change (6), $\mathbb{E}_{\widehat{\Delta}}[f] = \int_{x=0}^1 \int_{y=0}^1 f(y-x) d\widehat{R}_{\text{mle}}(x, y)$. Therefore the Wasserstein-1 distance between the true distribution change Δ^* and the estimated change distribution $\widehat{\Delta}$ can be written as follows,

$$W_1(\Delta^*, \widehat{\Delta}) = \sup_{f \in \text{Lip}(1)} (\mathbb{E}_{\Delta^*}[f] - \mathbb{E}_{\widehat{\Delta}}[f]) \quad (17)$$

$$= \sup_{f \in \text{Lip}(1)} \int_{x=0}^1 \int_{y=0}^1 f(y-x) (dR^*(x, y) - d\widehat{R}_{\text{mle}}(x, y)). \quad (18)$$

The estimator error for magnitude of change can be upper bounded as follows:

$$\left| \|\widehat{\Delta^*}\|_1 - \|\Delta\|_1 \right| = \left| \int \int |y-x| (dR^*(x, y) - d\widehat{R}_{\text{mle}}(x, y)) \right|. \quad (19)$$

By the definition the Wasserstein-1 distance between the true joint distribution R^* and \widehat{R}_{mle} is,

$$W_1(R^*, \widehat{R}_{\text{mle}}) = \sup_{f \in \text{Lip}(1)} \int_0^1 \int_0^1 f(x, y) (dR^*(x, y) - d\widehat{R}_{\text{mle}}(x, y)). \quad (20)$$

Noting that $|y-x|$ is a Lipschitz-1 function, all the estimator error expressions in Equations (8), (9) and (10) involve Lipschitz-1 functions. The following bound holds for any $f \in \text{Lip}(1)$ supported on $[0, 1]^2$ and any polynomial approximation \widehat{f} of f on $[0, 1]^2$.

$$\begin{aligned} & \left| \int_{x=0}^1 \int_{y=0}^1 f(x, y) (dR^*(x, y) - d\widehat{R}_{\text{mle}}(x, y)) \right|, \\ &= \left| \int_0^1 \int_0^1 (f(x, y) - \widehat{f}(x, y)) (dR^* - d\widehat{R}_{\text{mle}}) + \int_0^1 \int_0^1 \widehat{f}(x, y) (dR^* - d\widehat{R}_{\text{mle}}) \right|, \\ &\leq 2\|f - \widehat{f}\|_\infty + \left| \int_0^1 \int_0^1 \widehat{f}(x, y) (dR^* - d\widehat{R}_{\text{mle}}) \right| = 2\|f - \widehat{f}\|_\infty + \left| \mathbb{E}_{R^*}[\widehat{f}] - \mathbb{E}_{\widehat{R}_{\text{mle}}}[\widehat{f}] \right|, \quad (21) \end{aligned}$$

where $\|f - \widehat{f}\|_\infty := \max_{(x, y) \in [0, 1]^2} |f(x, y) - \widehat{f}(x, y)|$ is the uniform polynomial approximation error.

Let \widehat{f} be the 2-d Bernstein polynomial of degree $2t$ defined by,

$$\widehat{f}_t(x, y) = \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} B_u^t(x) B_v^t(y), \quad (22)$$

where $B_u^t(x) := \binom{t}{u} x^u (1-x)^{t-u}$ is the u -th Bernstein polynomial of degree t . Using this in Equation (21), we can bound the EMD between Δ^* and $\widehat{\Delta}$ can be bounded as follows,

$$W_1(\Delta^*, \widehat{\Delta}) \leq \sup_{f \in \text{Lip}(1)} \left\{ 2\|f - \widehat{f}\|_\infty + \left| \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} (\mathbb{E}_{R^*}[h_{uv}] - \mathbb{E}_{\widehat{R}_{\text{mle}}}[h_{uv}]) \right| \right\}, \quad (23)$$

where h_{uv} is the fraction of coins with observation tuple (u, v) . Similarly, the error bound on estimator of l_1 -magnitude of change can be bounded as,

$$\left| \|\widehat{\Delta^*}\|_1 - \|\Delta\|_1 \right| \leq 2|y - x| - \widehat{f}|\infty + \left| \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} \left(E_{R^*}[h_{uv}] - E_{\widehat{R}_{\text{mle}}}[h_{uv}] \right) \right|. \quad (24)$$

The EMD between the true join distribution on (X, Y) , R^* and the MLE estimator \widehat{R}_{mle} can be bounded as,

$$W_1(R^*, \widehat{R}_{\text{mle}}) \leq \sup_{f \in \text{Lip}(1)} \left\{ 2\|f - \widehat{f}\|_\infty + \left| \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} \left(E_{R^*}[h_{uv}] - E_{\widehat{R}_{\text{mle}}}[h_{uv}] \right) \right| \right\}, \quad (25)$$

The second term in the bounds above (Equation (23), (24) and (25)) can be bounded as follows,

$$\begin{aligned} & \left| \sum_{u,v} \kappa_{uv} \left(E_{R^*}[h_{uv}] - E_{\widehat{R}_{\text{mle}}}[h_{uv}] \right) \right| = \left| \sum_{u,v} \kappa_{uv} \left(E_{R^*}[h_{uv}] - h_{uv}^{\text{obs}} + h_{uv}^{\text{obs}} - E_{\widehat{R}_{\text{mle}}}[h_{uv}] \right) \right|, \\ & \leq \underbrace{\left| \sum_{u,v} \kappa_{uv} \left(E_{R^*}[h_{uv}] - h_{uv}^{\text{obs}} \right) \right|}_{(a)} + \underbrace{\left| \sum_{u,v} \kappa_{uv} \left(h_{uv}^{\text{obs}} - E_{\widehat{R}_{\text{mle}}}[h_{uv}] \right) \right|}_{(b)}, \end{aligned} \quad (26)$$

where the term (a) is the sampling error and term (b) captures the error between the observed fingerprint the MLE solution. These two terms can be bounded by the following two lemmas which are extensions of their 1-d counterparts from [VKVK19],

Lemma B.1. (Concentration of 2-d fingerprint) With probability at least $1 - \alpha$,

$$\left| \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} \left(E_{R^*}[h_{u,v}] - h_{u,v}^{\text{obs}} \right) \right| \leq \mathcal{O} \left(\max_{u,v} |\kappa_{uv}| \sqrt{\frac{\log 1/\alpha}{N}} \right). \quad (27)$$

Proof. Recall that $h_{u,v}^{\text{obs}}$ is the fraction of the population that sees u heads out of t tosses before and v heads out of t tosses after. $E_{R^*}[h_{u,v}]$ is the expected fingerprint under the true distribution which is exactly $E[h_{u,v}^{\text{obs}}]$, and we will use $E[h_{u,v}^{\text{obs}}]$ and $E_{R^*}[h_{u,v}]$ interchangeably. Define, $\phi(X, Y) := \sum_{u=0}^t \sum_{v=0}^t \kappa_{u,v} (h_{u,v}^{\text{obs}} - E[h_{u,v}^{\text{obs}}])$, that is,

$$\phi(X, Y) = \frac{1}{N} \sum_{i=1}^N \sum_{u=0}^t \sum_{v=0}^t \kappa_{u,v} \left(\mathbb{1}_{\{(X_i, Y_i)=(u,v)\}} - E[h_{u,v}^{\text{obs}}] \right).$$

Note that $E[\phi(X, Y)] = 0$. Let $\phi_{i'}(X, Y)$ be ϕ with (X_i, Y_i) being replaced by re-drawing (X'_i, Y'_i) , with $X'_i \sim \text{Bin}(t, p_i)$ and $Y'_i \sim \text{Bin}(t, q_i)$. We can bound the difference in $\phi(X, Y)$ and $\phi_{i'}(X, Y)$ as follows,

$$|\phi(X, Y) - \phi_{i'}(X, Y)| \leq \max_{(u,v)} |\kappa_{uv}| \frac{2}{N}.$$

By McDiarmid's Inequality, for some absolute constants $C, c > 0$,

$$\Pr(|\phi(X, Y)| \geq \epsilon) \leq C \exp \left(-\frac{cN\epsilon^2}{4(\max_{(u,v)} |\kappa_{uv}|)^2} \right).$$

Hence, with probability at least $1 - \delta$,

$$\left| \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} \left(E_{R^*}[h_{u,v}] - h_{u,v}^{\text{obs}} \right) \right| \leq \mathcal{O} \left(\max_{u,v} |\kappa_{uv}| \sqrt{\frac{\log 1/\delta}{N}} \right). \quad (28)$$

□

Lemma B.2. (Bound on error in 2-d fingerprint due to MLE solution) For $3 \leq (t+1)^2 \leq \sqrt{C_0 N} + 2$, where $C_0 > 0$ is a constant, w. p. $1 - \alpha$,

$$\begin{aligned} \left| \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} (h_{u,v}^{\text{obs}} - E_{R_{\text{mle}}}[h_{u,v}]) \right| &\leq \max_{(u,v)} |\kappa_{uv}| \sum_{u=0}^t \sum_{v=0}^t |(h_{u,v}^{\text{obs}} - E_{R_{\text{mle}}}[h_{u,v}])|, \\ &\leq \max_{(u,v)} |\kappa_{uv}| \sqrt{2 \ln(2)} \sqrt{\frac{(t+1)^2}{2N} \log \frac{4N}{(t+1)^2} + \frac{1}{N} \log \frac{3e}{\alpha}}. \end{aligned} \quad (29)$$

Proof. Let \hat{R}_{mle} be an optimal solution to the MLE and R^* be the true distribution. By optimality of the MLE solution, we have the following inequality,

$$\text{KL}(\mathbf{h}^{\text{obs}}, E_{\hat{R}_{\text{mle}}}[\mathbf{h}]) \leq \text{KL}(\mathbf{h}^{\text{obs}}, E_{R^*}[\mathbf{h}]). \quad (30)$$

Proposition B.1 (Pinsker's Inequality [CT12]). For discrete distributions P and Q :

$$\text{KL}(P, Q) \geq \frac{1}{2 \ln 2} \|P - Q\|_1^2. \quad (31)$$

Note that the fingerprint vector can be seen as a discrete distribution with support size $(t+1)^2$. Using Pinsker's inequality and the optimality of the MLE solution,

$$\begin{aligned} \left| \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} (h_{u,v}^{\text{obs}} - E_{\hat{R}_{\text{mle}}}[h_{u,v}]) \right| &\leq \sum_{u=0}^t \sum_{v=0}^t |\kappa_{uv}| |h_{u,v}^{\text{obs}} - E_{\hat{R}_{\text{mle}}}[h_{u,v}]|, \\ &\leq \max_{(u,v)} |\kappa_{uv}| \sum_{j=0}^t |h_j^{\text{obs}} - E_{\hat{R}_{\text{mle}}}[h_j]|, \\ &\leq \max_{(u,v)} |\kappa_{uv}| \sqrt{2 \ln(2) \text{KL}(\mathbf{h}^{\text{obs}}, E_{\hat{R}_{\text{mle}}}[\mathbf{h}])}, \end{aligned} \quad (32)$$

$$\leq \max_{(u,v)} |\kappa_{uv}| \sqrt{2 \ln(2) \text{KL}(\mathbf{h}^{\text{obs}}, E_{R^*}[\mathbf{h}])}. \quad (33)$$

Using the results on bounds on KL divergence between empirical observations and the true distribution for discrete distributions [MJT⁺18], for $3 \leq (t+1)^2 \leq \sqrt{C_0 N} + 2$, w. p. $1 - \delta$,

$$\left| \sum_{u=0}^t \sum_{v=0}^t \kappa_{uv} (h_{u,v}^{\text{obs}} - E_{R_{\text{mle}}}[h_{u,v}]) \right| \leq \max_{(u,v)} |\kappa_{uv}| \sqrt{2 \ln(2)} \sqrt{\frac{(t+1)^2}{2N} \log \frac{4N}{(t+1)^2} + \frac{1}{N} \log \frac{3e}{\delta}}.$$

□

Note that the term $\max_{(u,v)} |\kappa_{uv}|$ appears in the bounds on both the terms (a) and (b). Therefore, it is crucial to bound the uniform approximation error $\|f - \hat{f}\|_\infty$ while keeping the coefficients of the approximation (12) from being too large. Polynomial approximation of Lipschitz-1 functions using Bernstein polynomials play a crucial role in obtaining the minimax optimal rates in Theorem 2.2 and Proposition 2.1.

Lemma B.3. (Bernstein polynomial approximation of 2-d Lipschitz-1 functions) Any Lipschitz-1 function on $[0, 1]^2$ can be approximated using Bernstein polynomials (Equation (12)) of degree $2t$, with an uniform approximation error of

- $\mathcal{O}(\frac{1}{t})$ with $\max_{(u,v)} |\kappa_{uv}| \leq t2^{2t}$.
- $\mathcal{O}(\frac{1}{k})$ with $\max_{(u,v)} |\kappa_{uv}| \leq k(t+1)^2 e^{\frac{2k^2}{t}}$, for $k < t$.

The proof of Lemma B.3 uses the following results.

Lemma B.4. ([BBL02] Polynomial approximation of Continuous Functions on $[0, 1]^2$) Given any continuous function $f : [0, 1]^2 \rightarrow \mathbb{R}$, there exists a degree k polynomial,

$$p_k(x, y) := \sum_{\substack{i=0 \\ :i+j \leq k}}^k \sum_{j=0}^k c_{ij} x^i y^j, \quad \forall (x, y) \in [0, 1]^2 \quad (34)$$

that approximates f with error $\max_{(x,y) \in [0,1]^2} |f(x, y) - p_k(x, y)| = O(\frac{1}{k})$.

Let \tilde{T}_m denote Chebyshev polynomial of degree m shifted to $[0, 1]$ which satisfy the following recursive relation:

$$\tilde{T}_m(x) = (4x - 2)\tilde{T}_{m-1} - \tilde{T}_{m-2}(x), \quad m = 2, 3, \dots,$$

and $\tilde{T}_0(x) = 1$, $\tilde{T}_1(x) = 2x - 1$.

Lemma B.5. (Chebyshev polynomial approximation of Lipschitz-1 functions on $[0, 1]^2$) Given any Lipschitz-1 function $f : [0, 1]^2 \rightarrow \mathbb{R}$, there exists a degree k polynomial in the form of

$$f_k(x, y) = \sum_{\substack{m=0 \\ :m+l \leq k}}^k \sum_{l=0}^k \tau_{m,l} \tilde{T}_m(x) \tilde{T}_l(y), \quad (35)$$

that approximates $f(x, y)$ with uniform approximation error of $\|f - f_k\|_\infty = O(\frac{1}{k})$, where $\tilde{T}_m(x)$ denotes Chebyshev polynomial of degree m shifted to $[0, 1]$. Further, the coefficient vector satisfies $\|\tau\|_2 \leq 32$.

Proof. Chebyshev polynomials of degree up to k form a basis for polynomials of degree k . Therefore each term $x^i y^j$ in the approximation polynomial (34) can be written as a linear combination of Chebyshev polynomials of degree up to k ,

$$x^i y^j = \left(\sum_{m=0}^i a_{m,i} \tilde{T}_m(x) \right) \left(\sum_{n=0}^j b_{n,j} \tilde{T}_n(y) \right) = \sum_{m,n=0}^{i,j} a_{m,i} b_{n,j} \tilde{T}_m(x) \tilde{T}_n(y). \quad (36)$$

The approximation polynomial in Lemma B.4 (Equation (34)) can be transformed to Chebyshev polynomial basis,

$$p_k(x, y) = \sum_{\substack{m=0 \\ :m+l \leq k}}^k \sum_{l=0}^k \tau_{m,l} \tilde{T}_m(x) \tilde{T}_l(y) =: f_k(x, y), \quad (37)$$

where, $\tau_{m,l} = \sum_{\substack{i=m \\ :i+j \leq k}}^k \sum_{j=l}^k c_{i,j} a_{m,i} b_{l,j}$. Therefore, $f_k(x, y)$ is a uniform $\frac{1}{k}$ - approximation of $f(x, y)$. Chebyshev polynomials (shifted) form a sequence of orthogonal polynomials with respect to the weight $\frac{1}{\sqrt{4x-4x^2}}$:

$$\int_0^1 \tilde{T}_m(x) \tilde{T}_n(x) \frac{dx}{\sqrt{4x-4x^2}} = \begin{cases} 0 & \text{if } m \neq n \\ \frac{\pi}{2} & \text{if } m = n = 0 \\ \frac{\pi}{4} & \text{if } m = n \neq 0. \end{cases} \quad (38)$$

Therefore, we can bound the coefficients of Chebyshev polynomials in Equation (37),

$$\int \int |f_k(x, y)|^2 \frac{dx}{\sqrt{4x-4x^2}} \frac{dy}{\sqrt{4y-4y^2}} = \frac{\pi^2}{16} \left(4\tau_{0,0}^2 + 2 \sum_{l=0}^k \tau_{0,l}^2 + 2 \sum_{m=0}^k \tau_{m,0}^2 + \sum_{\substack{m=1 \\ :m+l \leq k}}^k \sum_{l=1}^k \tau_{m,l}^2 \right).$$

Since f is Lipschitz-1 on $[0, 1]^2$, w.l.o.g $f(x, y) \leq \sqrt{2}$ for all (x, y) in $[0, 1]^2$. Since $f_k(x, y)$ is a $\frac{1}{k}$ -uniform approximation of f on $[0, 1]^2$, we have $|f_k(x, y)| \leq \sqrt{2} + \frac{1}{k} \leq 2\sqrt{2}$ for all (x, y) in $[0, 1]^2$. Therefore,

$$\int \int |f_k(x, y)|^2 \frac{dx}{\sqrt{4x - 4x^2}} \frac{dy}{\sqrt{4y - 4y^2}} \leq 2\pi^2. \quad (39)$$

Thus we can bound the l_2 -norm of the coefficients τ in the approximation (37),

$$3 \tau_{0,0}^2 + \sum_{l=0}^k \tau_{0,l}^2 + \sum_{m=0}^k \tau_{m,0}^2 + \|\tau\|_2^2 \leq 32.$$

Therefore, $\|\tau\|_2^2 \leq 32$. □

Note that τ is a $\binom{k+2}{2}$ length vector. Hence, we have the following bound,

$$\|\tau\|_1 = \sum_{m=0}^k \sum_{\substack{l=0 \\ :m+l \leq k}}^k |\tau_{m,l}| \leq \sqrt{32} \sqrt{\binom{k+2}{2}} < 4(k+2). \quad (40)$$

Transforming to Bernstein polynomials: Bernstein polynomials of degree $m < t$ can be raised to degree t as:

$$B_i^m(x) = \sum_{j=i}^{i+t-m} \frac{\binom{m}{i} \binom{t-m}{j-i}}{\binom{t}{j}} B_j^t(x). \quad (41)$$

Using degree raising of Bernstein polynomials, we can write shifted Chebyshev polynomials of degree $m < t$ in terms of Bernstein polynomials of degree t as [Rab03],

$$\tilde{T}_m(x) = \sum_{i=0}^m (-1)^{m-i} \frac{\binom{2m}{2i}}{\binom{m}{i}} \sum_{j=i}^{i+t-m} \frac{\binom{m}{i} \binom{t-m}{j-i}}{\binom{t}{j}} B_j^t(x) =: \sum_{j=0}^t C(t, m, j) B_j^t(x), \quad (42)$$

where the coefficient of j -th Bernstein polynomial of degree t is given by,

$$C(t, m, j) := \sum_{l=0}^j (-1)^{m-l} \frac{\binom{2m}{2l} \binom{t-m}{j-l}}{\binom{t}{j}}, \quad (43)$$

with $\binom{a}{b} = 0$ when $b > a > 0$. The following lemma from [VKVK19] provides an upper bound on the coefficients $C(t, m, j)$,

Lemma B.6. The l_2 -norm of the coefficients of B_j^t can be bounded as follows,

$$\sqrt{\sum_{j=0}^t |C(t, m, j)|^2} \leq (t+1) e^{\frac{m^2}{t}}. \quad (44)$$

And, hence the coefficients of B_j^t can be bounded as follows,

$$|C(t, m, j)| \leq (t+1) e^{\frac{m^2}{t}}. \quad (45)$$

Every term in the Chebyshev approximation polynomial of degree k in Equation (35) can therefore be written using Bernstein polynomials of degree up to t ,

$$\begin{aligned} f_k(x, y) &= \sum_{m=0}^k \sum_{\substack{l=0 \\ :m+l \leq k}}^k \tau_{m,l} \tilde{T}_m(x) \tilde{T}_l(y) = \sum_{m=0}^k \sum_{\substack{l=0 \\ :m+l \leq k}}^k \tau_{m,l} \left(\sum_{j=0}^t C(t, m, j) B_j^t(x) \right) \left(\sum_{v=0}^t C(t, l, v) B_v^t(y) \right), \\ &=: \sum_{j=0}^t \sum_{v=0}^t \kappa_{jv} B_j^t(x) B_v^t(y), \end{aligned} \quad (46)$$

where the coefficients are given by,

$$\kappa_{jv} := \sum_{\substack{m=0 \\ :m+l \leq k}}^k \sum_{\substack{l=0 \\ :m+l \leq k}}^k \tau_{m,l} C(t, m, j) C(t, l, v). \quad (47)$$

Each of these coefficients can be bounded by,

$$\begin{aligned} |\kappa_{jv}| &\leq \sum_{\substack{m=0 \\ :m+l \leq k}}^k \sum_{\substack{l=0 \\ :m+l \leq k}}^k |\tau_{m,l}| |C(t, m, j)| |C(t, l, v)| \leq \max_{m,l} |C(t, m, j)| |C(t, l, v)| \left(\sum_{\substack{m=0 \\ :m+l \leq k}}^k \sum_{\substack{l=0 \\ :m+l \leq k}}^k |\tau_{m,l}| \right) \\ &\leq (t+1)^2 e^{\frac{2k^2}{t}} \left(\sum_{\substack{m=0 \\ :m+l \leq k}}^k \sum_{\substack{l=0 \\ :m+l \leq k}}^k |\tau_{m,l}| \right) \leq 4(k+2)(t+1)^2 e^{\frac{2k^2}{t}}. \end{aligned} \quad (48)$$

This concludes the proof of Lemma B.3.

Putting these results together,

$$\begin{aligned} W_1 \left(R^*, \widehat{R}_{\text{mle}} \right) &\leq \mathcal{O} \left(\frac{1}{k} \right) + \mathcal{O} \left(\max_{u,v} |\kappa_{uv}| \sqrt{\frac{\log 1/\delta}{N}} \right), \\ &\quad + \mathcal{O} \left(\max_{(u,v)} |\kappa_{uv}| \sqrt{\frac{(t+1)^2 \log \frac{4N}{(t+1)^2} + \frac{1}{N} \log \frac{3e}{\delta}}{2N}} \right), \end{aligned} \quad (49)$$

$$\begin{aligned} &\leq \mathcal{O} \left(\frac{1}{k} \right) + \mathcal{O} \left(\max_{u,v} |\kappa_{uv}| \sqrt{\frac{\log 1/\delta}{N}} \right), \\ &\quad + \mathcal{O} \left(\max_{(u,v)} |\kappa_{uv}| \sqrt{\frac{(t+1)^2 \log \frac{4N}{(t+1)^2} + \frac{1}{N} \log \frac{3e}{\delta}}{2N}} \right), \end{aligned} \quad (50)$$

where $k = t$ for the case when $t = \mathcal{O}(\log N)$, and $k = \sqrt{t \log N^c}$ when $t > \log N$. These bounds are also upper bounds for $W_1 \left(\Delta^*, \widehat{\Delta} \right)$ and $\| \|\Delta^*\|_1 - \|\widehat{\Delta}\|_1 \|$. This completes the proof of Proposition 2.1 and the main results stated in Theorem 2.2.