

# Testing (Conditional) Mutual Information\*

Jan Seyfried  
 Centre for Quantum Technologies,  
 National University of Singapore  
 jan.seyfried@u.nus.edu

Sayantana Sen  
 Centre for Quantum Technologies,  
 National University of Singapore  
 sayantan789@gmail.com

Marco Tomamichel  
 Department of Electrical and Computer Engineering,  
 Centre for Quantum Technologies,  
 National University of Singapore  
 marco.tomamichel@nus.edu.sg

## Abstract

We investigate the sample complexity of mutual information and conditional mutual information testing. For conditional mutual information testing, given access to independent samples of a triple of random variables  $(A, B, C)$  with unknown distribution, we want to distinguish between two cases: (i)  $A$  and  $C$  are conditionally independent, i.e.,  $I(A : C|B) = 0$ , and (ii)  $A$  and  $C$  are conditionally dependent, i.e.,  $I(A : C|B) \geq \varepsilon$  for some threshold  $\varepsilon$ . We establish an upper bound on the number of samples required to distinguish between the two cases with high confidence, as a function of  $\varepsilon$  and the three alphabet sizes. We conjecture that our bound is tight and show that this is indeed the case in several parameter regimes. For the special case of mutual information testing (when  $B$  is trivial), we establish the necessary and sufficient number of samples required up to polylogarithmic terms.

Our technical contributions include a novel method to efficiently simulate weakly correlated samples from the conditionally independent distribution  $P_{A|B}P_{C|B}P_B$  given access to samples from an unknown distribution  $P_{ABC}$ , and a new estimator for equivalence testing that can handle such correlated samples, which might be of independent interest.

---

\*Accepted for presentation at the Conference on Learning Theory (COLT) 2025

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Overview of Results</b>	<b>5</b>
2.1	Problem Setup . . . . .	5
2.2	Our Results in Context . . . . .	8
2.3	Methods . . . . .	9
<b>3</b>	<b>Preliminaries</b>	<b>14</b>
<b>4</b>	<b>Reducing CMI Testing to Testing in <math>D_H^2</math></b>	<b>17</b>
<b>5</b>	<b>Independence Testing in <math>D_H^2</math></b>	<b>20</b>
5.1	Some Preliminary Results . . . . .	21
5.2	Our Independence Tester . . . . .	23
5.2.1	Sample Complexity of the Heavy Regime . . . . .	26
5.2.2	Sample Complexity of the Mixed Regime . . . . .	27
<b>6</b>	<b>Lower Bounds for Independence Testing</b>	<b>29</b>
6.1	Description of the Hard Distributions . . . . .	30
6.2	Preliminary Properties of the Mutual Information . . . . .	30
6.3	Bounding the Mutual Information . . . . .	32
<b>7</b>	<b>Conditional Independence Testing</b>	<b>37</b>
7.1	Testing The Small Regime . . . . .	39
7.1.1	Description of the Algorithm . . . . .	39
7.1.2	Intuition for the Sample Complexity . . . . .	41
7.1.3	Formal Analysis . . . . .	42
7.2	Testing the Large Regime . . . . .	46
7.2.1	Subroutines . . . . .	47
7.2.2	Description of the Algorithm . . . . .	49
7.2.3	Sample Complexity of Heavy Categories . . . . .	54
7.2.4	Sample Complexity of Mixed Categories . . . . .	55
7.2.5	Sample Complexity of Light Categories . . . . .	57
<b>8</b>	<b>Lower Bounds for Conditional Independence Testing</b>	<b>60</b>
<b>9</b>	<b>Equivalence Testing of Distributions</b>	<b>62</b>
<b>A</b>	<b>Calculations</b>	<b>70</b>
A.1	Remaining Proof from Section 3 . . . . .	70
A.2	Remaining Proof from Section 4 . . . . .	70
A.3	Remaining Proofs from Section 6 . . . . .	72
A.4	Remaining Proofs from Section 7 . . . . .	74

# 1 Introduction

Distribution testing is a central problem in computer science and statistics [Rub12, Gol17, Can20, Can22]. Given access to independent samples from an unknown distribution, the goal of distribution testing is to efficiently determine whether the distribution has some specific property. In particular, we are interested in the sample complexity of this problem; the minimal number of samples needed to correctly identify the property with high confidence. Distribution testing has been a vibrant topic of research recently due to the fact that testing algorithms can succeed with far fewer samples than would be needed for completely learning the unknown distribution. Indeed, for many testing problems, the sample complexity grows sub-linearly with the size of the distribution, making such algorithms amenable to problems involving big data sets [Rub12].

Independence testing and conditional independence testing are two distribution testing problems that are of particular relevance and have found various applications ranging from the study of causality [Gra80, Pea14, SGS01, ZPJS11], information theory [TH18] and economics [WH18, dMASdBP14] to machine learning and graphical models [N+04, CDKS20, DP17, BGP+23, DP21, CYBC24, WGT+24]. In particular, in the context of graphical models such as Bayesian Networks, learning tree-structured distributions [CL68] (an important subclass of Bayesian networks) reduces to the problem of testing the conditional independence of the underlying distribution [DP21]. Recently, [CYBC24] studied the more general class of *polytree* structured distributions (a Bayesian network is a polytree if the underlying undirected graph is a tree), which again uses a reduction to the conditional independence testing of the associated distribution.

In independence testing, we are given access to i.i.d. samples of a pair of random variables  $(A, C)$  and our goal is to test if the two random variables are independent of each other, or if they are correlated beyond some threshold. Similarly, in conditional independence testing, we want to test if a triple of random variables  $(A, B, C)$  is conditionally independent and forms a Markov chain  $A - B - C$ , or if they are conditionally dependent above some threshold. We can always consider independence testing as a special case of conditional independence testing where  $B$  is trivial, and hence in the following, we will focus our attention on the latter problem.

The first question to ask then is how we measure conditional dependence. An uncontroversial choice is to measure this in terms of the total variation distance to Markov chains, i.e., to devise an algorithm which can distinguish between the case where  $A - B - C$  form a Markov chain and the case where the distribution of  $(A, B, C)$  is  $\varepsilon$ -far from any Markov chain in total variation distance. This problem has been studied recently [CDKS18] and more extensively in the special case of independence testing [BFF+01, ADK15, DK16, DGK+21]. We will discuss these results and their relation to our work in more detail in the next section. However, we argue that measuring distance in Kulback-Leibler (KL) divergence is more natural in this context since the distance to the closest Markov chain in KL divergence, the conditional mutual information [Wyn78], has well-defined operational meaning in information theory and is a widely used measure of conditional dependence.

This brings us to conditional mutual information (CMI) testing, which was (to the best of our knowledge) first studied in [CDKS18, BGP+23]. In CMI testing, we are given a promise that the variables  $(A, B, C)$  are either conditionally independent,  $I(A : C|B) = 0$ , or that they are indeed sufficiently conditionally dependent,  $I(A : C|B) \geq \varepsilon$ . The promise gap, or threshold,  $\varepsilon > 0$ , plays an important role in the analysis of this problem, and we measure it in terms of conditional mutual information due to its wide use in information theory, statistics and computer science. The goal of CMI testing is to devise an efficient tester that can correctly distinguish these two cases with high probability (for distributions in between, with  $0 < I(A : C|B) < \varepsilon$ , we accept either output). Mutual information (MI) testing and CMI testing will be formally introduced in the next section.

In this work we investigate the sample complexity of both MI and CMI testing, i.e., we find the minimal number of samples required from an unknown distribution to succeed in these tasks with high probability. It is worth noting at this point that any algorithm that can learn the underlying distribution  $P_{ABC}$  of  $(A, B, C)$  with sufficient precision will also succeed at this task—but we will see that this is not efficient, as it would necessitate at least  $\Omega(d^3)$  samples, where  $d$  is the alphabet size, assumed to be the same for all three variables. Hence, the growth is linear in the alphabet size of the triple  $(A, B, C)$ . On the other hand, a sub-linear scaling of  $O(d^{7/4})$  samples is sufficient to solve this problem [CDKS18].

Our results are much more fine-grained in terms of their dependence on the three alphabet sizes,  $d_A$ ,  $d_B$  and  $d_C$ , as well as the threshold  $\varepsilon$ , and reveal a complex landscape of scaling regimes depending on the ratios between the different problem parameters. We solve this problem completely for MI testing, where we give matching upper and lower bounds on the sample complexity in Result 1. We find that it scales as

$$\tilde{\Theta} \left( \min \left\{ \frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \right). \quad (1)$$

For CMI testing we show an upper bound on the sample complexity in Result 2. Namely, the sample complexity scales as

$$\tilde{O} \left( \max \left\{ \frac{d_A^{1/2} d_B^{3/4} d_C^{1/2}}{\varepsilon}, \min \left\{ \frac{d_A^{1/4} d_B^{7/8} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/7} d_B^{6/7} d_C^{2/7}}{\varepsilon^{8/7}} \right\}, \min \left\{ \frac{d_A^{3/4} d_B^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_B^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \right\} \right). \quad (2)$$

We conjecture this to be tight in all regimes but our Result 3 only proves this partially. We provide an in-depth discussion of our results in the next section. They exhibit some similarity with the bounds achieved for total variation distance [CDKS18], but differ in some interesting ways that we will discuss. Our treatment is also more complete since we can show matching upper and lower bounds for MI testing and some regimes of CMI testing for general parameters.

To achieve these tight bounds we need to introduce several new techniques that we believe will be of independent interest. For MI testing, we reduce the problem to equivalence testing in Hellinger distance between  $P_{AC}$  and  $P_A P_C$ , and then leverage the fact that our reference distribution is product to improve upon existing techniques [DK16, CDVV14] to optimally solve this problem by relating it to  $\ell_2$  distance equivalence testing. For CMI, we dig deeper as we first need to find a way to efficiently sample from the Markov distribution  $P_{A|B} P_{C|B} P_B$  given access only to the joint distribution  $P_{ABC}$ . This is difficult because sampling from  $P_{A|B} P_{C|B} P_B$  requires collisions in the  $B$  coordinate, but we fail to see sufficiently many collisions for all low probability outcomes in  $B$ . We devise such an efficient sampler, which however gives us samples that are not independent. This then necessitates the introduction of a new equivalence tester that can deal with such weakly correlated samples. We describe our methods in detail in the next section. We note that there have also been some works in the statistics literature, where it is assumed that the number of samples is asymptotically large, independent of the other parameters [NBW21, MV19, KNBW22, KNBW24], which is a different setting than our work.

The remainder of this paper is structured as follows. In Section 2, we present a detailed overview of our results and the methods we employed to prove them. In Section 3, we discuss the necessary preliminaries for this work. Then, in Section 4, we show the reduction from (conditional) mutual information testing to independence testing in squared Hellinger distance. In Section 5, we present our tester for independence testing of bipartite distributions, followed by the proof of the lower

bound of our independence tester in Section 6. Next, in Section 7, we present our conditional mutual information tester, and in Section 8, we discuss the lower bounds for conditional mutual information testing. In Section 9, we present our result for equivalence testing of distributions. Finally, in Appendix A, we present the proofs of some lemmas and claims that are omitted previously.

## 2 Overview of Results

### 2.1 Problem Setup

Before we can formally state our results we will introduce our model and formal problem statements. To do this, let us first introduce some notation.

**Definition 2.1** (Independence and Conditional Independence). Let  $(A, B, C)$  be discrete random variables defined over discrete alphabets  $\mathcal{A} := \{1, 2, \dots, d_A\}$ ,  $\mathcal{B} := \{1, 2, \dots, d_B\}$  and  $\mathcal{C} := \{1, 2, \dots, d_C\}$ , respectively. Then,  $A$  and  $C$  are said to be *independent* if

$$P_{AC}(a, c) = P_A(a)P_C(c) \quad \forall a \in \mathcal{A}, c \in \mathcal{C}, \quad (3)$$

or, in short,  $P_{AC} = P_A P_C$ , where  $P_A$  and  $P_C$  are marginal distributions of  $P_{AC}$ . We then call  $P_{AC}$  a *product distribution*. Moreover,  $A$  and  $C$  are said to be *conditionally independent given B* if

$$P_{AC|B}(a, c|b) = P_{A|B}(a|b)P_{C|B}(c|b) \quad \forall a \in \mathcal{A}, b \in \mathcal{B}, c \in \mathcal{C}, \quad (4)$$

or, in short,  $P_{AC|B} = P_{A|B}P_{C|B}$ , where  $P_{AC|B}$ ,  $P_{A|B}$  and  $P_{C|B}$  are conditional distributions. In this case we also say that  $P_{ABC}$  forms a *Markov chain*  $A - B - C$ .

In the absence of independence, we have correlation between the random variables. Correlation between discrete random variables is naturally measured in terms of mutual information. The mutual information has various operational interpretations in information theory, most prominently in the study of noisy channel coding [Sha48]. The mutual information,  $I(A:C)$ , can be seen as the minimal Kullback-Leibler divergence between the joint distribution of  $A$  and  $C$  and any product distribution on the same alphabets. It is positive and zero only for product distributions. Similarly, conditional mutual information,  $I(A:C|B)$ , is the minimal Kullback-Leibler divergence between the joint distribution of  $A$ ,  $B$  and  $C$  and any Markov chain  $A - B - C$ . It is positive and zero only for Markov chains. We will use these observations for our formal definitions.

**Definition 2.2** (Mutual information and Conditional Mutual information). Let  $P_{AC}$  be the joint distribution of  $A$  and  $C$ . The *mutual information (MI)* of  $A$  and  $C$  is defined as

$$I(A:C)_P := \min_{Q_A, Q_C} D(P_{AC} \| Q_A Q_C), \quad (5)$$

where the minimization is over all product distributions  $Q_{AC} = Q_A Q_C$ , and

$$D(P \| Q) := \sum_x P(x) \log \frac{P(x)}{Q(x)} \quad (6)$$

is the *Kullback-Leibler (KL) divergence* between  $P$  and  $Q$ .<sup>1</sup> Moreover, the *conditional mutual information (CMI)* of  $A$  and  $C$  given  $B$  is defined as

$$I(A:C|B)_P := \min_{Q_B, Q_{A|B}, Q_{C|B}} D(P_{ABC} \| Q_{A|B} Q_B Q_{C|B}), \quad (7)$$

where the minimization is over all Markov chains  $Q_{ABC} = Q_{A|B} Q_B Q_{C|B}$  of the form  $A - B - C$ .

---

<sup>1</sup>We will assume here and throughout that the support of  $Q$  contains the support of  $P$ , and use the convention  $0 \log 0 = 0$  to deal with zeros, so that the KL-divergence is always finite.

We will omit the subscript  $P$  from  $I(A : C)_P$  and  $I(A : C|B)_P$  when it is clear from the context.

We note that the minimum in the above expression is achieved by the marginal distributions of  $P_{AC}$ , i.e., it holds that  $I(A : C)_P = D(P_{AC} \| P_A P_C)$ . Similarly, one can show that  $I(A : C|B)_P = D(P_{ABC} \| P_{A|B} P_B P_{C|B})$ . While these expressions are more commonly used as the definition of (conditional) mutual information, our alternative choice highlights that:

$$\begin{aligned} I(A : C)_P \geq \varepsilon &\iff P_{AC} \text{ is } \varepsilon\text{-far from the set of product distributions,} && \text{and} \\ I(A : C|B)_P \geq \varepsilon &\iff P_{ABC} \text{ is } \varepsilon\text{-far from the set of } A - B - C \text{ Markov chains.} \end{aligned}$$

Here and throughout “ $P$  is  $\varepsilon$ -far from  $Q$ ” refers to the condition  $D(P \| Q) \geq \varepsilon$ . With this in mind we can now state the sample complexity problems this paper is concerned with. We will now define different testing problems, in which we always want to correctly assign a distribution to one of two classes. We refer to an algorithm performing such a test as a *tester*. Throughout, we say that a tester succeeds if it correctly identifies when the distribution  $P$  is in one of the two classes. If neither condition is satisfied, either outcome is accepted. The *sample complexity* of a problem is the minimal number  $N$  of samples for which there exists a tester that succeeds with probability at least  $2/3$ , where  $N$  depends on the parameters of the problem.

**Problem 1** (Independence testing). Fix a distance measure  $\Delta$ , a threshold  $\varepsilon > 0$  and alphabet sizes  $d_A$  and  $d_C$ . Consider the following decision problem: Given access to  $N$  i.i.d. samples from an unknown distribution  $P_{AC}$ , distinguish between the classes

- (i)  $\Delta(P_{AC} \| P_A P_C) = 0$ , and
- (ii)  $\Delta(P_{AC} \| P_A P_C) \geq \varepsilon$ .

We denote the sample complexity of this problem by  $\text{SC}_{\text{Indep}}(\Delta, \varepsilon, d_A, d_C)$ .

We are interested in solving a special case of this problem, where the distance measure  $\Delta$  is chosen to equal the Kullback-Leibler divergence.

**Problem 2** (MI testing). Consider the setting of independence testing (Problem 1), where  $\Delta(\cdot \| \cdot) = D(\cdot \| \cdot)$ , which is equivalent to deciding between the classes

- (i)  $I(A : C)_P = 0$ , and
- (ii)  $I(A : C)_P \geq \varepsilon$ .

We denote the sample complexity by  $\text{SC}_{\text{MI}}(\varepsilon, d_A, d_C) := \text{SC}_{\text{Indep}}(D, \varepsilon, d_A, d_C)$ .

This problem was first studied in this form in [BGP<sup>+</sup>23], albeit only for the case when  $d_A = d_C$ . We would like to note that a variant of this problem has been studied in [BFF<sup>+</sup>01, LRR13, ADK15, DK16, DGK<sup>+</sup>21]. There, the authors designed an independence testing algorithm that tests if  $P_{AC}$  is a product distribution or if  $P_{AC}$  is at least  $\varepsilon$ -far in total variation distance from any product distribution.

Now we consider the problem of conditional mutual information testing, which we first introduce in a more general form:

**Problem 3** (Conditional independence testing). Fix a threshold  $\varepsilon$  and alphabet sizes  $d_A, d_B$  and  $d_C$ . Consider the following decision problem: Given access to  $N$  i.i.d. samples from an unknown distribution  $P_{ABC}$ , distinguish between the cases<sup>2</sup>

<sup>2</sup>Note that for the definition of the decision gap, we fix the ‘reference distribution’ to be  $P_{AB} P_{C|B}$ , which might not necessarily be the closest conditionally independent distribution. In general,  $\Delta(P_{ABC} \| P_{AB} P_{C|B}) \neq \min_{Q_{ABC}} \Delta(P_{ABC} \| Q_{AB} Q_{C|B})$ . However, as mentioned before, this equality does hold if we set  $\Delta(\cdot \| \cdot) = D(\cdot \| \cdot)$ .

- (i)  $\Delta(P_{ABC} \| P_{AB}P_{C|B}) = 0$  and
- (ii)  $\Delta(P_{ABC} \| P_{AB}P_{C|B}) \geq \varepsilon$ .

We denote the sample complexity by  $\text{SC}_{\text{CIndep}}(\Delta, \varepsilon, d_A, d_B, d_C)$ .

Similar to before, when the distance measure  $\Delta$  is the Kullback-Leibler divergence, we have the following problem.

**Problem 4** (CMI testing). Consider the setting of conditional independence testing (Problem 3), where  $\Delta(\cdot \| \cdot) = D(\cdot \| \cdot)$ , which is equivalent to deciding between the classes

- (i)  $I(A:C|B)_P = 0$  and
- (ii)  $I(A:C|B)_P \geq \varepsilon$ .

We denote the sample complexity by  $\text{SC}_{\text{CMI}}(\varepsilon, d_A, d_B, d_C) = \text{SC}_{\text{CIndep}}(D, \varepsilon, d_A, d_B, d_C)$ .

This problem was first studied in [BGP<sup>+</sup>23] for  $d_A = d_B = d_C$  in the context of learning a special case of Bayesian networks, namely learning tree-structured distributions. A related work in this context is the work of [CDKS18]. The authors in [CDKS18] studied a related problem of testing whether  $I(A:C|B) = 0$ , or the underlying distribution is far in variation distance from all such distributions. It is clear that MI testing is a special case of CMI testing when  $d_B = 1$ , i.e., when  $B$  is trivial.

Now we consider the problem of equivalence testing of distributions.

**Problem 5** (Sample complexity of equivalence testing). Fix a threshold  $\varepsilon$ , a distance measure  $\Delta$ , an integer  $n$ , and a positive real number  $b$ . Consider the following decision problem: Given access to  $N$  i.i.d. samples from two unknown distributions  $P$  and  $Q$  each defined over  $[n]$ , with the additional guarantee that  $\max\{\|P\|_2, \|Q\|_2\} \leq b$ , distinguish between the classes

- (i)  $P = Q$  and
- (ii)  $\Delta(P \| Q) \geq \varepsilon$ .

We denote the sample complexity of this problem by  $\text{SC}_{\text{EQIV}}(\Delta, \varepsilon, b, n)$ .

Equivalence testing has been studied for a long time in the literature, starting with the work of [BFR<sup>+</sup>13], who studied the problem when the distance measure is the  $\ell_1$  distance and there is no bound on the 2-norm of the distributions. There have been several follow-up works such as [BV15, AJOS14, DGK<sup>+</sup>21]. For the particular setting of Problem 5, where  $b$  is treated as an input parameter, it is known that  $\text{SC}_{\text{EQIV}}(\ell_2, \varepsilon, b, n) = \Theta(b/\varepsilon^2)$ ,  $\text{SC}_{\text{EQIV}}(\ell_1, \varepsilon, b, n) = O(nb/\varepsilon^2)$  [CDVV14, Thm. 2]. We note that we can replace the ‘max’ with a ‘min’ at the cost of an additive factor  $O(\sqrt{d})$  in the sample complexity, in which case the algorithm first tests whether  $\|P\|_2 = \Theta(\|Q\|_2)$  [DK16, Lemma 2.3].

Note that in all of these problem statements, the probability 2/3 can be replaced by any target probability larger than 1/2 as we are not interested in the dependence on this confidence parameter and the success probability can be amplified by repeating the tester a constant number of times.

## 2.2 Our Results in Context

In the following, we present our main results. The first result establishes the exact sample complexity for MI testing.

**Result 1** (MI testing). *Consider Problem 2 and assume  $d_A \geq d_C$ . Then,*

$$\text{SC}_{\text{MI}}(\varepsilon, d_A, d_C) = \tilde{\Theta} \left( \min \left\{ \frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \right). \quad (8)$$

Here  $\tilde{\Theta}$  hides terms that are poly-logarithmic in the problem parameters. The initial study of this problem [BGP<sup>+</sup>23] considered only the case  $d_A = d_C = d$  with a sample complexity in  $\tilde{O}(d^2/\varepsilon)$ . Later, [FO24, Theorem 1.18] improved this bound to  $\tilde{O}(d/\varepsilon)$ . If we adopt the computations in [FO24] with independent alphabet sizes, this gives us a bound of  $\tilde{O}((d_A + d_C)/\varepsilon)$  samples. Using the techniques described in Section 4, it is easy to see that equivalence testing in the squared Hellinger distance can also be used for (conditional) independence testing. For two unknown distributions of dimension  $d$ , equivalence testing has a sample complexity of  $\Theta(\min\{d_A^{3/4} d_C^{3/4}/\varepsilon, d_A^{2/3} d_C^{2/3}/\varepsilon^{4/3}\})$ , [DKW18, Thm. 4.2] and [DK16, Prop. 3.8], which results in a bound of  $\tilde{O}(\min\{d_A^{3/4} d_C^{3/4}/\varepsilon, d_A^{2/3} d_C^{2/3}/\varepsilon^{4/3}\})$  for independence testing. The similarity between this sample complexity and our Result 1 is not surprising, as independence testing is equivalent to equivalence testing where one of the distributions is guaranteed to having product structure (see Corollary 5.2).

Moreover, combining the independence testing result of [DK16] in total variation distance along with [CDKS18, Lemma A.1] and assuming  $d_A \geq d_C$ , this bound can be further improved to

$$\tilde{O} \left( \min \left\{ \frac{d_A}{\varepsilon}, \frac{d_A^{3/4} d_C^{3/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{2/3}}{\varepsilon^{4/3}}, \max \left\{ \frac{d_A^{1/2} d_C^{1/2}}{\varepsilon^2}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \right\} \right). \quad (9)$$

Comparing this best previous bound with our bound in Result 1, we notice an improvement unless  $d_C = \Theta(d_A)$ , in which case the result reported in [FO24] is already optimal. Moreover, we also prove that the dependencies on  $d_A$  and  $d_C$  are tight, ignoring poly-logarithmic factors. For the lower bounds, we use an information theoretic approach akin to the technique used in [DK16] to prove lower bounds on the problem in variation distance. To the best of our knowledge, this is the first sample optimal mutual information tester in the literature.

Now let us proceed to present our main results on CMI testing.

**Result 2** (CMI testing, upper bound). *Consider Problem 4 and assume  $d_A \geq d_C$ . Then,*

$$\text{SC}_{\text{CMI}}(\varepsilon, d_A, d_B, d_C) = \tilde{O} \left( \max \{ f_{\text{sym}}(\varepsilon, d_A, d_B, d_C), f_{\text{asym}}(\varepsilon, d_A, d_B, d_C) \} \right), \quad (10)$$

where we distinguish between terms that are symmetrical and asymmetrical in  $A$  and  $C$ ,

$$f_{\text{sym}}(\varepsilon, d_A, d_B, d_C) := \max \left\{ \frac{d_A^{1/2} d_B^{3/4} d_C^{1/2}}{\varepsilon}, \min \left\{ \frac{d_A^{1/4} d_B^{7/8} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/7} d_B^{6/7} d_C^{2/7}}{\varepsilon^{8/7}} \right\} \right\}, \quad \text{and} \quad (11)$$

$$f_{\text{asym}}(\varepsilon, d_A, d_B, d_C) := \min \left\{ \frac{d_A^{3/4} d_B^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_B^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\}. \quad (12)$$

The prior work in [BGP<sup>+</sup>23] showed a sample complexity of  $\tilde{O}(d^3/\varepsilon)$ , assuming all alphabets to be of the same size. As for mutual information testing, prior work on testing for conditional



independence in variation distance [CDKS18] can again be translated to a bound on conditional mutual information testing [CDKS18, Lemma A.1], and yields a sample complexity of

$$\tilde{O} \left( \max \left\{ \frac{d_B^{1/2} d_A^{1/2} d_C^{1/2}}{\varepsilon^2}, \frac{d_B^{2/3} d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}}, \frac{d_B^{3/4} d_A^{1/2} d_C^{1/2}}{\varepsilon}, \min \left\{ \frac{d_B^{7/8} d_A^{1/4} d_C^{1/4}}{\varepsilon}, \frac{d_B^{6/7} d_A^{2/7} d_C^{2/7}}{\varepsilon^{8/7}} \right\} \right\} \right). \quad (13)$$

The resulting algorithm will in principle show an  $\varepsilon$ -scaling in  $\tilde{O}(1/\varepsilon^2)$ , but has the advantage of a sublinear dimensional scaling compared to [BGP+23]. Our algorithm manages to materialize the best of both worlds, achieving a linear scaling in  $\tilde{O}(1/\varepsilon)$ , as well as a sublinear dimensional scaling.

An interesting immediate observation is that in the regimes which are not dominated by  $\varepsilon$ , our sample complexity coincides with (13). It is important to stress that, analogous to the case of independence testing, the underlying approaches to solve the two problems vary significantly: the authors in [CDKS18] studied the behavior of a polynomial defined over a probability distribution, and used this polynomial to design an efficient tester for their problem. In a sense, they follow an algebraic approach. While their technique provides a certain generality and can in principle be applied to other learning problems, it is currently not known whether their approach can be modified to obtain bounds for our problem which improve beyond (13). This is not unexpected when we compare to the simpler problem of independence testing: the known (sample optimal) algorithms for MI-testing and testing independence with respect to variational distance use very different approaches as well. Our approach reduces the problem to a polylogarithmic number of instances of equivalence testing in  $\ell_2$  distance between an unknown distribution and a reference distribution simulated from the unknown distribution.

To obtain lower bounds, we can combine a reduction to MI-testing with existing lower bounds on independence testing coming from variation distance [CDKS18] to derive the following lower bound for CMI testing.

**Result 3** (CMI testing, lower bound). *Consider Problem 4 and assume  $d_A \geq d_C$ . Then,*

$$\text{SC}_{\text{CMI}}(\varepsilon, d_A, d_B, d_C) = \tilde{\Omega} \left( f_{\text{sym}}(\varepsilon, 1, d_B, 1), f_{\text{asym}}(\varepsilon, d_A, d_B, d_C) \right). \quad (14)$$

Comparing this with Result 2, we see that in the symmetric term, we are missing the dependence on  $d_A$  and  $d_C$ . For the special case where  $\varepsilon = \Omega(1)$  and  $d = d_A = d_B = d_C$ , [CDKS18] provides a lower bound of  $\Omega(d^{7/4})$ . While there remain open questions on the exact form of the lower bounds, these findings strongly indicate that the complicated structure of the sample complexity we report in Result 2 is indeed required.

Our algorithm for CMI testing uses a new estimator, and we show that we can recover the aforementioned bounds for equivalence testing in the  $\ell_1$  and  $\ell_2$  distance [CDVV14, Thm. 2] and [DK16, Lemma 2.3], Problem 5, in Section 9.

$$\text{SC}_{\text{EQIV}}(\ell_1, \varepsilon, b, n) = O \left( \frac{bn}{\varepsilon^2} \right), \quad \text{SC}_{\text{EQIV}}(\ell_2, \varepsilon, b, n) = O \left( \frac{b}{\varepsilon^2} \right). \quad (15)$$

We expect that our estimator will show favorable properties when considering data which violates the i.i.d. assumption.

## 2.3 Methods

Here we would like to give an overview of the methods we used to prove our results.

**Reduction to testing in Hellinger distance:** The first step in solving the problem of testing for (C)MI is to reduce it to the testing for (conditional) independence with respect to the squared Hellinger distance ( $D_H^2$  in short). For this purpose, we use the following inequalities from [FO24, Proposition 2.12] and [GS02, p. 429] (see Lemma 3.1):

$$D_H^2(P, Q) \leq D(P\|Q) \leq \left(2 + \log \left( \max_{i \in [d], P(i) \neq 0} \frac{P(i)}{Q(i)} \right)\right) D_H^2(P, Q). \quad (16)$$

In order to use this inequality gainfully, we need to ensure a lower bound on  $Q_{\min} := \min_{i \in [d]} Q[i]$ . We achieve this by taking a mixture of the reference distribution with the uniform distribution such that  $Q_{\min}$  does not become too small, while also preserving (conditional) independence and  $\varepsilon$ -farness between distributions (up to a constant factor). This requires a specific continuity result for both the KL-divergence and the conditional mutual information which we prove in Section 4.

**MI Testing:** To perform independence testing in  $D_H^2$ , we note that sample access to  $P_{AC}$  allows us to directly simulate samples from  $P_A P_C$ . The question is then whether these two distributions are the same or far from each other. In fact, we solve a more general problem by testing for equivalence between  $P_{AC}$  and an arbitrary product distribution  $Q_A Q_C$ . Our matching lower bounds show that the problems have the same asymptotic sample complexity. A known idea for equivalence testing is to reduce the problem from the  $D_H^2$  metric to testing in  $\ell_2$  distance (see [CDVV14, Theorem 2], [DK16]), which is achieved by splitting a distribution into a logarithmic number of buckets such that all elements in a bucket have *roughly* the same weight. This allows us to tightly bound  $D_H^2$  in terms of  $\ell_2$  distance. Each bucket can then individually be tested for equivalence in  $\ell_2$  distance. One bucket,  $S_{\text{small}}$  contains all small weights and requires special treatment.

We introduce a more fine grained version of this approach by leveraging the fact that our reference distribution is a product distribution, illustrated in Figure 1. This allows us to refine the bucketing, by effectively creating a two-dimensional grid of buckets, one dimension for  $A$  and  $C$  each, which ultimately leads to an improved sample complexity. Again, buckets with small weights (i.e. either  $p_a$  or  $p_c$  are small) need to be treated separately. The entire argument is laid out and formally proved in Section 5. Our approach explains the origin of the two regimes in the sample complexity: the difficulty of equivalence testing on a bucket  $S$  depends on the  $\ell_2$  norm of  $P$  on  $S$ ,  $(\sum_{i \in S} P(i)^2)^{1/2}$ , which we can in general bound in two ways,  $(\sum_{i \in S} P(i)^2)^{1/2} \leq \min\{\sqrt{|S|}p_{\max}, \sqrt{p_{\max}}\}$ , where  $p_{\max} := \max_{i \in S} \{P(i)\}$ , giving rise to the two regimes in our results. Since the threshold for which indices are placed in the bucket of small weights,  $S_{\text{small}}$ , depends on the number of samples  $N$  we use, the  $\ell_2$  norm of  $P$  on  $S_{\text{small}}$  depends on  $N$  again: solving the resulting expression for  $N$  then gives us the sample complexity in Result 1 and explains the seemingly counterintuitive exponents in the sample complexity of our results.




Fig. 1: Partition of  $d_A \times d_C$  based on  $\hat{Q}_A$  and  $\hat{Q}_C$ . Indices  $(a, c)$  of similar weight  $\hat{q}_a \hat{q}_c$  are grouped together in buckets  $S_{ij}$ , which are used to perform piecewise equivalence testing with  $P_{AC}$ . The axes are labeled according to the corresponding category, which is inverse logarithmic to the weight of the probabilities. The color of the categories indicate a different analysis of the sample complexity of the categories. The red regimes dominate the sample complexity.

**CMI Testing:** Similar to MI testing, we reduce CMI testing to testing for conditional independence in  $D_H^2$ . However, unlike for MI, we can no longer sample directly from the reference distribution  $Q_{ABC} := P_{AB}P_{C|B}$ , which is crucial for us. We overcome this bottleneck by performing a case distinction depending on whether  $p_b$  for a given  $b \in B$  is large or small. Both regimes are tested separately to see whether the conditional mutual information is zero or at least  $\varepsilon/2$ . The sample complexity will be dominated by the fact that we need to generate enough samples for both regimes. This is illustrated in Figure 3.

- To generate samples from the first, ‘large’ regime, we use two phases of taking samples. In the first phase, we take  $\tilde{O}(N)$  samples and sort their  $A$ -coordinates into queues depending on their  $b$ -values. In the second phase, we take  $O(N)$  more samples: if we draw a sample  $(a, b, c)$  from the large regime, we remove an element from the corresponding  $b$ -queue, say  $a'$ , and we output  $(a', b, c)$ . The probability to observe  $(a', b, c)$  is easily seen to equal  $p_{a'|b}p_{bc}$ , which corresponds exactly to our reference distribution  $Q_{ABC}$ . Concentration bounds guarantee that with high probability, enough large samples are available in the respective queues from the reference distribution.

For testing, we follow a similar approach as for independence testing. Now, instead of the two-dimensional bucketing from MI testing, here the bucketing needs to be performed in three dimensions, as shown in Figure 2. As for the MI-testing, the sample complexity will be dominated by accounting for elements with smaller probability masses. The case distinction in the sample complexity follows from bounding the  $\ell_2$  norm of specific buckets in two different ways.




Fig. 2: Partition of  $d_A \times d_B \times d_C$  based on  $\hat{P}_{AB}$ ,  $\hat{P}_{BC}$ , and  $\hat{P}_B$ . Indices  $(a, b, c)$  of similar weight  $\hat{p}_{ab}\hat{p}_{bc}/\hat{p}_b$  are grouped together in categories  $L_{ij}^k$ , which are used to perform piecewise equivalence testing with  $P_{ABC}$ . The axes are labeled according to the corresponding category, which is inverse logarithmic to the weight of the probabilities. The color of the categories indicate a different analysis of the sample complexity of testing the categories. The red and orange regimes dominate the sample complexity. The small regime is treated separately.

- Unlike the previous case, in the ‘small’ regime, we can no longer guarantee how often a *particular* rare  $b$ -value will appear. However, we know that if we take enough samples, *some* of the rare elements will appear more than once, and the *overall* number of rare collisions can be quantified using concentration bounds. This motivates us to take enough samples to obtain multiple collisions with rare elements, and use them to (approximately) simulate samples from  $P_{AB}P_{C|B}$  by combining  $(a, b, c)$ ,  $(a', b, c')$  into  $(a, b, c')$ . The procedure is described in Figure 3, and comes with a few caveats when compared to the approach used in the large regime, both due to the conditioning on seeing the respective  $b$  twice.

First, the resulting samples are not independent, but slightly correlated. As a result, we can no longer use the same technique for equivalence testing which we used in the large regime. Instead, we introduce a novel tester for equivalence testing which is robust with respect to the correlations induced by our sampling approach. This is one of the technical novelties in this work, and we believe it will be of independent interest to the community, which we describe in more detail below. Second, the frequency at which we see a particular  $b$  will be biased. This is due to the fact that a sample with a rare  $b$  element appears with probability  $p_b$ , but a collision appears only with a probability proportional to  $p_b^2$ . Unsurprisingly, skewing  $P_B$  will affect the statistics of the resulting samples, which makes testing for conditional independence more costly, as an improved precision is required. However, the conditional probabilities of witnessing a specific  $(a, c)$  given  $b$ ,  $P_{AC|B=b}$ , remain unchanged.

One might try to avoid using a case distinction, and extend the first regime to cover all  $p_b$  which might be relevant in testing for conditional mutual information. However, it is easy to

construct examples where this would necessitate the approximate learning of probabilities  $p_b$  in  $O(\varepsilon/d_B)$ , which would require  $O(d_B/\varepsilon)$  samples and is as costly as learning  $P_B$ . Our case distinction allows us to achieve a sample complexity with sublinear scaling in  $d_B$ .

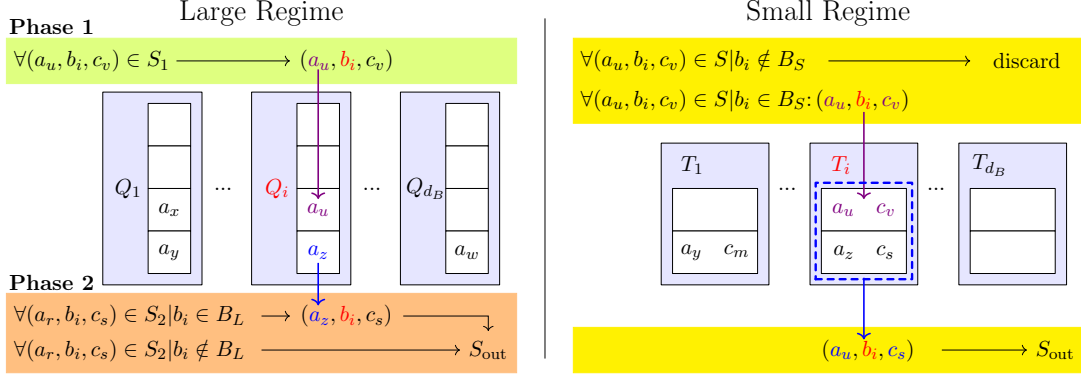


Fig. 3: (Comparison of sampling methods). In the large regime ( $b \in B_L$ ), sampling is split into two phases using separate sets  $S_1$  and  $S_2$  of samples. In the first phase, we assign  $A$ -coordinates of samples to queues  $Q_i$ , depending on the  $B$ -coordinate. In the second phase, the actual output is generated by taking samples and replacing the  $A$ -coordinate with an element from the respective  $B$ -queue, if  $b_i \in B_L$ . The samples in  $S_{out}$  can be seen as drawn i.i.d. from a distribution which coincides with  $P_{AB}P_{C|B}$ . In the small regime ( $b \in B_S$ ), we process only samples from  $B_S$ , which we sort into tuples  $T_i$ , dependent on the  $B$ -coordinate. As soon as a tuple is filled, we generate an output. This way of sampling skews the statistics of  $B$  and the resulting samples are not independent.

In general, our new tester for equivalence testing takes two pairs of independent multisets of samples  $X, X'$  and  $Y, Y'$  generated by two random processes  $\mathcal{X}$  and  $\mathcal{Y}$ , respectively, which in particular might cause correlations within the respective multisets. Let  $X_i$  denote the number of times the element  $i \in A \times B \times C$  has appeared in set  $X$ , and define  $Y_i, X'_i$  and  $Y'_i$  analogously. Our estimator then constructs the following variable  $Z$ , which is compared against a threshold.

$$Z = \sum_i Z_i, \quad Z_i := X_i X'_i - 2X_i Y_i + Y_i Y'_i. \quad (17)$$

A crucial property of this estimator is that if  $X$  and  $Y$  follow the same statistics, then  $\mathbb{E}[Z] = \sum_i (\mathbb{E}[X_i] - \mathbb{E}[Y_i])^2 = 0$ . This property is very useful, as it allows us to build a tester by bounding  $\gamma \leq \mathbb{E}[Z]$  for the case where  $\mathcal{X}$  and  $\mathcal{Y}$  differ, and defining the threshold at  $\gamma/2$ . The sample complexity is determined by the variance of  $Z$ . In our application in the small regime, a detailed analysis of the specific correlations between the samples is necessary to tightly bound the variance of  $Z$ .

This tester can also be used for general equivalence testing, in which case  $\mathcal{X}$  and  $\mathcal{Y}$  correspond to sampling from distributions  $P$  and  $Q$ . This is discussed in Section 9, recovering the best known bounds in  $\ell_1$  distance, [DK16, Lemma 2.3], and sample optimal bounds in  $\ell_2$  distance [CDVV14, Theorem 2]. To the best of our knowledge, this estimator is not present in the current literature and will be of independent interest.

**Lower Bounds:** The primary idea behind proving lower bounds for MI testing is to use Le Cam's two-point method. We define two sets of distributions over distributions  $S_{yes}$  and  $S_{no}$  such that the

distributions in  $S_{\text{yes}}$  are independent, and the distributions in  $S_{\text{no}}$  are “*far*” from being independent in the squared Hellinger distance. Moreover, the distributions are designed such that given a random distribution from  $S_{\text{yes}} \cup S_{\text{no}}$ , deciding whether it came from  $S_{\text{yes}}$  or  $S_{\text{no}}$  is difficult. Lower bounds to this task imply bounds on MI testing, since an MI tester could be used to solve this problem.

To obtain a bound on the sample complexity, we follow the mutual information route, introduced in [DK16]. In a nutshell, the goal is to prove that the mutual information between a set of samples obtained from the unknown distribution  $P$  (chosen randomly from either  $S_{\text{yes}}$  or  $S_{\text{no}}$ ), and the distribution itself is small. As a consequence, unless a sufficient number of samples is drawn, it is difficult to distinguish whether  $P$  was drawn from either  $S_{\text{yes}}$  or  $S_{\text{no}}$ . To connect the mutual information between the set of samples, we rely on the following inequality:

$$2I(X : A) \leq \sum_{a \in S_A} \frac{(\Pr[A = a|X = 0] - \Pr[A = a|X = 1])^2}{\Pr[A = a|X = 0] + \Pr[A = a|X = 1]} \leq 12I(X : A). \quad (18)$$

where  $X$  is a uniform random bit, and  $A$  is a random variable taking values in a set  $S_A$ . The first inequality was known (e.g. [DK16]), but to our knowledge, the second inequality does not appear in the literature in this form, and might be of independent interest. The specific construction we use extends distributions utilized in [DK16] for proving lower bounds on equivalence testing. There, each index of the distribution is assigned one of two types, with one being identical for both instances, and thus not carrying information (this, in a sense, acts as noise which makes distinguishing them harder). We keep this structure in the dimension  $d_A$ , and simply expand it in an additional direction  $d_C$ , such that, for a fixed  $a$ , all  $(a, c)$  belong to the same type. The asymmetry of  $d_A$  and  $d_C$  in the sample complexity is thus also reflected in the construction leading to our lower bounds.

Our lower bound result for CMI testing, presented in Result 3, is a reduction to our approach of independence testing described above.

### 3 Preliminaries

In this work, we will be using the following notations. We denote the set  $\{1, \dots, n\}$  as  $[n]$ , and  $[n] \cup \{0\}$  as  $[n]_0$ . Similarly, let  $\mathbb{N}_0$  denote the set  $\mathbb{N} \cup \{0\}$ . For a positive integer  $t \in \mathbb{N}^+$ , and a set  $\mathcal{S} \subset \mathbb{N}_0^+$ , a vector  $a \in \mathcal{S}^{\times t}$  denotes an ordered list  $(a_1, \dots, a_t)$  where for every  $i \in [t]$ ,  $a_i \in \mathcal{S}$ . For a vector  $v \in \mathcal{S}^{\times t}$ ,  $\|v\|_1 := \sum_{i \in [t]} v_i$  denotes the sum of the entries in  $v$ . For a parameter  $p \in (0, 1)$ ,  $\text{Ber}(p)$  denotes the Bernoulli distribution over  $\{0, 1\}$ , where for a random variable  $X$  drawn from  $\text{Ber}(p)$ , it holds that  $\Pr[X = 1] = p$ .

Throughout this work,  $P_{ABC}$  denotes an unknown discrete distribution defined over  $A \times B \times C$ , and we denote  $|A| = d_A$ ,  $|B| = d_B$ ,  $|C| = d_C$ . When studying conditional independence, we denote the conditioning system by  $B$ . Due to the symmetry of the problem with respect to the non-conditioning systems, we may assume  $d_A \geq d_C$  without loss of generality. For a set  $\mathcal{S}$ ,  $U_{\mathcal{S}}$  denotes the uniform probability distribution over  $\mathcal{S}$ . For a probability distribution  $P_{\mathcal{D}}$  defined on a set  $\mathcal{D}$ , we will often interchangeably denote  $p_x := P_{\mathcal{D}}(x)$  for some  $x \in \mathcal{D}$ . By stating i.i.d. samples from a distribution  $P_{\mathcal{D}}$ , we mean independently and identically distributed samples from  $P_{\mathcal{D}}$ . We also introduce the following notation for a probability distribution reduced to a subset  $\mathcal{S}$  of its domain  $\mathcal{D}$ :

$$P_{\mathcal{D}}^{\mathcal{S}}(x) := \begin{cases} P_{\mathcal{D}}(x) & \text{if } x \in \mathcal{S}, \\ 0 & \text{otherwise.} \end{cases} \quad (19)$$

The weight of  $P_{\mathcal{D}}$  on  $\mathcal{S}$  is denoted by  $P_{\mathcal{D}}[\mathcal{S}] := \sum_{i \in \mathcal{S}} p_i$ . Note in particular that  $P_{\mathcal{D}}^{\mathcal{S}}$  is in general not normalized, i.e.,  $\|P_{\mathcal{D}}^{\mathcal{S}}\|_1 = P_{\mathcal{D}}[\mathcal{S}]$ .

For concise expressions and readability, we use the asymptotic complexity notion of  $\tilde{O}(\cdot)$ ,  $\tilde{\Omega}(\cdot)$ , and  $\tilde{\Theta}(\cdot)$ , where we hide poly-logarithmic dependencies on the parameters. Throughout this work the logarithm  $\log(\cdot)$  will denote the natural logarithm, and  $\sinh(\cdot)$  and  $\cosh(\cdot)$  denote the sine and cosine hyperbolic functions, respectively.

Next we define the different distance measures we will use to relate distributions. Let  $\mathcal{P}(\mathcal{D})$  be the set of probability distributions over the elements of a set  $\mathcal{D}$ . Let  $P$  and  $Q$  be two such distributions over  $\mathcal{D}$  (we drop the  $\mathcal{D}$  in the subscript when it is clear from the context):

- The *KL-divergence* between  $P$  and  $Q$  is defined as:  $D(P\|Q) = \sum_{x \in \mathcal{D}} P(x) \log \frac{P(x)}{Q(x)}$ . We assume that the support of  $Q$  contains the support of  $P$ , and use the convention  $0 \log 0 = 0$  to deal with zeros, so that the KL-divergence is always finite.
- The *squared Hellinger* distance between  $P$  and  $Q$  is defined as:

$$D_H^2(P, Q) = \frac{1}{2} \sum_{x \in \mathcal{D}} \left( \sqrt{P(x)} - \sqrt{Q(x)} \right)^2. \quad (20)$$

- The  $\ell_p$  *distance* between  $P$  and  $Q$  is defined as  $\|P - Q\|_p$ , where the  $\ell_p$ -*norm* is given as  $\|A\|_p := \left( \sum_{x \in \mathcal{D}} |A(x)|^p \right)^{1/p}$  for any  $p \geq 1$ .

Let us next state a lemma which connects the KL-divergence and the squared Hellinger distance, which will be crucial for our work.

**Lemma 3.1** ([GS02, p. 429] & [FO24, Proposition 2.12]). *Let  $P$  and  $Q$  be two probability distributions over  $[d]$ . Then we have*

$$D_H^2(P, Q) \leq D(P\|Q) \leq \left( 2 + \log \left( \max_{i \in [d], P(i) \neq 0} \frac{P(i)}{Q(i)} \right) \right) D_H^2(P, Q). \quad (21)$$

Note that we can simply bound this further to  $D(P\|Q) \leq 3 \log(1/Q_{\min}) D_H^2(P, Q)$ , where  $Q_{\min} := \min_{i \in [d]} Q(i)$  if we assume  $d \geq 3$ . We also make use of the following folklore relations between different distance measures (see e.g. [GS02, Eq. 8] for (i), and (ii) follows from Jensen's inequality).

**Fact 3.2.** *For arbitrary distributions  $P$  and  $Q$  of dimension  $d$ , it holds that*

$$(i) \quad \frac{1}{2} D_H^2(P, Q) \leq \|P - Q\|_1 \leq D_H(P, Q),$$

$$(ii) \quad \|P - Q\|_1 \leq \sqrt{d} \|P - Q\|_2.$$

In our analysis, we will also use different bounds on the  $\ell_2$  norm of a distribution.

**Fact 3.3.** *Let  $Q$  be an arbitrary probability distribution defined over  $[N]$  and  $\mathcal{S} \subseteq [N]$  be a subset of the support. Then, we have*

$$\|Q^{\mathcal{S}}\|_2^2 \leq \begin{cases} |\mathcal{S}| \max_{i \in \mathcal{S}} q_i^2 & [\text{Case 1}], \\ \max_{i \in \mathcal{S}} q_i & [\text{Case 2}]. \end{cases}$$

We will use the following two concentration bounds (see, e.g., [DP09, Thm. 1.1] and [MR95, Thm. 3.3, 4.1 & 4.2]).

**Lemma 3.4** (Chernoff Bound). *Let  $X_1, \dots, X_n$  be independent random variables such that  $X_i \in [0, 1]$ . For  $X = \sum_{i=1}^n X_i$ , the following holds for any  $0 \leq \delta \leq 1$ .*

$$\Pr[X \geq (1 + \delta)\mathbb{E}[X]] \leq e^{-\delta^2\mathbb{E}[X]/3}, \quad \Pr[X \leq (1 - \delta)\mathbb{E}[X]] \leq e^{-\delta^2\mathbb{E}[X]/2}. \quad (22)$$

**Lemma 3.5** (Chebyshev's inequality). *Let  $X$  be a random variable with  $\mathbb{E}[X^2] < \infty$ . The following holds for any  $t > 0$ .*

$$\Pr[|X - \mathbb{E}[X]| \geq t] \leq \frac{\text{Var}[X]}{t^2}. \quad (23)$$

Besides concentration bounds, another crucial ingredient of our techniques is the notion of Poissonization technique. Let us start with the definition of Poisson random variables.

**Definition 3.6** (Poisson random variable). Let  $\lambda > 0$ . A discrete random variable  $X$  is said to be a Poisson random variable with parameter  $\lambda$ , denoted as  $\text{Poi}(\lambda)$  if the following holds:

$$\forall k \in \mathbb{N}, \Pr[X = k] = e^{-\lambda} \frac{\lambda^k}{k!}. \quad (24)$$

Now we are ready to describe the Poissonization technique (see, e.g., [Can20, Fact D.10]).

**Lemma 3.7** (Poissonization technique). *Let  $\Omega$  be a discrete domain,  $D \in \Delta(\Omega)$  be a distribution and  $m \in \mathbb{N}$ . Suppose  $M' \sim \text{Poi}(m)$  independent samples  $s_1, \dots, s_{M'}$  have been obtained from  $D$ . Suppose  $X_t$  denotes the number of times  $t \in \Omega$  appears among the samples  $s_1, \dots, s_{M'}$ . Then the following hold:*

- (i)  $(X_t)_{t \in \Omega}$  are independent.
- (ii)  $X_t \sim \text{Poi}(mD(t))$ .

Poisson random variables follow the following concentration bound (see, e.g., [Can22, Theorem A.8]).

**Lemma 3.8** (Poisson concentration bound). *Let  $X$  be a  $\text{Poi}(\lambda)$  random variable for some  $\lambda > 0$ . Then for any  $t > 0$ , the following holds:*

$$\Pr[|X - \lambda| \geq t] \leq 2e^{-\frac{t^2}{2(\lambda+t)}}. \quad (25)$$

Finally, we will use the following bounds on the exponential function in our analysis, proved in Appendix A.1.

**Lemma 3.9.** *For  $0 \leq x < 1$ , it holds that*

$$1 - x + \frac{x^2}{4} \leq e^{-x} \leq 1 - \frac{x}{2}. \quad (26)$$

Moreover, for any  $x \geq 0$ , we have that

$$1 - x + \frac{x^2}{2} - \frac{x^3}{6} \leq e^{-x} \leq 1 - x + \frac{x^2}{2}. \quad (27)$$



## 4 Reducing CMI Testing to Testing in $D_H^2$

In this section, we present how (conditional) mutual information testing can be reduced to the problem of (conditional) independence testing with respect to the squared Hellinger distance, which we formalize in the following two problems. The reduction itself is the subject of Theorem 4.2.

**Problem 6** (Conditional independence testing in  $D_H^2$ ). Consider the setting of conditional independence testing (Problem 3), where  $\Delta(\|\cdot\|) = D_H^2(\cdot, \cdot)$ , that is, for fixed a threshold  $\nu > 0$ , we want to distinguish between the classes

- (i)  $D_H^2(P_{ABC}, P_{AB}P_{C|B}) = 0$ , and
- (ii)  $D_H^2(P_{ABC}, P_{AB}P_{C|B}) \geq \nu$ .

We denote the sample complexity of this problem by  $SC_{CI,H}(\nu, d_A, d_B, d_C)$ .

We can also define the corresponding independence testing problem by choosing  $B$  to be trivial.

**Problem 7** (Independence testing in  $D_H^2$ ). Consider the setting of conditional independence testing (Problem 1), where  $\Delta(\|\cdot\|) = D_H^2(\cdot, \cdot)$ , that is, for fixed a threshold  $\nu > 0$ , we want to distinguish between the classes

- (i)  $D_H^2(P_{AC}, P_AP_C) = 0$ , and
- (ii)  $D_H^2(P_{AC}, P_AP_C) \geq \nu$ .

We denote the sample complexity of this problem by  $SC_{I,H}(\nu, d_A, d_C)$ .

The reduction from (conditional) independence testing to a testing problem in  $D_H^2$  distance only requires a preprocessing of the samples. We can then apply any algorithm to solve Problem 6, by slightly increasing the precision  $\nu$  to which we test. Thus, the sample complexity of CMI testing is, up to logarithmic factors, the same as the sample complexity of conditional independence testing in the squared Hellinger distance.

As we will show in the following, the key idea is to use the following inequality (see Lemma 3.1) which relates the KL-divergence and the squared Hellinger distance, with the choice  $P = P_{AC|B}$  and  $Q = P_{A|B}P_{C|B}$ ,

$$D(P\|Q) \leq \left(2 + \log\left(\frac{1}{Q_{\min}}\right)\right) D_H^2(P, Q). \quad (28)$$

Here,  $Q_{\min}$  denotes the minimum probability mass of any element in the distribution  $Q$ . In order to use this inequality, we need to have a guarantee on the minimum probability mass of every element of  $Q_{AC|B}$ . We ensure this by modifying our sampling approach: for both  $A$  and  $C$ , we independently take a weak mixture with the uniform distribution. This action preserves conditional independence and ensures the minimum probability mass on  $P_{AC|B=b}$ . Our sampling is as follows:

---

**Algorithm 1:** Sampling from  $\tilde{P}$ 

---

**Input:** A triplet  $(a, b, c)$  from  $A \times B \times C$ ,  $\eta \in (0, 1)$ .

**Output:** A triplet from  $A \times B \times C$

- 1  $(a', b', c') \leftarrow (a, b, c)$
  - 2 Sample  $X_A, X_C \stackrel{\$}{\leftarrow} \text{Ber}(\eta)$
  - 3 **if**  $X_A = 1$  **then**
  - 4      $a' \stackrel{\$}{\leftarrow} U_A$
  - 5 **if**  $X_C = 1$  **then**
  - 6      $c' \stackrel{\$}{\leftarrow} U_C$
  - 7 **return**  $(a', b', c')$ .
- 

It is immediate that the above algorithm preserves Markovianity since we can see  $\tilde{P}_{ABC}$  as the result of an application of a stochastic map acting on  $A$  and  $C$ , respectively. While our preprocessing ensures minimum probability mass in the conditional distributions, it might also change the conditional mutual information when it is non-zero. For this reason, we also need to guarantee that  $I(A:C|B)_{\tilde{P}} \geq I(A:C|B)_P/4$ . Our argument will use the following continuity result.

**Lemma 4.1.** *Let  $P_{AC}$  and  $T_{AC}$  be two arbitrary distributions over  $A \times C$  and  $\varepsilon \in (0, 1)$  be a threshold. Then, for any  $\alpha \leq \left(\frac{\varepsilon}{48d_A d_C \log(d_A d_C / \varepsilon)}\right)^2$ , it holds that*

$$D(P_{AC} \| P_A P_C) \geq \varepsilon \implies D((1 - \alpha)P_{AC} + \alpha T_{AC} \| (1 - \alpha)P_A P_C + \alpha T_{AC}) \geq \varepsilon/2. \quad (29)$$

The proof is a direct calculation which bounds the difference between the two terms in (29) using a case distinction and basic bounds on the logarithm. We prove this in Appendix A.2. We can now formally argue for the reduction.

**Theorem 4.2.** *CMI testing (Problem 3) can be solved by preprocessing samples of  $P$  using Algorithm 1 with  $\eta = \tilde{O}(\varepsilon^2 / (d_A d_C)^2)$ , and then applying any algorithm for Problem 6 to precision  $\nu = \varepsilon / (8 \log(d_A d_C / \eta^2))$ . In particular, we have*

$$\text{SC}_{\text{CMI}}(\varepsilon, d_A, d_B, d_C) \leq O(\text{SC}_{\text{CI}, H}(\nu, d_A, d_B, d_C)). \quad (30)$$

Analogously,  $\text{SC}_{\text{MI}}(\varepsilon, d_A, d_C) \leq O(\text{SC}_{\text{I}, H}(\nu, d_A, d_C))$ .

*Proof of Theorem 4.2.* We will only argue for the conditional case, as the result for MI testing follows from the case when  $B$  is trivial. We first note that  $I(A:C|B)_P = 0$  is equivalent to  $D_H^2(\tilde{P}_{ABC} \| \tilde{P}_{AB} \tilde{P}_{C|B}) = 0$ , since Algorithm 1 preserves Markovianity.

Now assume  $I(A:C|B)_P \geq \varepsilon$ . The proof consists of two steps. We will first examine the effect of our preprocessing on  $I(A:C|B)_P$  using Lemma 4.1, and then apply Lemma 3.1 to arrive at a testing problem in  $D_H^2$  distance. For brevity, let us denote  $P_{AC}^b := P_{AC|B=b}$ ,  $P_A^b := P_{A|B=b}$  and  $P_C^b = P_{C|B=b}$ . We first split  $B$  into two sets,  $B_+ := \{b \in B : D(P_{AC}^b \| P_A^b P_C^b) \geq \varepsilon/2\}$  and  $B_- := B \setminus B_+$ . Then

$$I(A:C|B)_P = \sum_{b \in B} p_b D(P_{AC|B=b} \| P_{A|B=b} P_{C|B=b}) \quad (31)$$

$$= \sum_{b \in B_+} p_b D(P_{AC}^b \| P_A^b P_C^b) + \sum_{b \in B_-} p_b D(P_{AC}^b \| P_A^b P_C^b) \quad (32)$$

$$\leq \sum_{b \in B_+} p_b D(P_{AC}^b \| P_A^b P_C^b) + \frac{\varepsilon}{2}, \quad (33)$$

Thus,  $I(A:C|B)_P \geq \varepsilon$  implies

$$\sum_{b \in B_+} p_b D(P_{AC}^b \| P_A^b P_C^b) \geq \frac{\varepsilon}{2}. \quad (34)$$

For each  $b \in B_+$ , we would now like to individually apply Lemma 4.1. In order to do so, we first introduce a probability distribution  $T_{ABC}$  over  $A \times B \times C$ , defined as

$$t_{abc} := \frac{p_b}{2 - \eta} \left( (1 - \eta) \left( \frac{p_{c|b}}{d_A} + \frac{p_{a|b}}{d_C} \right) + \frac{\eta}{d_A d_C} \right), \quad \text{for all } a \in A, b \in B, c \in C. \quad (35)$$

Note that  $T_{ABC}$  is a valid probability distribution, since  $\forall (a, b, c) \in A \times B \times C : t_{abc} \geq 0$  and  $\sum_{a,b,c} t_{abc} = 1$ . The conditioned distribution  $T_{AC|B}$  has a guarantee on the minimum probability mass since  $t_{ac|B=b} \geq \eta / ((2 - \eta) d_A d_C)$ . We can now rewrite  $\tilde{Q}$  (obtained by applying Algorithm 1 to the distribution  $Q$ ) using  $T_{ABC}$ ,

$$\tilde{P}_{AC}^b = (1 - \eta)^2 P_{AC}^b + \eta(1 - \eta) \left( \frac{P_C^b}{d_A} + \frac{P_A^b}{d_C} \right) + \frac{\eta^2}{d_A d_C} = (1 - \eta)^2 P_{AC}^b + \eta(2 - \eta) T_{AC|B=b}, \quad (36)$$

$$\tilde{P}_A^b \tilde{P}_C^b = \left( (1 - \eta) P_A^b + \frac{\eta}{d_A} \right) \left( (1 - \eta) P_C^b + \frac{\eta}{d_C} \right) = (1 - \eta)^2 P_A^b P_C^b + \eta(2 - \eta) T_{AC|B=b}. \quad (37)$$

Introducing  $\nu = 2\eta - \eta^2$  (note that  $(1 - \eta)^2 + \eta(2 - \eta) = 1$ ), we can now rewrite

$$D(\tilde{P}_{AC}^b \| \tilde{P}_A^b \tilde{P}_C^b) = D\left( (1 - \nu) P_{AC}^b + \nu T_{AC|B=b} \middle\| (1 - \nu) P_A^b P_C^b + \nu T_{AC|B=b} \right). \quad (38)$$

We choose  $\nu = \tilde{O}(\varepsilon^2 / (d_A d_C)^2)$ , and apply Lemma 4.1, implying

$$\forall b \in B_+ : D(P_{AC}^b \| P_A^b P_C^b) = \varepsilon_b \implies D(\tilde{P}_{AC}^b \| \tilde{P}_A^b \tilde{P}_C^b) \geq \frac{\varepsilon_b}{2}, \quad (39)$$

which, together with (34), shows that

$$I(A:C|B)_P \geq \varepsilon \implies I(A:C|B)_{\tilde{P}} = D(\tilde{P}_{ABC} \| \tilde{P}_{AB} \tilde{P}_{C|B}) \geq \frac{\varepsilon}{4}. \quad (40)$$

The decision gap is reduced by only a constant factor, while we obtained a guarantee on the minimum probability mass of  $P_{A|B} P_{C|B}$ , since

$$\forall a, b, c : \tilde{P}_A^b(a) \tilde{P}_C^b(c) \geq \eta(2 - \eta) T_{AC|B=b}(a, c) \geq \frac{\eta^2}{d_A d_C}, \quad (41)$$

We now want to use Lemma 3.1 to move from KL-divergence to  $D_H^2$  distance.

$$I(A:C|B)_{\tilde{P}} = \sum_b p_b D(\tilde{P}_{AC}^b \| \tilde{P}_A^b \tilde{P}_C^b) \quad (42)$$

$$\leq \sum_b p_b \left( 2 + \log \left( \frac{d_A d_C}{\eta^2} \right) \right) D_H^2(\tilde{P}_{AC}^b, \tilde{P}_A^b \tilde{P}_C^b) \quad (43)$$

$$\leq 2 \log \left( \frac{d_A d_C}{\eta^2} \right) D_H^2(\tilde{P}_{ABC}, \tilde{P}_{AB} \tilde{P}_{C|B}), \quad (44)$$

which completes our reduction, since

$$I(A:C|B)_P \geq \varepsilon \implies D_H^2(P_{ABC}, P_{AB} P_{C|B}) \geq \frac{\varepsilon}{8 \log \left( \frac{d_A d_C}{\eta^2} \right)}. \quad (45)$$

□

## 5 Independence Testing in $D_H^2$

In this section, we study the problem of independence testing of a distribution  $P_{AC}$ , with respect to the squared Hellinger distance. We will first study a related problem of testing the equivalence of distributions when one distribution is promised to be a product distribution. Our problem is formally defined as follows.

**Problem 8** (Sample complexity of equivalence testing for product distribution). Fix a threshold  $\varepsilon$ , distance measure  $\Delta$  and alphabet sizes  $d_A$  and  $d_C$ . Consider a tester that, given access to  $N$  samples from two unknown distributions  $P_{AC}$  and  $Q_{AC}$  each with  $Q_{AC}$  being a product distribution, distinguishes between the classes

- (i)  $P_{AC} = Q_A Q_C$  and
- (ii)  $\Delta(P_{AC}, Q_A Q_C) \geq \varepsilon$ .

We denote the sample complexity of this problem by  $\text{SC}_{\text{EqProd}}(\Delta, \varepsilon, d_A, d_C)$ .

We first prove that this problem can be solved efficiently. Our result is described below.

**Theorem 5.1.** *Consider equivalence testing with product structure (Problem 8). Assuming w.l.o.g. that  $d_A \geq d_C$ , we have*

$$\text{SC}_{\text{EqProd}}(D_H^2, \varepsilon, d_A, d_C) = \tilde{O} \left( \min \left\{ \frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \right). \quad (46)$$

Note in particular that two samples from  $P_{AC}$  can be used to simulate a sample from  $P_A P_C$ , by taking  $(a_1, c_1), (a_2, c_2)$  from  $P_{AC}$  and returning  $(a_1, c_2)$  as a sample from  $P_A P_C$ .

As a corollary to the above theorem, we have the following results, where the formalized upper bounds of Result 1 come from combining Theorem 5.1 with Theorem 4.2 in Section 4.

**Corollary 5.2.** *Consider independence testing in  $D_H^2$  as defined in Problem 7. Assuming w.l.o.g. that  $d_A \geq d_C$ , we have*

$$\text{SC}_{\text{I,H}}(\varepsilon, d_A, d_C) = O \left( \text{SC}_{\text{EqProd}}(D_H^2, \varepsilon, d_A, d_C) \right). \quad (47)$$

Consider MI testing as defined in Problem 2. Assuming w.l.o.g. that  $d_A \geq d_C$ , we have

$$\text{SC}_{\text{MI}}(\varepsilon, d_A, d_C) = \tilde{O} \left( \min \left\{ \frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \right). \quad (48)$$

It is clear that an upper bound for Problem 8 implies an upper bound for Problem 7, and vice versa for the lower bound. We derive upper bounds for the former, and lower bounds for the latter to show equivalence of the sample complexity up to logarithmic factors. For the connection to the mutual information, we note that for any distributions  $P$  and  $Q$ ,  $D_H^2(P, Q) \leq D(P||Q)$  such that lower bounds for testing with respect to the squared Hellinger distance carry over to mutual information testing as well (see Corollary 6.2).

## 5.1 Some Preliminary Results

We first collect a few results which we will need along the way. Let us start with the result of [CDVV14] about equivalence testing between distributions.

**Lemma 5.3.** [CDVV14, Theorem 2] *Let  $P$  and  $Q$  be two unknown distributions defined over  $[d]$  such that  $\|P\|_2, \|Q\|_2 \leq b$  for some  $b$ . Then in order to distinguish between  $\|P - Q\|_2 \leq \varepsilon$  and  $\|P - Q\|_2 \geq 2\varepsilon$  with probability at least  $2/3$ ,  $\Theta(b/\varepsilon^2)$  samples are necessary and sufficient.*

Now we state a result from [BFF<sup>+</sup>01] regarding estimating the  $\ell_2$  norm of a distribution.

**Lemma 5.4.** [BFF<sup>+</sup>01, Theorem 12] *Let  $P$  be an unknown distribution defined over  $[d]$ . Given  $\eta \in (0, 1)$ , in order to estimate  $\|P\|_2^2$  up to a multiplicative precision of  $1 \pm \eta$  with probability at least  $2/3$ ,  $O(\sqrt{d}/\eta^2)$  i.i.d. samples from  $P$  are sufficient.*

We need to generalize both the results stated above, since we will need analogous statements about subsets of distributions. Note that variants of both our Lemma 5.5 and Lemma 5.6 results already appeared implicitly in the proofs of [DK16].

**Lemma 5.5.** *Let  $P$  be an unknown distribution over  $[d]$ ,  $S \subseteq [d]$  and  $\varepsilon \in (0, 1)$  be a threshold. Given i.i.d. sample access to  $P$ , in order to determine  $c$  such that  $\frac{\|P^S\|_2}{2} \leq c \leq 2(\|P^S\|_2 + \varepsilon)$  holds with probability at least  $2/3$ ,  $O(\sqrt{d} + \frac{1}{\varepsilon})$  samples from  $P$  are sufficient.*

*Proof.* Following Lemma 5.4, we know that estimating  $\|P\|_2$  up to a multiplicative factor 2 requires  $O(\sqrt{d})$  samples from  $P$ . However, in our case, we are only interested in a subset of the support of  $P$ , and samples that are not coming from  $S$  might not be directly useful to determine the weight of  $P^S$ .

For this reason we create a dummy distribution  $P'$  using  $P$ , such that  $\|P^S\|_2 \approx \|P'\|_2$ , up to an additive constant  $\varepsilon$ . In order to define  $P'$ , we create a set of indices  $T = S \cup N$ , with  $|N| = O(1/\varepsilon^2)$  new indices disjoint from  $[d]$ .

The main idea is now to sample from  $P'$  via a flattening technique. Let us take a sample  $s$  from  $P$ . If  $s \in S$ , we keep it as it is. Otherwise, if  $s \notin S$ , we assign it one of the  $N$  new indices, uniformly at random. Our new distribution  $P'$  satisfies

$$\|P\|_2 \leq \|P'\|_2 \leq \|P^S\|_2 + \sqrt{\sum_{i \notin S} \frac{1}{N^2}} \leq \|P^S\|_2 + \frac{1}{\sqrt{N}} = \|P^S\|_2 + O(\varepsilon), \quad (49)$$

where the last equality follows from the fact that  $|N| = O(1/\varepsilon^2)$ .

Note that the distribution  $P'$  is defined on  $S \cup N$ . Since  $S \subseteq [d]$ , the support size of  $P'$  is  $O(d + 1/\varepsilon^2)$ . Using Lemma 5.4, we can determine  $\|P'\|_2$  up to a small multiplicative constant factor, say  $1/2$ , using  $O(\sqrt{d + 1/\varepsilon^2}) = O(\sqrt{d} + 1/\varepsilon)$  samples from  $P$ , which directly implies our result.  $\square$

Now let us prove an analogous statement to Lemma 5.3 for subsets of distributions. We will first use a trick to reduce the constraint  $\max\{\|P\|_2, \|Q\|_2\} \leq b$  to  $\|Q\|_2 \leq b$ , which was used in the proof of [DK16, Lemma 2.3].

**Lemma 5.6.** *Let  $P$  and  $Q$  be two unknown distributions over  $[d]$ , and  $S \subseteq [d]$ . Moreover, let us assume that  $\|Q^S\|_2 \leq b$  for some  $b \in \mathbb{R}$ . Given i.i.d. sample access to  $P$  and  $Q$ , in order to distinguish if  $P^S = Q^S$  or  $\|P^S - Q^S\|_2 \geq \varepsilon$  with probability at least  $2/3$ ,  $O(\max\{\frac{b}{\varepsilon^2}, \frac{1}{\varepsilon}, \sqrt{d}\})$  samples are sufficient.*

*Proof.* We will first use Lemma 5.5 to obtain a  $c$  such that  $\|P^S\|_2/2 \leq c \leq 2\|P^S\|_2 + \max\{\varepsilon, b\}$  holds, which takes  $O(\sqrt{d} + \min\{1/\varepsilon, 1/b\})$  samples. We return ‘Far’ in case  $[c - (b + \varepsilon)]/2 > b$ , i.e., if our lower bound on  $\|P^S\|_2$  is larger than  $b$  (our upper bound on  $\|Q^S\|_2$ ). If we did not reject, then  $[c - (b + \varepsilon)]/2 \leq b$ , which implies that  $c \leq 3b + \varepsilon$ . This guarantees that  $\max\{\|P^S\|_2, \|Q^S\|_2\} \leq 6b + 2\varepsilon$ .

We now perform equivalence testing between  $P^S$  and  $Q^S$ . For this, we will create two dummy distributions  $P'$  and  $Q'$ , consisting of the indices  $S$  and a set  $N$  of new indices, for  $|N| = \Omega(1/\varepsilon^2)$ . Let us obtain a sample  $s$  from  $P$ . If  $s \in S$ , we keep it as it is. Otherwise, if  $s \notin S$ , we will assign it one of the  $N$  new indices, uniformly at random. We perform the analogous operation for  $Q$  as well to obtain samples from  $Q'$ . Note that

$$\|P' - Q'\|_2^2 = \|P^S - Q^S\|_2^2 + N \left( \frac{1 - P[S]}{N} - \frac{1 - Q[S]}{N} \right)^2 \quad (50)$$

$$= \|P^S - Q^S\|_2^2 + \frac{(P[S] - Q[S])^2}{N}. \quad (51)$$

If we choose  $N \geq 4/\varepsilon^2$ , then

$$P^S = Q^S \implies \|P' - Q'\|_2 \leq \varepsilon/2, \quad \text{and} \quad \|P^S - Q^S\|_2 \geq \varepsilon \implies \|P' - Q'\|_2 \geq \varepsilon. \quad (52)$$

We can thus use robust equivalence testing between  $P'$  and  $Q'$  (Lemma 5.3) to test for equivalence between  $P^S$  and  $Q^S$ . The precision to which we test is  $\varepsilon/2$ , and it remains to upper bound  $\max\{\|P'\|_2, \|Q'\|_2\}$ . By construction,  $\|P'\|_2 \leq \|P^S\|_2 + O(\varepsilon)$ , and analogous for  $Q'$  (see (49)). Since we showed  $\max\{\|P^S\|_2, \|Q^S\|_2\} \leq 6 \max\{b, \varepsilon\}$ , this results in a sample complexity of  $O(\max\{b, \varepsilon\}/\varepsilon^2)$  using Lemma 5.3. Recall that, at the beginning, we use  $O(\sqrt{d} + \min\{1/\varepsilon, 1/b\})$  samples to obtain a  $c$  such that  $\|P^S\|_2/2 \leq c \leq 2\|P^S\|_2 + \max\{\varepsilon, b\}$  holds. Combining the above, we have the result.  $\square$

In order to boost the success probability to  $1 - \delta$  for any  $\delta \in (0, 1)$ , we run the above algorithm  $O(\log(1/\delta))$  times and return the majority answer. The analysis follows from a standard success probability amplification argument. This is formalized in the following corollary.

**Corollary 5.7.** *Given a  $\delta \in (0, 1)$ , two multisets  $\mathcal{S}_P, \mathcal{S}_Q$  of i.i.d. samples from  $P$  and  $Q$ , and  $S, b$  and  $\varepsilon$  as in Lemma 5.6. There exists an algorithm `EquivalenceL2`( $\mathcal{S}_P, \mathcal{S}_Q, S, b, \varepsilon, \delta$ ) which performs equivalence testing between  $Q^S$  and  $P^S$  in  $\ell_2$  distance, which succeeds with probability at least  $1 - \delta$  if*

- $|\mathcal{S}_P|, |\mathcal{S}_Q| \geq c_{eq} \max\{\frac{b}{\varepsilon^2}, \frac{1}{\varepsilon}, \sqrt{d}\} \log(1/\delta)$  in case  $|S| < |D|$ ,
- $|\mathcal{S}_P|, |\mathcal{S}_Q| \geq c_{eq} \frac{b}{\varepsilon^2}$  if  $S = D$ ,

where  $c_{eq}$  is an instance independent constant.

We will repeatedly use simple concentration bounds based on the Chernoff inequalities, which we prove here for completeness.

**Lemma 5.8.** *Given a threshold  $\tau$ , a parameter  $\zeta \in (0, 1)$  and  $N \geq 8 \log(4d/\zeta)/\tau$  i.i.d. samples from an unknown distribution  $P$  of dimension  $d$ , the empirical estimator  $\hat{P}$  satisfies the following properties with probability at least  $1 - \zeta$ ,*

- (i)  $\forall i : p_i \geq \tau \implies 1/2 \leq \hat{p}_i/p_i \leq 2$ ,
- (ii)  $\forall i : p_i \leq \tau \implies \hat{p}_i \leq 2\tau$ .

where  $p_i$  and  $\hat{p}_i$  denote the probability mass of the element  $i \in [d]$  in the distributions  $P$  and  $\hat{P}$ , respectively.

*Proof.* This is a direct application of the Chernoff and union bounds. First, we assume  $p_b \leq \tau$ . We set  $\delta$  such that  $(1 + \delta)p_b = 2\tau$ , and denote by  $X$  how often  $b$  was observed among the  $N$  samples. Since  $\delta \geq 1$ , we have  $\delta^2 \mathbb{E}[X] \geq (1 + \delta)\mathbb{E}[X]/2$  such that

$$\Pr[X \geq 2\tau] \leq e^{-\delta^2 \mu/3} \leq e^{-\tau N/3} \leq \frac{\zeta}{4d}. \quad (53)$$

For the other direction, where  $p_b \geq \tau$ , we set  $\delta = 1/2$  and  $\delta = 2$ , respectively,

$$\Pr[X \leq \mathbb{E}[X]/2] \leq e^{-\mathbb{E}[X]/8} \leq e^{-\tau N/8} \leq \frac{\zeta}{4d}, \quad \text{and} \quad \Pr[X \geq 2\mathbb{E}[X]] \leq e^{-\mathbb{E}[X]/3} \leq \frac{\zeta}{4d}. \quad (54)$$

Finally, we apply a union bound over all cases.  $\square$

## 5.2 Our Independence Tester

We now have all the preliminaries in place to prove the main result of this section:

**Theorem 5.1.** *Consider equivalence testing with product structure (Problem 8). Assuming w.l.o.g. that  $d_A \geq d_C$ , we have*

$$\text{SC}_{\text{EqProd}}(D_H^2, \varepsilon, d_A, d_C) = \tilde{O} \left( \min \left\{ \frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \right). \quad (46)$$

The proof of correctness follows from the discussion of Algorithm 2 (Lemma 5.12) and the associated subroutines. Our algorithm takes two multisets  $\mathcal{S}_P$  and  $\mathcal{S}_Q$  of  $\text{SC}_{\text{EqProd}}(D_H^2, \varepsilon, d_A, d_C)$  samples from unknown distributions  $P_{AC}$  and  $Q_{AC}$ , where  $Q_{AC} = Q_A Q_C$  is guaranteed to have product structure. With high probability, our algorithm outputs ‘Yes’ if  $P_{AC} = Q_{AC}$  and ‘No’ if  $D_H^2(P_{AC}, Q_{AC}) \geq \varepsilon$ . As described in the overview of our methods in Section 2.3, the underlying idea is to partition  $A \times C$  into a polylogarithmic number of categories, denoted by  $S_{ij}$ , on which we then individually perform equivalence testing with respect to the  $\ell_2$  distance. The partitioning used in Algorithm 2 is visualized in Figure 4, where we also define the categories  $S_{ij}$ .




Fig. 4: Partition of  $d_A \times d_C$  based on  $\hat{Q}_A$  and  $\hat{Q}_C$ . Indices  $(a, c)$  of similar weight  $\hat{q}_a \hat{q}_c$  are grouped together in categories  $S_{ij}$ , which are used to perform piecewise equivalence testing with  $P_{AC}$ . The axes are labeled according to the corresponding category, which is inverse logarithmic to the weight of the probabilities. The color of the categories indicate a different analysis of the sample complexity of the categories. The red regimes dominate the sample complexity.

In the following, we define some quantities which will be used in Algorithm 2. Further, let  $k_{AC} := (\lceil \log(d_A d_C / \varepsilon) \rceil + 1)^2$ .

$$\gamma(i, j) := \begin{cases} \sqrt{\frac{\varepsilon e^{-(i+j)+2}}{e^6 k_{AC}}} & \text{if } i < k_A \text{ and } j < k_C, \\ \frac{\varepsilon}{4k_{AC} \sqrt{|S_{ij}|}} & \text{otherwise,} \end{cases} \quad (55)$$

$$b(i, j) := \min \left\{ \sqrt{|S_{ij}|} e^{-(i+j)+2}, \sqrt{e^{-(i+j)+2}} \right\}. \quad (56)$$

Algorithm 2 will use

$$N := 10^2 k_{AC} \log(k_{AC}) \max\{N_{\text{heavy}}, N_{\text{mixed}}\} \quad (57)$$

samples, where

$$N_{\text{heavy}} := 10^5 k_{AC} \frac{\sqrt{d_A d_C}}{\varepsilon}, \quad (\text{Claim 5.10}), \quad (58)$$

$$N_{\text{mixed}} := 10k_{AC}^2 \min \left\{ \frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\}, \quad (\text{Claim 5.11}), \quad (59)$$

which gives the sample complexity stated in Theorem 5.1.

Finally, we define  $k_A, k_C := \lceil \log(N_{\text{mixed}}) \rceil$ . Note that  $k_{AC} \geq (k_A + 1)(k_C + 1)$ .

---

**Algorithm 2:** Equivalence Product Testing in  $D_H^2$

---

**Input:** Multisets  $\mathcal{S}_P$  and  $\mathcal{S}_Q$  of size  $N$  of tuples in  $A \times C$  each,  $\varepsilon \in (0, 1) \triangleright N$ , see (57)

**Output:** ‘Yes’ or ‘No’

▷ Step 1: Partition  $A \times C$

1  $M \leftarrow N_{\text{mixed}}, \delta \leftarrow \frac{1}{10^3 k_{AC}}$  ▷  $N_{\text{mixed}}$ , see (59))

2  $\mathcal{S}_A, \mathcal{S}_C \leftarrow$  remove random multisets of size  $8M \log(10^3 d_A d_C)$  each from  $\mathcal{S}_Q$

3  $\hat{Q}_A \leftarrow$  empirical frequencies of  $a \in A$  in  $\mathcal{S}_A$

4  $\hat{Q}_C \leftarrow$  empirical frequencies of  $c \in C$  in  $\mathcal{S}_C$

5  $\forall i \in [k_A]_0, j \in [k_C]_0$ : define  $S_{ij}$  according to Figure 4 ▷ partition  $A \times C$

▷ Step 2: Equivalence Testing

6 for all  $i, j$  do

7  $\mathcal{S}_p, \mathcal{S}_q \leftarrow$  remove random multisets of size  $M$  each from  $\mathcal{S}_P$  and  $\mathcal{S}_Q$ , respectively

8 if EquivalenceL2( $\mathcal{S}_p, \mathcal{S}_q, S_{ij}, b(i, j), \gamma(i, j), \delta$ ) = ‘Far’ then

9 | return ‘No’. ▷ Corollary 5.7, distinction for heavy/mixed implicit in  $\gamma$

10 return ‘Yes’.

---

**Remark 5.9.** We note that for  $i < k_A, j < k_C$ , we can group together all groups for which  $i + j$  gives the same value. It is easy to see that this would reduce the number of equivalence tests from  $O(\log^2(k_A))$  to  $O(\log(k_A))$ . An analogous strategy can be applied later in Section 7.2 as well.

We will first approximately learn  $P_A$  and  $P_C$ . This allows us to group indices  $a \in A$  into  $k_A + 1$  many bins  $A_0, A_1, \dots, A_k$ , such that the weights of indices in the same bin are close to each other up to a constant factor. The only exception are the indices with particularly small weights, which we group together in the last bin,  $A_{k_A}$ . We perform the same approach for  $C$  as well. For a fixed bin  $A_i$  and a fixed bin  $C_j$ , we can define  $S_{ij} := \{(a, c) | a \in A_i, c \in C_j\}$ , which we call a *category*. If  $i < k_A$  and  $j < k_C$ , we call  $S_{ij}$  a *heavy category*. For any  $(a, c), (a', c')$  from the same heavy category  $S_{ij}$ , the corresponding weights of our reference distribution,  $p_a p_c$  and  $p_{a'} p_{c'}$ , will be close as well. The



remaining categories are called *light*, as the weights can be arbitrarily small. A guarantee analogous to heavy categories is not possible here. After partitioning the distribution in a two-dimensional grid, we perform equivalence testing for each of the  $k_{AC}$  bins using Lemma 5.6. In the process, we will need to bound terms of the form  $\|P_{AC}^{S_{ij}}\|_2$ . There are two ways to do so (see Fact 3.3), and which option we choose depends on the parameters  $d_A, d_C$ , and  $\varepsilon$  of our problem. The final sample complexity will reflect this choice in terms of the minimization.

The sample complexity of testing a category  $S_{ij}$  (or, equivalently, the precision to which we have to test) depends on the type of category.

1. **Heavy categories:** This includes all categories where neither  $a$  nor  $c$  come from the last bin, such that for all  $a, c$ , both  $p_a$  and  $p_c$  are bounded away from zero. That is we will consider categories  $S_{ij}$  where  $i < k_A$  and  $j < k_C$ . Our result in this setting is as follows.

**Claim 5.10.** *For any heavy category  $S_{ij}$  as defined in Figure 4, distinguishing whether  $P_{ABC}^{S_{ij}} = Q_{ABC}^{S_{ij}}$  or  $D_H^2(P_{ABC}^{S_{ij}}, Q_{ABC}^{S_{ij}}) \geq \varepsilon/k_{AC}$  can be done using*

$$N_{\text{heavy}} \geq 10^5 k_{AC} \frac{\sqrt{d_A d_C}}{\varepsilon} \quad (60)$$

*samples of  $P_{AC}$  and the product distribution  $Q_A Q_C$  each.*

2. **Mixed categories:** This includes all categories where at least one of  $a$  or  $c$  come from the last bin, that is, categories  $S_{ij}$  where  $i = k_A$  or  $j = k_C$ , or both.

**Claim 5.11.** *For any mixed category  $S_{ij}$  as defined in Figure 4, distinguishing whether  $P_{ABC}^{S_{ij}} = Q_{ABC}^{S_{ij}}$  or  $D_H^2(P_{ABC}^{S_{ij}}, Q_{ABC}^{S_{ij}}) \geq \varepsilon/k_{AC}$  can be done using*

$$N_{\text{mixed}} \geq 10k_{AC}^2 \min \left\{ \frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \quad (61)$$

*samples of  $P_{AC}$  and the product distribution  $Q_A Q_C$  each.*

These two claims will be proven in the following subsections, Section 5.2.1 and Section 5.2.2, respectively. Assuming these two claims hold, we can prove the correctness of Theorem 5.1, as follows.

**Lemma 5.12.** *Algorithm 2 performs equivalence testing for a product distribution with probability of success at least  $2/3$  when using the amount of samples specified in Theorem 5.1.*

*Proof.* Note that the first step in the algorithm produces a binning such that with probability at least  $9/10$ ,

$$\forall i \in [k_A]_0, \forall a \in A_i : e^{-(i+2)} \mathbb{1}[i < k_A] \leq q_a \leq e^{-i+1}, \quad (62)$$

$$\forall j \in [k_C]_0, \forall c \in C_j : e^{-(j+2)} \mathbb{1}[j < k_C] \leq q_c \leq e^{-j+1}. \quad (63)$$

The guarantee follows directly from Lemma 5.8 and a union bound, analogously for both  $A$  and  $C$ : since we learn all  $q_a \geq 1/N_{\text{mixed}}$  up to a factor two, the smallest possible  $q_a$  for which the estimate  $\hat{q}_a$  satisfies  $\hat{q}_a \geq e^{-i+1}$  is  $e^{-(i+1)}/2$ . This implies that no  $a$  with  $q_a < e^{-i+1}/2$  can be assigned to  $A_i$ . Analogously, the largest  $p_a$  with  $a \in A_i$  has to satisfy  $p_a \leq 2e^{-i+1}$ . In particular, it holds with high probability that

$$\forall i \in [k_A - 1]_0 : \max_{a, a' \in A_i} \frac{p_a}{p_{a'}} \leq e^3, \quad \forall j \in [k_C - 1]_0 : \max_{c, c' \in C_j} \frac{p_c}{p_{c'}} \leq e^3. \quad (64)$$

We will use these inequalities in the following subsections. Note that Lemma 5.8 also upper bounds  $\hat{q}_a$  and  $\hat{q}_c$  for small values from  $A_0$  and  $C_0$ .

Let us now consider the squared Hellinger distance between  $P_{AC}$  and  $Q_{AC}$ .

$$D_H^2(P_{AC}, Q_{AC}) = \sum_{S_{ij} \in \mathcal{S}} \sum_{(a,c) \in S_{ij}} (\sqrt{p_{ac}} - \sqrt{q_{ac}})^2 = \sum_{i,j} D_H^2(P_{AC}^{S_{ij}}, Q_{AC}^{S_{ij}}) =: \sum_{i,j} \eta_{ij}. \quad (65)$$

The above equation implies that if  $P_{AC} = Q_{AC}$ , then  $P_{AC}^{S_{ij}} = Q_{AC}^{S_{ij}}$  will hold for every  $S_{ij} \in \mathcal{S}$ . On the other hand, if  $D_H^2(P_{AC}, Q_{AC}) \geq \varepsilon$ , then by the pigeonhole principle, at least one of the  $|\mathcal{S}| = k_A k_C \leq k_{AC}$  categories, say  $S^{i'j'}$ , needs to satisfy

$$\eta_{i'j'} = D_H^2(P_{AC}^{S^{i'j'}}, (Q_{AC})^{S^{i'j'}}) \geq \frac{\varepsilon}{k_{AC}}. \quad (66)$$

We will test this individually for each category. If a category has a value  $\eta_{ij}$  such that  $0 < \eta_{ij} < \varepsilon/k_{AC}$ , then the algorithm might either indicate that  $\eta_{ij} = 0$  or  $\eta_{ij} \geq \varepsilon/k_{AC}$ . If it indicates  $\geq \varepsilon/k_{AC}$ , then this is fine, since we ultimately only care whether  $\sum_{i,j} \eta_{ij} = 0$  or not. In both cases, the output can not cause a wrong answer overall, since we know that the output for  $S^{i'j'}$  will be correct with high probability.

The choice of  $\delta$  guarantees that we succeed with probability 9/10 over all categories, and a union bound guarantees a probability of success of at least 2/3 for the entire algorithm.

To show the sample complexity, we note that

- Step 1 takes  $16M \log(10^3 d_A d_C)$  samples,
- Step 2 takes at most  $k_{AC} \cdot 8 \log(10^4 k_{AC}^2) \max\{N_{\text{heavy}}, N_{\text{mixed}}\}$  samples.

This directly implies the sample complexity of Theorem 5.1. □

In the remainder of this section, we prove Claim 5.10 and Claim 5.11.

### 5.2.1 Sample Complexity of the Heavy Regime

**Claim 5.10.** *For any heavy category  $S_{ij}$  as defined in Figure 4, distinguishing whether  $P_{ABC}^{S_{ij}} = Q_{ABC}^{S_{ij}}$  or  $D_H^2(P_{ABC}^{S_{ij}}, Q_{ABC}^{S_{ij}}) \geq \varepsilon/k_{AC}$  can be done using*

$$N_{\text{heavy}} \geq 10^5 k_{AC} \frac{\sqrt{d_A d_C}}{\varepsilon} \quad (60)$$

*samples of  $P_{AC}$  and the product distribution  $Q_A Q_C$  each.*

*Proof.* Fix an arbitrary heavy category, say  $S_{ij}$ . Recall that the heavy regime includes all categories where neither  $a$  nor  $c$  come from the last bin  $A_{k_A}$  and  $C_{k_C}$ , such that both  $P_A[a]$  and  $P_C[c]$  are bounded away from zero. For better readability, we will simply write  $p_a$  for  $P_A[a]$  and  $p_c$  for  $P_C[c]$ , as well as  $T := S_{ij}$  (since we fixed  $i, j$ ), and analogously  $T_A := (S_{ij})_A$  and  $T_C := (S_{ij})_C$ . Since  $Q_{AC}$  is a product distribution, we will write  $q_{ac}$  and  $q_a q_c$  interchangeably.

We will bound  $D_H^2(P_{AC}^T, Q_{AC}^T)$  as follows:

$$D_H^2(P_{AC}^T, Q_{AC}^T) = \sum_{(a,c) \in T} (\sqrt{p_{ac}} - \sqrt{q_{ac}})^2 \quad (67)$$

$$= \sum_{(a,c) \in T} \frac{(p_{ac} - q_{ac})^2}{(\sqrt{p_{ac}} + \sqrt{q_{ac}})^2} \quad (68)$$

$$\leq \sum_{(a,c) \in T} \frac{(p_{ac} - q_{ac})^2}{q_{ac}} \quad (69)$$

$$\leq \sum_{(a,c) \in T} \frac{(p_{ac} - q_{ac})^2}{\min_{(a,c) \in T} \{q_a q_c\}} \quad (70)$$

$$\leq e^6 \frac{\|P_{AC}^T - Q_{AC}^T\|_2^2}{\max_{(a,c) \in T} \{q_a q_c\}}. \quad (71)$$

Note that the last line, which bounds  $\max_{(a,c),(a',c') \in T} q_a q_c / (q_{a'} q_{c'})$  follows from (64). This follows from learning  $q_a$  up to a factor 2, and the definition of the bins. An analogous argument holds for  $T_C$  as well.

Following (71), the problem of testing heavy categories is reduced to equivalence testing between  $P_{AC}^T = Q_{AC}^T$  and  $D_H^2(P_{AC}^T, Q_{AC}^T) \geq \varepsilon/k_{AC}$ . Moreover, from (71), we can say that the problem is reduced to the following testing problem (using bounds from Lemma 5.12):

$$P_{AC}^T = Q_{AC}^T \quad \text{or} \quad \|P_{AC}^T - Q_{AC}^T\|_2 \geq \sqrt{\frac{\varepsilon \max_{(a,c) \in T} \{q_a q_c\}}{e^6 k_{AC}}} \geq \sqrt{\frac{\varepsilon e^{-(i+j+2)}}{e^6 k_{AC}}} =: \eta. \quad (72)$$

From Lemma 5.6 and Corollary 5.7, we know that equivalence testing with respect to the  $\ell_2$  distance with parameter  $\eta$  requires  $c_{\text{eq}} \max\{\|(Q_A Q_C)^T\|_2 / \eta^2, 1/\eta, \sqrt{d_A d_C}\}$  samples. Note that

$$\|(Q_A Q_C)^T\|_2 \leq \sqrt{\sum_{(a,c) \in T} q_a^2 q_c^2} \leq \sqrt{|T_A| |T_C|} \max_{(a,c) \in T} \{q_a q_c\} \leq \sqrt{|T_A| |T_C|} e^{-(i+j)+2}. \quad (73)$$

Hence, the sample complexity of the testing problem in (72) is  $e^{10} k_{AC} \sqrt{|T_A| |T_C|} / \varepsilon$ . Since  $|T_A| \leq d_A$  and  $|T_C| \leq d_C$ , the problem of testing a heavy category requires  $e^{10} k_{AC} \sqrt{d_A d_C} / \varepsilon$  samples. This completes the proof of Claim 5.10.  $\square$

### 5.2.2 Sample Complexity of the Mixed Regime

**Claim 5.11.** *For any mixed category  $S_{ij}$  as defined in Figure 4, distinguishing whether  $P_{ABC}^{S_{ij}} = Q_{ABC}^{S_{ij}}$  or  $D_H^2(P_{ABC}^{S_{ij}}, Q_{ABC}^{S_{ij}}) \geq \varepsilon/k_{AC}$  can be done using*

$$N_{\text{mixed}} \geq 10k_{AC}^2 \min \left\{ \frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \quad (61)$$

*samples of  $P_{AC}$  and the product distribution  $Q_A Q_C$  each.*

*Proof.* Our proof follows in similar line as Claim 5.10. We would like to upper bound  $D_H^2(P_{AC}^T, Q_{AC}^T)$  by  $\|P_{AC}^T - Q_{AC}^T\|_2$ . However, we can no longer use the same approach as (71) as there is no lower

bound on  $q_i$ . Instead, we denote by  $T_-$  the set of indices in  $T$  for which  $(\sqrt{p_{ac}} - \sqrt{q_{ac}})^2 \leq x$ , for some  $x \in (0, 1)$  to be determined later and consider  $T_+ := T \setminus T_-$ . Then we have

$$D_H^2(P_{AC}^T, Q_{AC}^T) = \sum_{(a,c) \in T_+} (\sqrt{p_{ac}} - \sqrt{q_{ac}})^2 + \sum_{(a,c) \in T_-} (\sqrt{p_{ac}} - \sqrt{q_{ac}})^2 \quad (74)$$

$$\leq \sum_{(a,c) \in T_+} \frac{(p_{ac} - q_{ac})^2}{(\sqrt{p_{ac}} + \sqrt{q_{ac}})^2} + |T_-| \cdot x \quad (75)$$

$$\leq \frac{1}{x} \|P_{AC}^{T_+} - Q_{AC}^{T_+}\|_2^2 + |T_-| \cdot x \quad (76)$$

$$\leq \frac{1}{x} \|P_{AC}^T - Q_{AC}^T\|_2^2 + |T_-| \cdot x, \quad (77)$$

where the second inequality holds since  $(\sqrt{p_{ac}} + \sqrt{q_{ac}})^2 \geq x$ . Thus, we can say that

$$D_H^2(P_{AC}^T, Q_{AC}^T) \geq \frac{\varepsilon}{k_{AC}} \implies \|P_{AC}^T - Q_{AC}^T\|_2 \geq \sqrt{\frac{x\varepsilon}{k_{AC}} - |T_-|x^2} \quad (78)$$

We set  $x = \varepsilon/(2|T_-|k_{AC})$ , and bound  $|T_-| \leq |T_A||T_C|$ . As a result, following (78), our goal is reduced to the following decision problem

$$P_{AC}^T = Q_{AC}^T \quad \text{or} \quad \|P_{AC}^T - Q_{AC}^T\|_2 \geq \frac{\varepsilon}{4k_{AC}\sqrt{|T_A||T_C|}} =: \eta. \quad (79)$$

Similar to the proof of the previous claim, we now need to bound  $\|Q_{AC}^T\|_2$ . Here we will use Fact 3.3, which allows us to bound  $\|Q_{AC}^T\|_2$  in the following two ways

$$\|Q_{AC}^T\|_2 \leq \begin{cases} \sqrt{|T_A||T_C| \left( \max_{(a,c) \in T} \{q_a q_c\} \right)^2} & \text{(Case 1)} \\ \sqrt{\max_{(a,c) \in T} \{q_a q_c\}} & \text{(Case 2)} \end{cases} \quad (80)$$

From Lemma 5.6, we know that equivalence testing with respect to the  $\ell_2$  distance with parameter  $\eta$  requires  $c_{\text{eq}}\|Q_{AC}^T\|_2/\eta^2$  samples (will we see that this term dominates the other terms  $1/\eta$  and  $\sqrt{|T|}$ ). Thus, to decide (79),  $c_{\text{eq}}\|Q_{AC}^T\|_2/\eta^2$  samples are sufficient.

Assume without loss of generality that the small probability mass comes from  $A$ , such that  $q_a \leq 1/M$ . Since  $\forall c \in T_C : |T_C|p_c \leq 1$  (independent of whether  $p_c$  is large or small), both options to bound  $\|Q_{AC}^T\|_2$  presented in (80) might be relevant, depending on the relations between the other parameters:

1. **Case 1:** Using (79) and (80), we have

$$c_{\text{eq}} \frac{\|Q_{AC}^T\|_2}{\eta^2} \leq 4^2 c_{\text{eq}} \frac{k_{AC}^2 (|T_A||T_C|)^{3/2}}{\varepsilon^2} \max_{(a,c) \in T} \{q_a q_c\} \quad (81)$$

$$\leq 4^2 c_{\text{eq}} \frac{k_{AC}^2 |T_A|^{3/2} |T_C|^{1/2}}{\varepsilon^2 M} \leq N_{\text{mixed}}. \quad (82)$$

Since  $N_L$  will be maximal if  $M = N_L$ , and since  $|T_A| \leq d_A$  and  $|T_C| \leq d_C$ , such that we find

$$4^2 c_{\text{eq}} \frac{k_{AC}^2 d_A^{3/2} d_C^{1/2}}{\varepsilon^2 N_{\text{mixed}}} \leq N_{\text{mixed}} \implies N_{\text{mixed}} \geq 4 \sqrt{c_{\text{eq}}} \frac{k_{AC} d_A^{3/4} d_C^{1/4}}{\varepsilon}, \quad (83)$$

2. **Case 2:** We use (79) and (80) to obtain:

$$c_{\text{eq}} \frac{\|Q_{AC}^T\|_2}{\eta^2} \leq 4^2 c_{\text{eq}} \frac{k_{AC}^2 |T_A| |T_C|}{\varepsilon^2} \sqrt{\max_{(a,c) \in T} \{q_a q_c\}} \quad (84)$$

$$\leq 4^2 c_{\text{eq}} \frac{k_{AC}^2 |T_A| |T_C|^{1/2}}{\varepsilon^2 \sqrt{M}} \leq N_{\text{mixed}}. \quad (85)$$

Again, assuming  $M = N_{\text{mixed}}$  gives us

$$4^2 c_{\text{eq}} \frac{k_{AC}^2 d_A d_C^{1/2}}{\varepsilon^2 \sqrt{N_{\text{mixed}}}} \leq N_{\text{mixed}} \implies N_{\text{mixed}} \geq (4k_{AC})^{4/3} c_{\text{eq}}^{2/3} \frac{d_A^{2/3} d_C^{2/3}}{\varepsilon^{4/3}}. \quad (86)$$

This completes the proof of Claim 5.11.  $\square$

## 6 Lower Bounds for Independence Testing

In this section, we prove lower bounds on independence testing in the squared Hellinger distance. Our result is formally stated below.

**Theorem 6.1.** *Consider independence testing in the squared Hellinger distance  $D_H^2$  (Problem 7). Assuming w.l.o.g.  $d_A \geq d_C$ , we have*

$$\text{SC}_{I,H}(\varepsilon, d_A, d_C) = \tilde{\Omega} \left( \min \left\{ \frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \right). \quad (87)$$

Since  $D_H^2(P, Q) \leq D(P||Q)$  (Lemma 3.1), this also implies lower bounds on testing for mutual information.

**Corollary 6.2** (Formalized lower bounds of Result 1). *Consider mutual information testing, Problem 2, where w.l.o.g.  $d_A \geq d_C$ . We have*

$$\text{SC}_{\text{MI}}(\varepsilon, d_A, d_C) = \tilde{\Omega} \left( \min \left\{ \frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \right). \quad (88)$$

In order to prove the above theorem, we construct a pair of distributions that are hard to distinguish, unless we take a sufficiently large number of samples. Our approach follows the techniques used in the proof of [DK16, Theorem 3.1], where lower bounds for independence testing in variation distance are derived. The definition of our hard instances is based on the construction used in [DK16, Prop. 3.8] to show bounds on equivalence testing in  $D_H^2$  distance, with according modifications to accommodate the product structure we require for independence testing.

We will use the following lemma.

**Lemma 6.3.** [DK16, Lemma 3.2] *Let  $X$  be a uniformly random bit and  $K$  be a correlated random variable. Then if  $f$  is a function such that  $\Pr[f(K) = X] > 51\%$ , then  $I(X : K) \geq 2 \cdot 10^{-4}$ .*

In order to apply Lemma 6.3, we will construct a decision problem with two important properties: first, independence testing will be a valid strategy to solve the problem, and second, Lemma 6.3 allows us to derive lower bounds on the sample complexity of the task. This then will also imply lower bounds on independence testing.

We will construct two sets of distributions, one set containing product distributions, and the other set consisting of distributions  $P_{AC}$  which satisfy  $D_H^2(P_{AC}, P_AP_C) \geq \varepsilon$ . Depending on a fair random bit  $X$ , we select one of the two sets, from which we pick a random distribution  $P$ . We then draw  $n$  samples from  $P$ . Since the samples are i.i.d all information about the underlying distribution can be obtained from the frequencies at which we observe the different tuples  $(a, c)$  in the samples. We create a matrix  $K$ , such that  $K[a, c]$  indicates how often we observed the tuple  $(a, c)$ . The task is now to identify the random bit  $X$  based on  $K$ , and Lemma 6.3 allows us to derive a lower bound for this task:  $I(X : K)$  is increasing with  $n$ , so let us choose  $n^*$  such that  $I(X : K)$  reaches the threshold as stated in Lemma 6.3. We can then only determine the random bit marginally better than random. However, we would like to correctly reconstruct  $X$  with high probability, implying that more than  $n^*$  samples are necessary to distinguish between  $P_{AC} = P_AP_C$  and  $D_H^2(P_{AC}, P_AP_C) \geq \varepsilon$ , thereby proving the desired lower bound. As mentioned, a valid strategy to reconstruct  $X$  would be a general independence tester, such that a lower bound on this problem implies a lower bound for deciding whether  $P_{AC} = P_AP_C$  or  $D_H^2(P_{AC}, P_AP_C) \geq \varepsilon$ .

We will use the Poissonization technique (see Lemma 3.7), in particular, we will take  $\text{Poi}(n)$  samples from the distribution. With the Poissonization technique, we will not get exactly  $n$  samples, but with high probability, we will obtain  $\Theta(n)$  samples. Moreover, it is sufficient to have a (pseudo) distribution whose total probability mass is  $\Theta(1)$  with high probability, instead of exactly 1.

## 6.1 Description of the Hard Distributions

To construct a distribution, we first assign all letters  $a \in \{2, \dots, d_A\}$  to one of two sets,  $S_1$  and  $S_2$ . With probability  $\alpha := \min\{n/d_A, 1/2\}$ , we assign  $a$  to  $S_1$ , and to  $S_2$  otherwise. Using this assignment, we then define a distribution:

- $\forall a \in S_1, \forall c: P_{AC}[a, c] := 1/(2nd_C)$
- $\forall a \in S_2$ : for each  $c$  individually, we set

$$P_{AC}[a, c] = \begin{cases} \frac{\varepsilon}{d_A d_C} & \text{if } X = 0, \\ \text{uniformly at random } \frac{\varepsilon}{2d_A d_C} \text{ or } \frac{3\varepsilon}{2d_A d_C} & \text{if } X = 1. \end{cases} \quad (89)$$

The remaining probability mass is distributed uniformly over  $(a = 1, c)$ , for all  $c \in C$ . Note that  $P_{AC}$  is a product distribution if  $X = 0$ , and  $P_C$  is always uniform.

We will prove that a distribution sampled with  $X = 1$ , is far from being independent.

**Lemma 6.4.** *With high probability, for a distribution  $P_{AC}$  generated using  $X = 1$ , it holds that*

$$D_H^2(P_{AC}, P_AP_C) \geq \Omega(\varepsilon). \quad (90)$$

We will prove the above lemma in Appendix A.3.

## 6.2 Preliminary Properties of the Mutual Information

In order to prove our desired lower bounds, we first state a few statements about the mutual information. Note that these do not directly depend on our specific construction for the hardness distribution, and thus might be of independent interest.

**Lemma 6.5.** *Let  $X$  be a uniformly random bit, and  $A$  be a random variable taking values in a set  $S_A$ . Then*

$$2I(X : A) \leq \sum_{a \in S_A} \frac{(\Pr[A = a | X = 0] - \Pr[A = a | X = 1])^2}{\Pr[A = a | X = 0] + \Pr[A = a | X = 1]} \leq 12I(X : A). \quad (91)$$

The first inequality is obtained by bounding the KL-divergence by the  $\chi^2$ -divergence, and appeared in various forms in the literature (e.g., [DK16, App. A.1]). We include a proof here for completeness. To the best of our knowledge, the second inequality is new and shows that the first bound is essentially tight, up to constant factors.

Before proceeding to prove Lemma 6.5, let us first prove the following lemma, which will be used in the proof later.

**Lemma 6.6.** *Let  $a$  and  $b$  be non-negative numbers. Then*

$$a \log \left( \frac{2a}{a+b} \right) + b \log \left( \frac{2b}{a+b} \right) \geq \frac{1}{6} \frac{(a-b)^2}{a+b}. \quad (92)$$

*Proof.* If at least one of  $a$  and  $b$  is zero, the result follows trivially. Assume without loss of generality that  $a > b > 0$ . Let

$$u := \frac{a-b}{a+b} > 0. \quad (93)$$

Note that  $1-u = 2b/(a+b)$  and  $u < 1$ . We will use the following inequalities [Top07, Eq. (2) & (3)], which hold for  $x > -1$ ,

$$\log(1+x) \geq \frac{x}{1+x} \quad \text{and} \quad \frac{\log(1+x)}{x} \geq \frac{2}{2+x}. \quad (94)$$

The later implies  $\log(1+x) \geq 2x/(2+x)$  if  $x > 0$ . We can now bound

$$a \log \left( \frac{2a}{a+b} \right) + b \log \left( \frac{2b}{a+b} \right) = a \log(1+u) + b \log(1-u) \quad (95)$$

$$\geq a \frac{2u}{2+u} - b \frac{u}{1-u} \quad (96)$$

$$= \frac{1}{2+u} \frac{1}{2} \frac{(a-b)^2}{a+b} \geq \frac{1}{6} \frac{(a-b)^2}{a+b}. \quad (97)$$

This completes the proof of the lemma.  $\square$

Now we are ready to prove Lemma 6.5.

*Proof of Lemma 6.5.* We can use that  $\Pr[X = x] = 1/2$ , as well as  $\log(x) \leq x - 1$ . Then

$$I(X : A) := \sum_{\substack{a \in S_A \\ x \in \{0,1\}}} \Pr[A = a, X = x] \log \left( \frac{\Pr[A = a, X = x]}{\Pr[A = a] \Pr[X = x]} \right) \quad (98)$$

$$= \sum_{\substack{a \in S_A \\ x \in \{0,1\}}} \Pr[A = a | X = x] \Pr[X = x] \log \left( \frac{\Pr[A = a | X = x]}{\Pr[A = a]} \right) \quad (99)$$

$$\leq \sum_{\substack{a \in S_A \\ x \in \{0,1\}}} \frac{\Pr[A = a | X = x]}{2} \cdot \frac{\Pr[A = a | X = x] - \Pr[A = a]}{\Pr[A = a]} \quad (100)$$

$$= \sum_{\substack{a \in S_A \\ x \in \{0,1\}}} \frac{2\Pr[A = a | X = x]^2 - \Pr[A = a | X = x](\Pr[A = a | X = 0] + \Pr[A = a | X = 1])}{2(\Pr[A = a | X = 0] + \Pr[A = a | X = 1])} \quad (101)$$

$$= \frac{1}{2} \sum_{a \in S_A} \frac{(\Pr[A = a | X = 0] - \Pr[A = a | X = 1])^2}{\Pr[A = a | X = 0] + \Pr[A = a | X = 1]}. \quad (102)$$

For the other direction, we use Lemma 6.6 to bound the following

$$\frac{1}{2} \frac{(\Pr[A = a|X = 0] - \Pr[A = a|X = 1])^2}{\Pr[A = a|X = 0] + \Pr[A = a|X = 1]} \quad (103)$$

$$\leq 3 \sum_{x \in \{0,1\}} \Pr[A = a|X = x] \log \left( \frac{2 \Pr[A = a|X = x]}{\Pr[A = a|X = 0] + \Pr[A = a|X = 1]} \right) \quad (104)$$

$$= 6 \sum_{x \in \{0,1\}} \Pr[A = a|X = x] \Pr[X = x] \log \left( \frac{\Pr[A = a|X = x]}{\Pr[A = a]} \right). \quad (105)$$

Summing over all  $a$  yields the desired bound.  $\square$

Another property of the mutual information we will use is the following folklore result (for example, apply the chain rule,  $I(X:Y,Z) = I(X:Z) + I(X:Y|Z)$ , and use that for Markov chains,  $I(X:Y|Z) \leq I(X:Y)$  [CT06, p. 35]).

**Lemma 6.7.** *Given random variables forming a Markov chain  $Y - X - Z$ , i.e.,  $I(Y:Z|X) = 0$ , we have*

$$I(X:Y,Z) \leq I(X:Y) + I(X:Z). \quad (106)$$

### 6.3 Bounding the Mutual Information

As introduced earlier,  $K$  denotes the matrix indicating how often we witnessed the different samples. In particular,  $K[a,c]$  denotes how often the tuple  $(a,c)$  was observed in the provided samples. In the following, for an arbitrary  $a \in A$ , we denote by  $K_a$  the vector in  $\mathbb{N}^{\times d_C}$ , where  $K_a[c]$  denotes the number of times the pair  $(a,c)$  has occurred in  $K$ .

By the construction of our distributions defined in Eq. (89), the different  $K_a$ 's are independent since the  $a$ 's are assigned independently to either  $S_1$  or  $S_2$ . Further,  $P_C$  is uniform, such that we have by Lemma 6.7 that

$$I(X:K) \leq \sum_a I(X:K_a). \quad (107)$$

Note that due to dependencies within a specific  $a$ , we cannot directly apply a similar bound for dimension  $C$ , since knowing about a certain  $K_a[c]$  gives us information about the type of  $a$ , that is, whether  $a \in S_1$  or  $a \in S_2$ , which in turn provides information about other elements  $K_{ac'}$ . Following the ideas used by [DK16, Lemma 3.7], we perform a case distinction for each  $a \in A$ : (i) if there are less than two samples in  $K$  with a specific  $a$ , the corresponding mutual information is small and we calculate it directly, (ii) if there are at least two samples in  $K$  with the same  $a$ , we want to find a bound roughly of the following form, for some suitable  $\beta$ :

$$I(X:K_a) \leq \beta I(X:K_a|a \in S_2). \quad (108)$$

The exact statement we will derive, see (133), will be a bit more technical. Once we know the type of  $a$ , the elements within  $a$  are independent again, and we will be able to bound

$$I(X:K_a|a \in S_2) \leq \sum_c I(X:K_a[c]|a \in S_2). \quad (109)$$

Note that we intuitively expect the terms with  $a \in S_2$  to dominate, as the distribution over  $a \in S_1$  is independent of  $X$ .



In the following, we will define two sets:

$$U := \{\Lambda \in \mathbb{N}_0^{\times d_C} : \|\Lambda\|_1 \geq 2\}, \quad \text{and} \quad V := \{\Lambda \in \mathbb{N}_0^{\times d_C} : \|\Lambda\|_1 < 2\}. \quad (110)$$

For any  $\Lambda \in \mathbb{N}_0^{\times d_C}$ , let us define:

$$s_a(\Lambda) := \frac{(\Pr[K_a = \Lambda|X = 0] - \Pr[K_a = \Lambda|X = 1])^2}{\Pr[K_a = \Lambda|X = 0] + \Pr[K_a = \Lambda|X = 1]}. \quad (111)$$

We can use Lemma 6.5 to bound

$$I(X : K) \leq \sum_{a \in A} I(X : K_a) \leq \sum_{a \in A} \left[ \sum_{\Lambda \in U} s_a(\Lambda) + \sum_{\Lambda \in V} s_a(\Lambda) \right]. \quad (112)$$

We will separately bound the two terms in (112), and determine the maximal sample complexity such that they are still upper bounded by a constant, see , which will then directly imply our lower bounds.

To bound the first term involving  $U$ , we first show two supplementary statements.

**Claim 6.8.** *There exists a constant  $c$  such that for all  $a \in A$ ,  $x \in \{0, 1\}$ , and  $\Lambda \in U$ ,*

$$\frac{\Pr[K_a = \Lambda, a \in S_2|X = x]}{\Pr[K_a = \Lambda, a \in S_1|X = x]} \leq c \frac{\varepsilon^2 n^2}{\alpha d_A^2}. \quad (113)$$

*Proof.* We start with finding  $\beta$ , where we use that  $X$  and whether  $a \in S_2$  or not are independent of each other. First, note that

$$\Pr[K_a = \Lambda, a \in S_1|X = x] = \prod_{c=1}^{d_C} \frac{\binom{\frac{n}{2nd_C}}{K_a[c]} e^{-\frac{n}{2nd_C}}}{K_a[c]!} = \left(\frac{1}{2d_C}\right)^{\|K_a\|_1} e^{-\frac{1}{2}} \prod_{c=1}^{d_C} \frac{1}{K_a[c]!}, \quad (114)$$

$$\Pr[K_a = \Lambda, a \in S_2|X = 0] = \prod_{c=1}^{d_C} \frac{\binom{\frac{n\varepsilon}{d_A d_C}}{K_a[c]} e^{-\frac{n\varepsilon}{d_A d_C}}}{K_a[c]!} = \left(\frac{n\varepsilon}{d_A d_C}\right)^{\|K_a\|_1} e^{-\frac{n\varepsilon}{d_A}} \prod_{c=1}^{d_C} \frac{1}{K_a[c]!}, \quad (115)$$

$$\Pr[K_a = \Lambda, a \in S_2|X = 1] = \prod_{c=1}^{d_C} \frac{\frac{1}{2} \left( \binom{\frac{n\varepsilon}{2d_A d_C}}{K_a[c]} e^{-\frac{n\varepsilon}{2d_A d_C}} + \binom{\frac{3n\varepsilon}{2d_A d_C}}{K_a[c]} e^{-\frac{3n\varepsilon}{2d_A d_C}} \right)}{K_a[c]!} \quad (116)$$

$$\leq \left(\frac{3n\varepsilon}{2d_A d_C}\right)^{\|K_a\|_1} e^{-\frac{n\varepsilon}{2d_A}} \prod_{c=1}^{d_C} \frac{1}{K_a[c]!}. \quad (117)$$

Note that  $\Pr[K_a = \Lambda, a \in S_2|X = 0]$  is also upper bounded by (117). We can use the above to bound

$$\frac{\Pr[K_a = \Lambda, a \in S_2|X = x]}{\Pr[K_a = \Lambda, a \in S_1|X = x]} = \frac{\Pr[K_a = \Lambda|a \in S_2, X = x] \Pr[a \in S_2]}{\Pr[K_a = \Lambda|a \in S_1, X = x] \Pr[a \in S_1]} \quad (118)$$

$$\leq \frac{(1 - \alpha) \left(\frac{3n\varepsilon}{2d_A d_C}\right)^{\|K_a\|_1} e^{-\frac{n\varepsilon}{2d_A}} \prod_{c=1}^{d_C} \frac{1}{K_a[c]!}}{\alpha \left(\frac{1}{2d_C}\right)^{\|K_a\|_1} e^{-\frac{1}{2}} \prod_{c=1}^{d_C} \frac{1}{K_a[c]!}} \leq \frac{1}{\alpha} \left(\frac{3\varepsilon n}{d_A}\right)^{\|K_a\|_1}. \quad (119)$$

Note that we impose the restriction  $n\varepsilon/d_A \leq 1$ . We will find that this does not contradict with the bound we derive for  $n$ . This means  $\beta$  will take the maximal value of this expression, which by assumption is achieved for  $\|K_a\|_1$  minimal,  $\|K_a\|_1 = 2$ . Let  $c$  be a large enough constant, then the choice

$$\beta := c \frac{\varepsilon^2 n^2}{\alpha d_A^2} \quad (120)$$

satisfies the desired inequality.  $\square$

We will also need the following statement on the mutual information between  $X$  and  $K_a$ , conditioned that  $a \in S_2$ .

**Claim 6.9.** *It holds that  $I(X : K_a[c] | a \in S_2) = O\left(\frac{\varepsilon^2 n^2}{d_A^2 d_C^2}\right)$ .*

*Proof.* In the following, use the shorthand notation  $\gamma := \varepsilon n / (d_A d_C)$ . Recall that  $K_a[c]$  denotes the number of times the pair  $(a, c)$  has appeared in  $K$ . Assuming  $n \ll d_A d_C / \varepsilon$ , we find:

$$\Pr[K_a[c] = k | X = 0, a \in S_2] = \frac{e^{-\gamma}}{k!} \gamma^k = \frac{\gamma^k}{k!} [1 - \gamma + O(\gamma^2)] \quad (121)$$

$$\Pr[K_a[c] = k | X = 1, a \in S_2] = \frac{e^{-3\gamma/2}}{2 \cdot k!} \left(\frac{3\gamma}{2}\right)^k + \frac{e^{-\gamma/2}}{2k!} \left(\frac{\gamma}{2}\right)^k \quad (122)$$

$$= \frac{\gamma^k}{2 \cdot k! 2^k} \left[ 3^k + 1 - \frac{3^k 3\gamma}{2} - \frac{\gamma}{2} + O(\gamma^2) \right] \quad (123)$$

We then use Lemma 6.5 to bound

$$I(X : K_a[c] | a \in S_2) \leq \sum_{k=0}^{\infty} \frac{(\Pr[K_a[c] = k | X = 0] - \Pr[K_a[c] = k | X = 1])^2}{\Pr[K_a[c] = k | X = 0] + \Pr[K_a[c] = k | X = 1]} =: \sum_{k=0}^{\infty} s_{ac}(k). \quad (124)$$

We will bound the  $s_{ac}(k)$  values, and we will see that one of them dominates the other strong enough to bound  $I(X : K_a[c] | a \in S_2) \leq O(s_{ac}(k_{\max}))$ .

$$s_{ac}(k) = \frac{\gamma^k (1 - \gamma + O(\gamma^2) - [3^k + 1 - 3^k 3\gamma/2 - \gamma/2 + O(\gamma^2)]/2^{k+1})^2}{k! (1 - \gamma + [3^k + 1 - 3^k 3\gamma/2 - \gamma/2]/2^{k+1} + O(\gamma^2))} \quad (125)$$

$$s_{ac}(0) = \frac{(1 - \gamma - [2 - 2\gamma]/2 + O(\gamma^2))^2}{\Theta(1)} = O(\gamma^4), \quad (126)$$

$$s_{ac}(1) = \gamma \frac{(1 - \gamma - [4 - 5\gamma]/4 + O(\gamma^2))^2}{\Theta(1)} = O(\gamma^3), \quad (127)$$

$$(k \geq 2) \quad s_{ac}(k) \leq \frac{\gamma^k}{k!} \left(\frac{3}{2}\right)^{2k} = O(\gamma^k). \quad (128)$$

We see that there is a single dominant term,  $\gamma^2$ , compared to which the contributions of the other terms can be neglected (without loss of generality, we can assume, say,  $\gamma < 1/10$ ). Thus we have the following

$$I(X : K_a[c] | a \in S_2) = O\left(\frac{\varepsilon^2 n^2}{d_A^2 d_C^2}\right). \quad (129)$$

$\square$

We can now proceed to bound  $\sum_{\Lambda \in U} s_a(\Lambda)$ .

**Lemma 6.10.** For

$$n \leq \min \left\{ \frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \quad (130)$$

we have with high probability that  $\sum_{\Lambda \in U} s_a(\Lambda) \leq O(1)$ .

*Proof.* Using Claim 6.8 in the first inequality, as well as Lemma 6.5, we can now bound

$$\sum_{\Lambda \in U} s_a(\Lambda) = \sum_{\Lambda \in U} \frac{\left( \sum_{M \in \{S_1, S_2\}} \Pr[K_a = \Lambda, a \in M | X = 0] - \Pr[K_a = \Lambda, a \in M | X = 1] \right)^2}{\sum_{M \in \{S_1, S_2\}} \Pr[K_a = \Lambda, a \in M | X = 0] + \Pr[K_a = \Lambda, a \in M | X = 1]} \quad (131)$$

$$= \sum_{\Lambda \in U} \frac{(\Pr[K_a = \Lambda, a \in S_2 | X = 0] - \Pr[K_a = \Lambda, a \in S_2 | X = 1])^2}{\sum_{M \in \{S_1, S_2\}} \Pr[K_a = \Lambda, a \in M | X = 0] + \Pr[K_a = \Lambda, a \in M | X = 1]} \quad (132)$$

$$\leq \beta \sum_{\Lambda \in U} \frac{(\Pr[K_a = \Lambda, a \in S_2 | X = 0] - \Pr[K_a = \Lambda, a \in S_2 | X = 1])^2}{\Pr[K_a = \Lambda, a \in S_2 | X = 0] + \Pr[K_a = \Lambda, a \in S_2 | X = 1]} \quad (133)$$

$$\leq \beta \Pr[a \in S_2] \sum_{\Lambda} \frac{(\Pr[K_a = \Lambda | a \in S_2, X = 0] - \Pr[K_a = \Lambda | a \in S_2, X = 1])^2}{\Pr[K_a = \Lambda | a \in S_2, X = 0] + \Pr[K_a = \Lambda | a \in S_2, X = 1]} \quad (134)$$

$$\leq 12\beta \Pr[a \in S_2] I(X : K_a | a \in S_2) \quad (135)$$

$$\leq 12\beta \sum_c \Pr[a \in S_2] I(X : K_a[c] | a \in S_2). \quad (136)$$

We bound (136) further using Claim 6.9:

$$\beta \sum_{a,c} \Pr[a \in S_2] I(X : K_a[c] | a \in S_2) \leq \beta \sum_{a,c} (1 - \alpha) \frac{\varepsilon^2 n^2}{d_A^2 d_C^2} \leq d_A d_C \frac{\varepsilon^2 n^2}{\alpha d_A^2 d_C^2}. \quad (137)$$

By Lemma 6.3 and (112), we require  $\sum_{\Lambda \in U} s_a(\Lambda)$  to be upper bounded by a small constant (for our asymptotic purposes, let this constant be one). Recall from the construction of the hard distribution  $\alpha := \min\{n/d_A, 1/2\}$ , so now we consider two cases for  $\alpha$ , and solve for  $n$  under the constraint that  $\sum_{\Lambda \in U} s_a(\Lambda) \leq O(1)$ :

$$\alpha = \Theta(1) : \quad \frac{n^4 \varepsilon^4}{d_A^3 d_C} \leq O(1) \implies n \leq O\left(\frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}\right), \quad (138)$$

$$\alpha = \frac{n}{d_A} : \quad \frac{n^3 \varepsilon^4}{d_A^2 d_C} \leq O(1) \implies n \leq O\left(\frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}}\right). \quad (139)$$

□

Bounding the remaining term,  $\sum_a \sum_{\Lambda \in V} s_a(\Lambda)$ , is done by a direct calculation. The following lemma contains the results of the calculations, the proof can be found in Appendix A.3.

**Claim 6.11.** For the hard distribution as defined in Section 6.1, for a fixed  $a \in A$ , the following hold:

(i) If there is no occurrence of an index  $a$ , i.e., the occurrence vector  $K_a$  is equal to  $\ell_0 := \vec{0}$ , then we have:

$$\frac{(\Pr[K_a = \ell_0|X = 0] - \Pr[K_a = \ell_0|X = 1])^2}{\Pr[K_a = \ell_0|X = 0] + \Pr[K_a = \ell_0|X = 1]} \leq O\left(d_C^2 \left[\frac{\varepsilon n}{d_A d_C}\right]^4\right). \quad (140)$$

(ii) If a specific index  $a$  appears only once, i.e., the occurrence vector  $K_a$  satisfies  $\|K_a\|_1 = 1$ , we have:

$$\sum_{\ell, \|\ell\|_1=1} \frac{(\Pr[K_a = \ell|X = 0] - \Pr[K_a = \ell|X = 1])^2}{\Pr[K_a = \ell|X = 0] + \Pr[K_a = \ell|X = 1]} \leq O\left(d_C^3 \left[\frac{\varepsilon n}{d_A d_C}\right]^6 \frac{d_C}{\alpha}\right). \quad (141)$$

Using the above, finding the required bounds on  $n$  is direct.

**Lemma 6.12.** For

$$n \leq \min \left\{ \frac{d_A^{3/4} d_C^{2/4}}{\varepsilon}, \frac{d_A^{5/6} d_C^{1/3}}{\alpha \varepsilon} \right\} \quad (142)$$

we have with high probability that  $\sum_{\Lambda \in V} s_a(\Lambda) \leq O(1)$ .

*Proof.* First, note that by Claim 6.11,

$$\sum_a \sum_{\Lambda \in V} s_a(\Lambda) = \sum_a \frac{(\Pr[K_a = \ell_0|X = 0] - \Pr[K_a = \ell_0|X = 1])^2}{\Pr[K_a = \ell_0|X = 0] + \Pr[K_a = \ell_0|X = 1]} \quad (143)$$

$$+ \sum_a \sum_{\ell, \|\ell\|_1=1} \frac{(\Pr[K_a = \ell|X = 0] - \Pr[K_a = \ell|X = 1])^2}{\Pr[K_a = \ell|X = 0] + \Pr[K_a = \ell|X = 1]} \quad (144)$$

$$\leq O\left(d_A \left[ d_C^2 \left[\frac{\varepsilon n}{d_A d_C}\right]^4 + d_C^3 \left[\frac{\varepsilon n}{d_A d_C}\right]^6 \frac{d_C}{\alpha} \right]\right). \quad (145)$$

It remains to solve for  $n$ :

$$d_A \left[\frac{\varepsilon n}{d_A d_C}\right]^4 d_C^2 \leq O(1) \implies n \leq O\left(\frac{d_A^{3/4} d_C^{2/4}}{\varepsilon}\right), \quad (146)$$

$$d_A d_C^3 \left[\frac{\varepsilon n}{d_A d_C}\right]^6 \frac{d_C}{\alpha} \leq O(1) \implies n \leq O\left(\frac{d_A^{5/6} d_C^{1/3}}{\alpha \varepsilon}\right). \quad (147)$$

□

Using Lemma 6.10 and Lemma 6.12, our lower bound follows directly:

$$n = \Omega\left(\min\left\{\frac{d_A^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_C^{1/3}}{\varepsilon^{4/3}}\right\}\right). \quad (148)$$

This completes the proof of Theorem 6.1.

## 7 Conditional Independence Testing

In this section, we will prove our result on conditional independence testing in  $D_H^2$  distance, which implies a similar result for conditional mutual information testing. We will prove the following theorem.

**Theorem 7.1.** *The sample complexity of conditional independence testing in  $D_H^2$  distance, Problem 6, is*

$$\text{SC}_{\text{CI},H}(\varepsilon, d_A, d_B, d_C) = \tilde{O}\left(\max\{f_{\text{sym}}(\varepsilon, d_A, d_B, d_C), f_{\text{asym}}(\varepsilon, d_A, d_B, d_C)\}\right), \quad (149)$$

where we distinguish between terms that are symmetrical and asymmetrical in  $A$  and  $C$ , and assume w.l.o.g.  $d_A \geq d_C$ ,

$$f_{\text{sym}}(\varepsilon, d_A, d_B, d_C) := \max\left\{\frac{d_A^{1/2} d_B^{3/4} d_C^{1/2}}{\varepsilon}, \min\left\{\frac{d_A^{1/4} d_B^{7/8} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/7} d_B^{6/7} d_C^{2/7}}{\varepsilon^{8/7}}\right\}\right\}, \quad \text{and} \quad (150)$$

$$f_{\text{asym}}(\varepsilon, d_A, d_B, d_C) := \min\left\{\frac{d_A^{3/4} d_B^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_B^{2/3} d_C^{1/3}}{\varepsilon^{4/3}}\right\}. \quad (151)$$

The proof of Theorem 7.1 follows from our discussion of Algorithm 3 which decomposes the testing problem into multiple subproblems solved by two subroutines `CMITestingSmallRegime` (Algorithm 6) and `CMITestingLargeRegime` (Algorithm 8), discussed in Section 7.1 and Section 7.2. The sample complexity is obtained by combining the sample complexities of these subroutines.

As a corollary to the above theorem, we have the following result.

**Corollary 7.2** (Formalized bound of Result 2). *The sample complexity of conditional independence testing, Problem 4, is*

$$\text{SC}_{\text{CMI}}(\varepsilon, d_A, d_B, d_C) = \tilde{O}(\text{SC}_{\text{CI},H}(\varepsilon, d_A, d_B, d_C)), \quad (152)$$

*Proof.* This follows directly from Theorem 7.1 and Theorem 4.2.  $\square$

The underlying idea is to reduce the problem of conditional independence testing with respect to  $D_H^2$  distance to a polylogarithmic number of instances of equivalence testing in the  $\ell_2$  distance. Compared to independence testing described in Section 5, we now face an additional obstacle: sampling from the reference distribution  $Q = P_{AB}P_{BC}/P_B$  is non-trivial. The basic idea to overcome this issue is as follows. Assume that besides  $P_{ABC}$ , we could also sample from  $P_{AC|B=b}$ , for some  $b \in B$  of our choice. Then we could first select a random sample from  $P_{ABC}$ ,  $(a, b, c)$ , followed by a second sample  $(a', b, c')$  from  $P_{AC|B=b}$ . The sample  $(a, b, c')$  can then be treated as a sample coming from  $P_{AB}P_{BC}/P_B$ , as is easily seen: the first sample we draw needs to be of the form  $(a, b, \cdot)$ , which happens with probability  $p_{ab}$ . The probability of obtaining  $c'$  from the second sample is given by  $p_{bc'}/p_b$ , since we conditioned on  $b$ . Together, this results in

$$\Pr[(a, b, c') \text{ is obtained}] = \frac{p_{ab}p_{bc'}}{p_b}. \quad (153)$$

Since we only have sample access to  $P_{ABC}$ , simulating samples from  $Q_{ABC}$  becomes more challenging. In our approach, we split the distribution into two regimes, depending on a threshold  $\nu$ .

- (i) A large regime  $L$ , consisting of all  $(a, b, c)$  for which  $\hat{p}_b$  is above a certain threshold  $\nu$ , that is,  $B_L := \{b | \hat{p}_b \geq \nu\}$ , such that  $L = \{(a, b, c) | b \in B_L\}$ .

(ii) A small regime  $S$ , for the remaining  $b$ , defined as  $B_S := \{b | \hat{p}_b < \nu\}$ ,  $S = \{(a, b, c) | b \in B_S\}$ .

The reason for the split is that we apply different techniques to simulate (approximate) samples from  $Q_{ABC}$  in the respective regimes. To test  $P_{ABC}$  for conditional independence, it is then sufficient to test individually whether  $D_H^2(P_{ABC}^S, Q_{ABC}^S) \geq \varepsilon/2$  or  $D_H^2(P_{ABC}^L, Q_{ABC}^L) \geq \varepsilon/2$ , see the respective subsections, Section 7.1 and Section 7.2. However, both techniques are based on the basic idea of combining two samples with the same coordinate in  $B$  to simulate conditional independence as in (153). We now introduce our algorithm for conditional independence testing in  $D_H^2$  distance, Algorithm 3. For this, let

$$N := 10^2 \log(d_B) \max\{N_S, N_L\}, \quad (154)$$

where

$$N_S := 10^{10} \min \left\{ \frac{d_A^{1/4} d_B^{7/8} d_C^{1/4}}{\varepsilon_S}, \frac{d_A^{2/7} d_B^{6/7} d_C^{2/7}}{\varepsilon_S^{8/7}} \right\}, \quad (\text{Lemma 7.4}) \quad (155)$$

$$N_L := 10^6 c_{\text{eq}} \log \left( \frac{d_A d_B d_C}{\varepsilon_L} \right)^7 \max \left\{ \min \left\{ \frac{d_A^{3/4} d_B^{3/4} d_C^{1/4}}{\varepsilon_L}, \frac{d_A^{2/3} d_B^{2/3} d_C^{1/3}}{\varepsilon_L^{4/3}} \right\}, \right. \quad (156)$$

$$\left. \frac{d_A^{1/2} d_B^{3/4} d_C^{1/2}}{\varepsilon_L}, N_S \right\}, \quad (\text{Lemma 7.10}). \quad (157)$$

Here  $c_{\text{eq}}$  denotes the implicit constant used in equivalence testing of distributions in Corollary 5.7. This gives the sample complexity reported in Theorem 7.1.

---

**Algorithm 3:** Conditional Independence Testing in  $D_H^2$  distance

---

- Input:** A multiset  $\mathcal{S}$  of  $N$  triplets in  $A \times B \times C$ ,  $\varepsilon, \nu \in (0, 1)$  ▷  $N$ : see (154)  
**Output:** ‘Yes’ or ‘No’ ▷ ‘Yes’ indicates conditional independence  
**1**  $\mathcal{S}_L, \mathcal{S}_S$ , and  $\mathcal{S}_\nu \leftarrow$  remove random disjoint multisets from  $\mathcal{S}$ , such that  $|\mathcal{S}_L| = N_L$ ,  
 $|\mathcal{S}_S| = N_S$  and  $|\mathcal{S}_\nu| = 32 \log(10^3 d_B) N_S$  ▷  $N_L$ : Lemma 7.10 &  $N_S$ : Lemma 7.4  
**2**  $\hat{P}_b \leftarrow$  empirical frequency of  $b \in B$  in  $\mathcal{S}_\nu$  for every  $b \in B$   
**3**  $\nu \leftarrow 1/(2N_S)$ ,  $B_S \leftarrow \{b \in B | \hat{p}_b < \nu\}$ ,  $B_L \leftarrow \{b \in B | \hat{p}_b \geq \nu\}$ .  
**4**  $\text{resSmall} \leftarrow \text{CMITestingSmallRegime}(\mathcal{S}_S, B_S, \varepsilon/2)$  ▷ Algorithm 6 & Lemma 7.4  
**5**  $\text{resLarge} \leftarrow \text{CMITestingLargeRegime}(\mathcal{S}_L, B_L, \varepsilon/2)$  ▷ Algorithm 8 & Lemma 7.10  
**6** **return** ‘Yes’ if  $\text{resSmall} = \text{resLarge} = \text{‘Yes’}$ , otherwise ‘No’.
- 

We have the following lemma about the correctness of Algorithm 3.

**Lemma 7.3.** *Given  $N$  samples (as defined in (154)), Algorithm 3 returns ‘Yes’ if  $P_{ABC}$  is conditionally independent, ‘No’ if  $D_H^2(P_{ABC}, P_{A|B}P_{C|B}P_B) \geq \varepsilon$ , with probability at least  $2/3$ .*

*Proof.* In order to prove that Algorithm 3 succeeds probability at least  $2/3$ , we first show that the following subtasks succeed with high probability each:

- $B_S$  and  $B_L$  satisfy the conditions mentioned in Algorithm 6 ( $\forall b \in B_L : p_b \geq \nu/2$ ) and Algorithm 8 ( $\forall b \in B_S : p_b < 1/N_S$ ) with probability at least  $99/100$ . This follows directly from Lemma 5.8.
- Algorithm 6 with probability at least  $99/100$ , returns ‘No’ (i.e. ‘Far’) in case  $D_H^2(P_{ABC}^S, Q_{ABC}^S) \geq \varepsilon/2$ , and ‘Yes’ if  $P_{ABC}^S = Q_{ABC}^S$ . We present the proof in Section 7.1 (Lemma 7.4).

- Algorithm 8, with probability at least 99/100, returns ‘No’ (i.e. ‘Far’) in case  $D_H^2(P_{ABC}^L, Q_{ABC}^L) \geq \varepsilon/2$ , and ‘Yes’ if  $P_{ABC}^L = Q_{ABC}^L$ . We present the proof in Section 7.2 (Lemma 7.10).

A union bound then guarantees that the entire algorithm succeeds with probability at least  $97/100 > 2/3$ . If  $D_H^2(P_{ABC}, Q_{ABC}) \geq \varepsilon$ , then by the pigeonhole principle, either  $D_H^2(P_{ABC}^S, P_{ABC}^S) \geq \varepsilon/2$  or  $D_H^2(P_{ABC}^L, P_{ABC}^L) \geq \varepsilon/2$ , which we will detect with high probability as argued above, and if  $P_{ABC} = Q_{ABC}$ , then we also have equivalence on  $S$  and  $L$ .

The sample complexity is then the direct result of Lemma 7.4 and Lemma 7.10, since it is clear from the description of Algorithm 3 that  $N$  needs to be chosen such that  $N = |\mathcal{S}_\nu| + |\mathcal{S}_S| + |\mathcal{S}_L|$ .  $\square$

We now present the subroutines testing the large (Algorithm 8) and the small regime (Algorithm 6). We will begin our discussion with the small regime, as this will define the threshold separating the two regimes. For the large regime, we will be able to simulate a distribution  $\tilde{Q}_{ABC}$  such that  $\tilde{Q}_{ABC}^L = Q_{ABC}^L$ , allowing us to reuse ideas from Section 5, an approach which will not be possible in the small regime. In the following, we do not optimize for logarithmic or constant factors. However, we avoid to abstract these factors using  $O$ -notation to provide explicit thresholds in the following algorithms.

## 7.1 Testing The Small Regime

Our result for the small regime is as follows:

**Lemma 7.4.** *CMITestingSmallRegime (Algorithm 6) is correct with probability at least 99/100 if called using*

$$N_S := 10^{10} \min \left\{ \frac{d_A^{1/4} d_B^{7/8} d_C^{1/4}}{\varepsilon_S}, \frac{d_A^{2/7} d_B^{6/7} d_C^{2/7}}{\varepsilon_S^{8/7}} \right\} \quad (158)$$

*samples, and a set  $B_S$  such that for all  $b \in B_S : p_b \leq 1/N_S$ .*

### 7.1.1 Description of the Algorithm

From the definition of  $B_S$  in Algorithm 3 it is clear that the small regime will consider  $b$ 's for which  $p_b N_S < 1$ , which means that in expectation, we see less than one sample with  $b$ . While we can say little about the statistics of a *specific*  $b \in B_S$ , concentration bounds can tell us how many rare collisions (that is, a  $b \in B_S$  appearing at least twice among the samples), we will witness *in total*. We now define two algorithms **SimABC** and **SimABC<sub>C1</sub>**, which simulate samples based on collisions similar to the idea described in (153). The statistics between the outputs of **SimABC** and **SimABC<sub>C1</sub>** coincide if  $P_{ABC} = P_{AB}P_{BC}/P_B$ , and can be distinguished if  $D_H^2(P_{ABC}, P_{AB}P_{BC}/P_B) \geq \varepsilon$ .

The algorithms each take a set of  $\text{Poi}(\tilde{N}_S)$  samples from  $P_{ABC}$ , and process them in a randomly chosen order. They discard samples  $(a, b, c) \in L$  and only store samples  $(a', b', c') \in S$ . As soon as there are two stored samples with the same letter  $b$ , the process produces an output:

- The first algorithm, **SimABC** (Algorithm 4), mimics the original distribution. If **SimABC** witnesses first  $(a_1, b, c_1)$  and then  $(a_2, b, c_2)$ , it returns  $(a_1, b, c_1)$ . The second sample,  $(a_2, b, c_2)$ , is discarded.
- The second algorithm, **SimABC<sub>C1</sub>** (Algorithm 5), simulates samples which correspond to a conditionally independent version of **SimABC**. Given sample  $(a_1, b, c_1)$  at first and then  $(a_2, b, c_2)$ , it returns  $(a_1, b, c_2)$ . This is described in Algorithm 5.

An equivalent way, which we adopt in the description of Algorithm 4 and Algorithm 5 is to say that the respective algorithm takes randomly chosen pairs for which  $B$ -coordinates and processes them as described above. In the following, the notation  $\mathcal{S}|_{B=b}$  denotes the multiset obtained from  $\mathcal{S}$  by only keeping samples whose  $B$ -coordinate is  $b$ .

---

**Algorithm 4: SimABC**

---

**Input:** A multiset  $\mathcal{S}$  of  $N$  triplets in  $A \times B \times C, B_S \subseteq [d_B]$   
**Output:**  $M \in \mathbb{N}^{d_A \times d_B \times d_C}$

- 1  $M \leftarrow 0^{d_A \times d_B \times d_C}$
- 2 **for**  $b \in B_S$  **do**
- 3     **while**  $|\mathcal{S}|_{B=b}| \geq 2$  **do**
- 4          $(a_1, b, c_1), (a_2, b, c_2) \leftarrow$  pick two random triplets from  $\mathcal{S}|_{B=b}$ , without replacement
- 5          $M_{a_1, b, c_1} \leftarrow M_{a_1, b, c_1} + 1$
- 6 **return**  $M$

---



---

**Algorithm 5: SimABC<sub>C1</sub>**

---

**Input:** A multiset  $\mathcal{S}$  of  $N$  triplets in  $A \times B \times C, B_S \subseteq [d_B]$   
**Output:**  $M \in \mathbb{N}^{d_A \times d_B \times d_C}$

- 1  $M \leftarrow 0^{d_A \times d_B \times d_C}$
- 2 **for**  $b \in B_S$  **do**
- 3     **while**  $|\mathcal{S}|_{B=b}| \geq 2$  **do**
- 4          $(a_1, b, c_1), (a_2, b, c_2) \leftarrow$  pick two random triplets from  $\mathcal{S}|_{B=b}$ , without replacement
- 5          $M_{a_1, b, c_2} \leftarrow M_{a_1, b, c_2} + 1$
- 6 **return**  $M$

---

While the analysis of these statistics enables us to test for conditional independence, it comes with two challenges:

- First, the samples simulated are not i.i.d. samples, which prevents us from using existing equivalence testing algorithms as a black box.
- Second, the information our samples carry about the conditional independence of the original distribution can be ‘skewed’: intuitively, this comes from the fact that in order to produce a sample with a specific  $b$ , we need to see two samples from  $P_{ABC}$  which have the same  $b$  coordinate, which happens with a probability proportional to  $p_b^2$ , as opposed to  $p_b$ , putting smaller  $p_b$  at a disadvantage. This is also the reason why we need SimABC, and cannot directly compare the statistics of SimABC<sub>C1</sub> against those of  $P_{ABC}$ . See Section 7.1.2 for a more detailed discussion.

Note that for independence testing, we solved the more general problem of testing whether  $P_{AC} = Q_A Q_C$  for arbitrary distributions  $P_{AC}$  and  $Q_{AC}$ . This was possible since simulating samples from  $Q_A Q_C$  using samples from  $Q_{AC}$  is direct. The analogue no longer holds for conditional independence testing: simulating samples from  $Q_{ABC} = P_{AB} P_{BC} / P_B$  is a bottleneck which can dominate the sample complexity. In general, the problem of deciding whether a distribution  $P_{ABC}$  is equal to a given conditionally independent distribution  $R_{ABC}$  could be a different (and, in terms of sample complexity, likely simpler) problem.

We can now compare the statistics between samples produced by Algorithm 4 and Algorithm 5 to test for conditional independence:



---

**Algorithm 6:** CMITestingSmallRegime

---

**Input:** A multiset  $\mathcal{S}_S$  of  $N_S$  triplets in  $A \times B \times C$ ,  $B_S \subseteq [d_B]$ ,  $\varepsilon_S \in (0, 1)$   
**Output:** ‘Yes’, ‘No’ or ‘Abort’

- 1  $\tilde{N}_S := N_S/8$
- 2  $\forall i \in [4]: M_i \leftarrow \text{Poi}(\tilde{N}_S)$
- 3 **if**  $\sum_{i=1}^4 M_i > N_S$  **then**
- 4 |   **return** ‘Abort’
- 5  $\forall i \in [4]:$  Pick disjoint multisets  $\mathcal{S}_i$  of size  $M_i$  from  $\mathcal{S}_S$
- 6  $X \leftarrow \text{SimABC}(\mathcal{S}_1, B_S)$ ,  $X' \leftarrow \text{SimABC}(\mathcal{S}_2, B_S)$
- 7  $Y \leftarrow \text{SimABCCl}(\mathcal{S}_3, B_S)$ ,  $Y' \leftarrow \text{SimABCCl}(\mathcal{S}_4, B_S)$
- 8  $Z \leftarrow \sum_{abc} X_{abc} X'_{abc} - 2X_{abc} Y_{abc} + Y_{abc} Y'_{abc}$
- 9 **return** ‘Yes’ if  $Z < \frac{\tilde{N}_S^4 \varepsilon_S^4}{2 \cdot 4^7 d_A d_B^3 d_C}$ , otherwise ‘No’

---

To prove that the algorithm succeeds with high probability, we need to determine  $\mathbb{E}[Z]$  for the two cases where  $P_{ABC}^S$  is either conditionally independent, or gapped away from conditional independence, and we also require a bound on the variance of  $Z$ .

Before proceeding further, we remark on the fact that there might be instances where we do not need to test the light regime.

**Remark 7.5.** *In some instances, it might be that we can infer that  $P[B_S] \leq \varepsilon/(2 \log(d_A d_C))$ , which implies  $I(A : C|B_S) \leq \varepsilon/2$ , since  $\forall b \in B : I(A : C|B = b) \leq \log(d_A d_C)$ . It might also be that  $B_S = \emptyset$ . While potentially interesting for instance-optimal approaches, we do not treat these cases explicitly: in such instances,  $Z$  will simply never cross the threshold in Algorithm 6.*

### 7.1.2 Intuition for the Sample Complexity

Before we start with the formal proof, let us start with a brief discussion on the intuition for the sample complexity presented in Lemma 7.4. We will make a few simplifying assumptions, which nevertheless results in the asymptotically same sample complexity as our formal analysis.

For simplicity, let us assume that we have taken  $N_I$  samples. Intuitively, note that for a ‘rare’  $b$ , the probability of collision among  $N_I$  samples is proportional to  $O(p_b^2 N_I^2)$ . Thus, in total there will be roughly  $O(\sum_{i \in S} p_i^2 N_I^2) = O(\|P_{ABC}^S\|_2^2 N_I^2)$  rare collisions. Assume that we have a multiset containing pairs of elements that collided. Ignoring the fact that we sample without replacement, the probability of drawing a specific  $b$  from this multiset is then roughly as follows

$$\Pr[\text{we draw } b] = \Pr[b \text{ appears in a collision}] \Pr[\text{we draw } b | b \text{ appears in a collision}] \quad (159)$$

$$= O\left((N_I p_b)^2 \frac{1}{(N_I \|P_{ABC}^S\|_2^2)^2}\right) \quad (160)$$

$$\geq p_b \cdot O\left(\frac{\varepsilon_S}{4 d_B \|P_{ABC}^S\|_2^2}\right), \quad (161)$$

where  $\varepsilon_S$  denotes the threshold we test for in Algorithm 6. In the last inequality above, we used that we can neglect  $p_b \leq \varepsilon_S/(2d_B)$ , as such  $b$  combined carries less than  $\varepsilon_S/2$  weight in total. Compared to the original distribution, where a specific  $b$  appears with probability  $p_b$ , the probability to draw a specific  $b$  from the multiset of ‘rare’ collisions is skewed by a factor  $\gamma = O(\varepsilon_S/(d_B \|P_{ABC}^S\|_2^2))$ . This effect needs to be taken into account when we want to test for properties of the original distributions: if we have  $D_H^2(P_{ABC}, Q_{ABC}) \geq \varepsilon_S$ , the bias our sampling introduces would require us to test for an

increased precision  $\eta := \varepsilon_S \cdot \tilde{O}(\varepsilon_S / (2d_B \|P_{ABC}^S\|_2^2))$  to surely detect if  $P$  is conditionally independent or not. We can bound this further using  $\|P_{ABC}^S\|_2^2 \leq O(\min\{d_B/N_I^2, 1/N_I\})$ , see Fact 3.3.

Performing the equivalence testing to precision  $\eta$  in the squared Hellinger distance requires  $\Theta(\min\{d^{2/3}/\eta^{4/3}, d^{3/4}/\eta\})$  i.i.d. samples [DKW18, Thm. 4.2]. For the moment, we ignore the fact that our samples are not independent. It is easy to check that here, we always have  $d^{2/3}/\eta^{4/3} \leq d^{3/4}/\eta$ , which is why we only consider the first option,  $O(d^{2/3}/\eta^{4/3})$ . The number of rare collisions we witness,  $O(N_I^2 \|P_{ABC}^S\|_2^2)$ , needs to be large enough to perform the testing, such that we require

$$N_I^2 \|P_{ABC}^S\|_2^2 \geq \frac{(d_A d_B d_C)^{2/3} (\|P_{ABC}^S\|_2^2 d_B)^{4/3}}{\varepsilon_S^{8/3}} \implies N_I \geq O\left(\left(\|P_{ABC}^S\|_2^2\right)^{1/6} \frac{d_B (d_A d_C)^{1/3}}{\varepsilon_S^{4/3}}\right). \quad (162)$$

We then insert the two bounds on  $\|P_{ABC}^S\|_2^2$  (see Fact 3.3), and find for the sufficient  $N_I$

$$\|P_{ABC}^S\|_2^2 \leq \frac{d_B}{N_I^2} : \quad N_I \geq O\left(\frac{d_B^{1/6}}{N_I^{1/3}} \frac{d_B (d_A d_C)^{1/3}}{\varepsilon_S^{4/3}}\right) \implies N_I \geq O\left(\frac{d_B^{7/8} (d_A d_C)^{1/4}}{\varepsilon_S}\right), \quad (163)$$

and

$$\|P_{ABC}^S\|_2^2 \leq \frac{1}{N_I} : \quad N_I \geq O\left(\frac{1}{N_I^{1/6}} \frac{d_B (d_A d_C)^{1/3}}{\varepsilon_S^{4/3}}\right) \implies N_I \geq O\left(\frac{d_B^{6/7} (d_A d_C)^{2/7}}{\varepsilon_S^{8/7}}\right), \quad (164)$$

resulting in the sample complexity reported in Lemma 7.4. We already pointed out that a few assumptions we made along the way do not hold. In the following, we will present an approach which addresses these issues. We will find that the actual sample complexity still coincides with our informal derivation here, which is why we included the above discussion to provide some intuition on the form of the sample complexity.

### 7.1.3 Formal Analysis

We will first argue in Lemma 7.6 that the condition we imposed on the size of  $\text{Poi}(\tilde{N}_S)$  (see Algorithm 6) can be neglected in our analysis.

**Lemma 7.6.** *The cut-off imposed on  $\text{Poi}(\tilde{N}_S)$  in Algorithm 6 decreases the probability of success by at most  $\Pr[\sum_i M_i > N_S]$ .*

*Proof.* Let  $A$  be the event that the algorithm succeeds. Then, if  $\Pr[A] \geq 1 - \delta_1$ ,  $\Pr[\sum_i M_i \leq N_S] \geq 1 - \delta_2$  and  $\delta_1 + \delta_2 < 1$ , Bayes rule implies that

$$\Pr\left[A \mid \sum_i M_i \leq N_S\right] = 1 - \frac{\Pr[\bar{A}, \sum_i M_i \leq N_S]}{\Pr[\sum_i M_i \leq N_S]} \quad (165)$$

$$\geq 1 - \frac{\Pr[\bar{A}]}{\Pr[\sum_i M_i \leq N_S]} \geq 1 - \frac{\delta_1}{1 - \delta_2} \geq 1 - \delta_1 - \delta_2. \quad (166)$$

□

We denote by  $X_{abc}$  the number of times  $\text{SimABC}$  returns the sample  $(a, b, c)$ , and define  $Y_{abc}$  analogously for  $\text{SimABC}_{C1}$ . Further, let  $X_b = \sum_{a,c} X_{abc}$ .

For simplicity, we write  $x_b := \tilde{N}_S p_b$  in the following. Importantly for our analysis, our choice for  $B_S$  guarantees that  $x_b \leq 1$ . In the following, for ease of readability, we will often omit the index

$abc$  when it is clear from the context. When providing Algorithm 4 and Algorithm 5 with  $\text{Poi}(\tilde{N}_S)$  samples each, we have the following.

$$\Pr[X_{abc} = k] = \sum_{\ell=k}^{\infty} \Pr[X_{abc} = k | X_b = \ell] \Pr[X_b = \ell] \quad (167)$$

$$= \sum_{\ell=k}^{\infty} \binom{\ell}{k} p_{ac|b}^k (1 - p_{ac|b})^{\ell-k} \left( \frac{x_b^{2\ell} e^{-x_b}}{(2\ell)!} + \frac{x_b^{2\ell+1} e^{-x_b}}{(2\ell+1)!} \right). \quad (168)$$

For SimABC (Algorithm 4), two samples from  $b$  result in an output for  $(a, b, c)$  if the first sample reads  $(a, b, c)$ , and the second one is arbitrary, resulting in  $p_{ac|b}$ . For SimABC<sub>CI</sub> (Algorithm 5), two samples from  $b$  result in an output for  $(a, b, c)$  if the first sample has the  $A$ -coordinate equal to  $a$ , and the second one has  $C$ -coordinate equal to  $c$ , resulting in  $q_{ac|b} := p_{a|b} p_{c|b}$ .

Note that our process outputs  $\ell$  samples of a specific  $b$ -value exactly if we witnessed either  $2\ell$  or  $2\ell + 1$  samples with this specific  $b$  from the original distribution. Analogously,

$$\Pr[Y_{abc} = k] = \sum_{\ell=k}^{\infty} \Pr[Y_{abc} = k | Y_b = \ell] \Pr[Y_b = \ell] \quad (169)$$

$$= \sum_{\ell=k}^{\infty} \binom{\ell}{k} q_{ac|b}^k (1 - q_{ac|b})^{\ell-k} \left( \frac{x_b^{2\ell} e^{-x_b}}{(2\ell)!} + \frac{x_b^{2\ell+1} e^{-x_b}}{(2\ell+1)!} \right). \quad (170)$$

It is clear that  $\forall b \in [d_B], \forall k \in \mathbb{N} : \mathbb{E}[X_b^k] = \mathbb{E}[Y_b^k]$ .

Now we will bound the expectations of the random variables  $X_{abc}, Y_{abc}, X_b, Y_b$  and other related quantities, whose proof will be presented in Appendix A.4.

**Lemma 7.7.** *For  $X_{abc}$  and  $Y_{abc}$  distributed according to (168), and (170), respectively, and  $q_{ac|b} = p_{a|b} p_{c|b}$  as defined before, the following hold:*

$$(i) \mathbb{E}[X_b] = \mathbb{E}[Y_b] = \frac{\tilde{N}_S p_b}{2} - \frac{e^{-\tilde{N}_S p_b}}{2} \sinh(\tilde{N}_S p_b), \text{ and } \mathbb{E}[X_b^2] = \mathbb{E}[Y_b^2] = \frac{\tilde{N}_S^2 p_b^2}{4} - \frac{\tilde{N}_S p_b e^{-2\tilde{N}_S p_b}}{4} + \frac{e^{-\tilde{N}_S p_b}}{4} \sinh(\tilde{N}_S p_b).$$

$$(ii) \mathbb{E}[X_{abc}] = p_{ac|b} \mathbb{E}[X_b], \text{ and } \mathbb{E}[X_{abc}^2] = p_{ac|b} (1 - p_{ac|b}) \mathbb{E}[X_b] + p_{ac|b}^2 \mathbb{E}[X_b^2].$$

$$\text{Similarly, } \mathbb{E}[Y_{abc}] = q_{ac|b} \mathbb{E}[X_b], \text{ and } \mathbb{E}[Y_{abc}^2] = q_{ac|b} (1 - q_{ac|b}) \mathbb{E}[X_b] + q_{ac|b}^2 \mathbb{E}[X_b^2].$$

$$(iii) \text{ For } (a, c) \neq (a', c'), \mathbb{E}[X_{abc} X_{a'bc'}] \leq 12 p_{ac|b} p_{a'c'|b} \tilde{N}_S^3 p_b^3, \text{ and } \mathbb{E}[Y_{abc} Y_{a'bc'}] \leq 12 q_{ac|b} q_{a'c'|b} \tilde{N}_S^3 q_b^3.$$

Our tester will use the following estimator, which uses two independent multisets of samples  $X, X'$ , generated by running two instances of SimABC (Algorithm 4) on different sets of samples from  $P_{ABC}$ . Analogously, our estimator also requires two sets  $Y, Y'$  generated by different instances of SimABC<sub>CI</sub> (Algorithm 5). Then, our tester determines the variable  $Z$  according to

$$Z = \sum_{abc} Z_{abc}, \quad Z_{abc} := X_{abc} X'_{abc} - 2X_{abc} Y_{abc} + Y_{abc} Y'_{abc}. \quad (171)$$

We will calculate  $\mathbb{E}[Z]$  and  $\text{Var}[Z]$ , which we will express in terms of  $X$  and  $Y$ . For better readability, we write  $i$  instead of  $abc$ . Note that a very interesting property of this estimator is that, no matter how  $X$  and  $Y$  are distributed, we have

$$\mathbb{E}[Z] = \sum_i \mathbb{E}[X_i] \mathbb{E}[X'_i] - 2\mathbb{E}[X_i] \mathbb{E}[Y_i] + \mathbb{E}[Y_i] \mathbb{E}[Y'_i] = \sum_i (\mathbb{E}[X_i] - \mathbb{E}[Y_i])^2, \quad (172)$$

which is zero if and only if  $\forall i : \mathbb{E}[X_i] = \mathbb{E}[Y_i]$ . For our specific algorithms  $\text{SimABC}$  and  $\text{SimABC}_{\text{Cl}}$ , we have the following results about the expectation of  $Z$ .

**Lemma 7.8.** *For  $i \in [4]$ , let  $\mathcal{S}_i$  be multisets of i.i.d. samples where  $|\mathcal{S}_i|$  is determined by drawing from  $\text{Poi}(\tilde{N}_S)$ . Let  $X = \text{SimABC}(\mathcal{S}_1, B_S)$ ,  $X' = \text{SimABC}(\mathcal{S}_2, B_S)$  (Algorithm 4), and  $Y = \text{SimABC}_{\text{Cl}}(\mathcal{S}_3, B_S)$ ,  $Y' = \text{SimABC}_{\text{Cl}}(\mathcal{S}_4, B_S)$  (Algorithm 5), and  $Z$  constructed from  $X$ ,  $X'$ ,  $Y$ , and  $Y'$  according to (171). Then*

$$(i) D_H^2(P_{ABC}^S, Q_{ABC}^S) \geq \varepsilon_S \text{ implies } \mathbb{E}[Z] \geq \tilde{N}_S^4 \varepsilon_S^4 / (4^7 d_A d_B^3 d_C),$$

$$(ii) P_{ABC}^S = Q_{ABC}^S \text{ implies } \mathbb{E}[Z] = 0.$$

*Proof.* Let us start with the proof of (i).

(i) Recall the definitions of  $Z$  and  $Z_{abc}$  from (171). Then, for a specific index  $abc$ , we have:

$$\mathbb{E}[Z_{abc}] = (\mathbb{E}[X_{abc}] - \mathbb{E}[Y_{abc}])^2. \quad (173)$$

Assume without loss of generality,  $p_{ac|b} > q_{ac|b}$  (otherwise flip the roles of  $p$  and  $q$ ). Then, using Lemma 7.7, we find that

$$\mathbb{E}[X_{abc}] - \mathbb{E}[Y_{abc}] = (p_{ac|b} - q_{ac|b}) \left[ \frac{\tilde{N}_S p_b}{2} - \frac{1}{2} \sinh(\tilde{N}_S p_b) e^{-\tilde{N}_S p_b} \right] \quad (174)$$

$$= (p_{ac|b} - q_{ac|b}) \left[ \frac{\tilde{N}_S p_b}{2} - \frac{1}{4} + \frac{1}{4} e^{-2\tilde{N}_S p_b} \right] \quad (175)$$

$$\geq \frac{p_{ac|b} - q_{ac|b}}{4} \tilde{N}_S^2 p_b^2, \quad (176)$$

where in the last step, we used that for  $x < 1$ ,  $e^{-x} \geq 1 - x + x^2/4$ , according to Lemma 3.9.

Next, we introduce  $S_+ := \{b \in S | p_b \geq \varepsilon_S / (4d_B)\}$  and  $S_- := \{b \in S | p_b < \varepsilon_S / (4d_B)\}$ . So combining the above with (173), and using  $q_{abc} = q_{ac|b} p_b$ , we get

$$\mathbb{E}[Z] = \sum_{b \in S, a, c} \frac{(p_{ac|b} - q_{ac|b})^2}{16} \tilde{N}_S^4 p_b^4 \quad (177)$$

$$= \sum_{\substack{a, c \\ b \in S_+}} \frac{(p_{ac|b} - q_{ac|b})^2}{16} \tilde{N}_S^4 p_b^4 + \sum_{\substack{a, c \\ b \in S_-}} \frac{(p_{ac|b} - q_{ac|b})^2}{16} \tilde{N}_S^4 p_b^4 \quad (178)$$

$$\geq \frac{\tilde{N}_S^4 p_b^2}{16} \|P_{ABC}^{S_+} - Q_{ABC}^{S_+}\|_2^2 \quad (179)$$

$$\geq \frac{\tilde{N}_S^4 \varepsilon_S^2}{16 \cdot 4^2 d_B^2} \|P_{ABC}^{S_+} - Q_{ABC}^{S_+}\|_2^2. \quad (180)$$

In the next step, we use Fact 3.2, which states that for any distributions  $P, Q$ ,  $D_H^2(P, Q) \leq 2\sqrt{d_A d_B d_C} \|P - Q\|_2^2$  holds. Thus we can write the above as the following

$$\mathbb{E}[Z] \geq \frac{\tilde{N}_S^4 \varepsilon_S^2}{4^6 d_A d_B^3 d_C} \left( D_H^2(P_{ABC}^{S_+}, Q_{ABC}^{S_+}) \right)^2. \quad (181)$$

Finally, using the fact that for any  $b \in S_-$ ,  $p_b \leq \varepsilon_S/(4d_B)$ , we can say that

$$D_H^2(P_{ABC}^{S_-}, Q_{ABC}^{S_-}) \leq \frac{\varepsilon_S}{4d_B} \sum_{b \in S_-, a, c} D_H^2(P_{AC|B=b}, Q_{AC|B=b}) \leq \frac{\varepsilon_S}{4d_B} \cdot 2d_B = \frac{\varepsilon_S}{2}. \quad (182)$$

This implies that if  $D_H^2(P_{ABC}^S, Q_{ABC}^S) \geq \varepsilon_S$ , then  $D_H^2(P_{ABC}^{S_+}, Q_{ABC}^{S_+}) \geq \varepsilon_S/2$ , and we find

$$\mathbb{E}[Z] \geq \frac{\tilde{N}_S^4 \varepsilon_S^4}{4^7 d_A d_B^3 d_C}. \quad (183)$$

- (ii) Now let us consider the case when  $P_{ABC}^S = Q_{ABC}^S$ . In this case,  $p_{ac|b} = p_{a|b}p_{c|b} = q_{ac|b}$ , such that (168) and (170) are identical. Following the definition of the estimator  $Z$  as stated in (172), we see that  $\mathbb{E}[Z] = 0$ .

□

The variance of  $Z$  (defined in Equation (171)) can be bounded as follows.

**Lemma 7.9.** *For  $i \in [4]$ , let  $\mathcal{S}_i$  be multisets of i.i.d. samples where  $|\mathcal{S}_i|$  is determined by drawing from  $\text{Poi}(\tilde{N}_S)$ . Let  $X = \text{SimABC}(\mathcal{S}_1, B_S)$ ,  $X' = \text{SimABC}(\mathcal{S}_2, B_S)$  (Algorithm 4), and  $Y = \text{SimABC}_{\text{Cl}}(\mathcal{S}_3, B_S)$ ,  $Y' = \text{SimABC}_{\text{Cl}}(\mathcal{S}_4, B_S)$  (Algorithm 5), and  $Z$  constructed from  $X$ ,  $X'$ ,  $Y$ , and  $Y'$  according to (171). Then*

$$\text{Var}[Z] \leq 2 \cdot 10^3 (\|P_{ABC}^S\|_2^2 + \|Q_{ABC}^S\|_2^2) \tilde{N}_S^2. \quad (184)$$

We will prove Lemma 7.9 in Appendix A.4. Assuming the above lemma holds, let us proceed to prove Lemma 7.4, which we restate for better readability.

**Lemma 7.4.** *CMITestingSmallRegime (Algorithm 6) is correct with probability at least 99/100 if called using*

$$N_S := 10^{10} \min \left\{ \frac{d_A^{1/4} d_B^{7/8} d_C^{1/4}}{\varepsilon_S}, \frac{d_A^{2/7} d_B^{6/7} d_C^{2/7}}{\varepsilon_S^{8/7}} \right\} \quad (158)$$

*samples, and a set  $B_S$  such that for all  $b \in B_S$ :  $p_b \leq 1/N_S$ .*

*Proof.* From Lemma 7.8, we know that when  $P_{ABC}^S$  and  $Q_{ABC}^S$  are far, we have

$$\mathbb{E}[Z] \geq \frac{\tilde{N}_S^4 \varepsilon_S^4}{4^7 d_A d_B^3 d_C}. \quad (185)$$

Next, we bound the variance of  $Z$ . For both the distributions  $P_{ABC}$  and  $Q_{ABC}$ , we can bound analogously

$$\|P_{ABC}^S\|_2^2 \leq \sum_{abc \in S} p_{ac|b}^2 p_b^2 \leq \sum_{b \in S} p_b^2 = \|P_B^S\|_2^2 \leq \min \left\{ \frac{1}{\tilde{N}_S}, \frac{d_B}{\tilde{N}_S^2} \right\}, \quad (186)$$

and

$$\|Q_{ABC}^S\|_2^2 \leq \sum_{abc \in S} q_{ac|b}^2 q_b^2 \leq \sum_{b \in S} q_b^2 = \|Q_B^S\|_2^2 \leq \min \left\{ \frac{1}{\tilde{N}_S}, \frac{d_B}{\tilde{N}_S^2} \right\}, \quad (187)$$

such that from Lemma 7.9, it follows that

$$\text{Var}[Z] \leq 2 \cdot 10^3 \min\{\tilde{N}_S, d_B\}. \quad (188)$$

We will apply Chebyshev's inequality (see Lemma 3.5). For this purpose, we simply set  $t = \mathbb{E}[Z]/2$ . Thus we need the following in order to ensure that the output of our estimator allows us to reliably distinguish the two cases:

$$\frac{\text{Var}[Z]}{\mathbb{E}[Z]^2/4} \leq \frac{1}{100} \implies \sqrt{\min\{\tilde{N}_S, d_B\}} \leq \frac{\tilde{N}_S^4 \varepsilon^4}{10 \cdot (2 \cdot 10)^{3/2} \cdot 2 \cdot 4^7 d_A d_B^3 d_C}. \quad (189)$$

For readability, let  $\alpha := 10 \cdot (2 \cdot 10)^{3/2} \cdot 2 \cdot 4^7 < 10^9$ . Combining (188) with (189), we have two different ways to choose  $\tilde{N}_S$ :

– **Case 1:**

$$\frac{\tilde{N}_S^4 \varepsilon^4}{\alpha d_A d_B^3 d_C} \geq d_B^{1/2} \implies \tilde{N}_S \geq \alpha \frac{d_A^{1/4} d_B^{7/8} d_C^{1/4}}{\varepsilon_S}, \quad (190)$$

– **Case 2:**

$$\frac{\tilde{N}_S^4 \varepsilon^4}{\alpha d_A d_B^3 d_C} \geq \tilde{N}_S^{1/2} \implies \tilde{N}_S \geq \alpha \frac{d_A^{2/7} d_B^{6/7} d_C^{2/7}}{\varepsilon_S^{8/7}}. \quad (191)$$

The total number of samples that is sufficient to correctly test the small regime for conditional independence with probability of success at least 99/100 is thus

$$N_S \geq 10^{10} \min \left\{ \frac{d_A^{1/4} d_B^{7/8} d_C^{1/4}}{\varepsilon_S}, \frac{d_A^{2/7} d_B^{6/7} d_C^{2/7}}{\varepsilon_S^{8/7}} \right\}, \quad (192)$$

where an additional factor 8 comes from the fact that  $N_S = \tilde{N}_S/8$ .  $\square$

## 7.2 Testing the Large Regime

In this subsection, we will prove the sample complexity of the large regime. Recall that  $c_{\text{eq}}$  is introduced in Corollary 5.7 to abstract the implicit constant in a testing subroutine.

**Lemma 7.10.** *CMITestingLargeRegime (Algorithm 8) correctly tests the large regime of  $P_{ABC}$  for conditional independence with probability at least 99/100 if called using*

$$N_L := 10^6 c_{\text{eq}} \log \left( \frac{d_A d_B d_C}{\varepsilon_L} \right)^7 \max \left\{ \min \left\{ \frac{d_A^{3/4} d_B^{3/4} d_C^{1/4}}{\varepsilon_L}, \frac{d_A^{2/3} d_B^{2/3} d_C^{1/3}}{\varepsilon_L^{4/3}} \right\}, \frac{d_A^{1/2} d_B^{3/4} d_C^{1/2}}{\varepsilon_L}, N_S \right\} \quad (193)$$

*samples from  $P_{ABC}$ , and a set  $B_L$  such that for all  $b \in B_L$  :  $p_b \geq 1/N_S$ , for  $N_S$  as defined in Lemma 7.4 and  $d_A \geq d_C$ .*

*Proof.* Lemma 7.10 follows directly from the proof of correctness of Algorithm 8, presented in Lemma 7.17.  $\square$

Before we describe the algorithm `CMITestingLargeRegime`, let us first present a few subroutines that will be used in `CMITestingLargeRegime`.

### 7.2.1 Subroutines

Unlike in the small regime, where we were not able to simulate the desired distribution exactly, we will simulate samples from a distribution  $\tilde{Q}_{ABC}$  which satisfies  $\tilde{Q}_{ABC}^L = (P_{AB}P_{C|B})^L$ . However, other than for independence testing, this is no longer straightforward. We now describe an algorithm `SimABCC|BLarge` to sample from  $P_{AB}P_{C|B}$  in the large regime and prove its correctness.

---

#### Algorithm 7: `SimABCC|BLarge`

---

**Input:** A multiset  $\mathcal{S}$  of triplets from  $A \times B \times C$ ,  $B_L \subseteq [d_B]$   
**Output:** A multiset  $\mathcal{S}_{\text{out}}$  of triplets from  $A \times B \times C$  or ‘**Abort**’

- 1 Pick disjoint multisets  $\mathcal{S}_1$  and  $\mathcal{S}_2$  from  $\mathcal{S}$  such that  $|\mathcal{S}_1| = 4|\mathcal{S}|/5$ ,  $|\mathcal{S}_2| = |\mathcal{S}|/5$
- 2 Define  $d_B$  empty queues  $Q[1], \dots, Q[d_B]$ ,  $\mathcal{S}_{\text{out}} \leftarrow \emptyset$
- 3 **while**  $\mathcal{S}_1 \neq \emptyset$  **do**
- 4      $(a, b, c) \leftarrow$  remove a random triplet from  $\mathcal{S}_1$
- 5     insert  $a$  into  $Q[b]$
- 6 **while**  $\mathcal{S}_2 \neq \emptyset$  **do**
- 7      $(a, b, c) \leftarrow$  remove a random triplet from  $\mathcal{S}_2$
- 8     **if**  $b \in B_L$  **then**
- 9         **if**  $Q[b] = \emptyset$  **then**
- 10             **return** ‘**Abort**’
- 11              $a' \leftarrow$  dequeue first element from  $Q[b]$
- 12              $\mathcal{S}_{\text{out}} \leftarrow \mathcal{S}_{\text{out}} \cup \{(a', b, c)\}$
- 13         **else**
- 14              $\mathcal{S}_{\text{out}} \leftarrow \mathcal{S}_{\text{out}} \cup \{(a, b, c)\}$   $\triangleright$  No modification required if  $b \notin B_L$
- 15 **return**  $\mathcal{S}_{\text{out}}$ .

---

The correctness of Algorithm 7 is presented in the following lemma.

**Lemma 7.11.** *Let  $\zeta \in (0, 1)$ ,  $\xi \in (0, 1)$ ,  $P_{ABC}$  be a tripartite distribution, and  $K := \{(a, b, c) | p_b \geq \xi\}$ . Using  $N \geq 8 \log(4d_B/\zeta)/\xi$  samples from  $P_{ABC}$ , Algorithm 7 simulates  $N/5$  samples from  $\tilde{Q}_{ABC}$ , a distribution which satisfies  $\tilde{Q}_{ABC}^K = (P_{AB}P_{C|B})^K$  and  $\tilde{Q}_B = Q_B$ , with probability at least  $1 - \zeta$ .*

*Proof.* To simulate samples from the reference distribution  $Q = P_{AB}P_{BC}/P_B$ , we perform two phases of sampling:

1. In the first phase, we take  $4N/5$  samples  $(a_1, b_1, c_1), \dots, (a_N, b_N, c_N)$  from  $P_{ABC}$ , and sort their  $A$ -coordinate into different queues according to their  $B$ -coordinate.
2. In the second phase, we take the remaining set of  $N/5$  samples from  $P_{ABC}$ . If we obtain  $(a, b, c)$  as a sample where  $b \in K$ , we pick the first element from the respective  $b$ -queue, say,  $a'$ , and we return  $(a', b, c)$ . Samples of the form  $(a, b, c)$  with  $b \notin K$  will be returned directly.

Since  $p_b$  for  $b \in K$  is lower bounded by  $\xi$ , Lemma 5.8 guarantees that for a given  $b$ , the number of samples we see in the two phases is within a factor two of the respective expected value. Since

$$\mathbb{E}[|Q[b]|] \geq 4\mathbb{E}[\#\text{queries to } Q[b]], \quad (194)$$

there are, with high probability, always enough samples in each queue.

To see that this procedure produces samples from  $\tilde{Q}_{ABC}^K$ , we ask for the probability to obtain  $(a', b, c) \in K$ . First, the sample we draw needs to be of the form  $(\cdot, b, c)$ , which happens with

probability  $p_{bc}$ . The probability to draw  $a'$  from the  $b$ -queue is given by  $p_{a'b}/p_b$ , since we conditioned on  $b$ . Together, this results in

$$\Pr[(a', b, c) \text{ is obtained}] = \frac{p_{a'b}p_{bc}}{p_b}. \quad (195)$$

Since we only test regions with non-rare  $b$ 's, it is irrelevant from where we sample them, as long as  $p_b$ 's for  $b \in K$  are not affected. The equality on  $B$ ,  $\tilde{Q}_B = Q_B$  follows directly, since if we take a random sample from  $\mathcal{S}_2$ , the value of  $b$  is never modified before adding the element to  $\mathcal{S}_{\text{out}}$ .  $\square$

We will also need a way to decide whether  $Q_D^K = P_D^K$ , in the regime where  $\|Q_D^K\|_2$  is considerably smaller than the error we are willing to tolerate. In this regime, it is more sample efficient to simply test whether  $\|P_D^K\|_2$  is sufficiently far below the threshold as well, instead of performing actual equivalence testing.

**Lemma 7.12.** *Consider two unknown distributions  $P_D$  and  $Q_D$  defined over a set  $D$ , and  $K \subseteq D$ . Let  $\eta \in (0, 1)$  be such that  $\|Q_D^K\|_2 \leq \eta/(10\sqrt{|K|})$ . Using  $O(\sqrt{|D|} + \sqrt{|K|}/\eta)$  samples, to give one of the following guarantees which holds with probability at least  $2/3$ :*

- (i)  $D_H^2(P_D^K, Q_D^K) \leq \eta$ ,
- (ii)  $P_D^K \neq Q_D^K$ .

*Proof.* For readability, let  $\tau := \eta/\sqrt{|K|}$ . We will use Lemma 5.5, to learn an estimate  $c_P$  of  $\|P_D^K\|_2$  with error parameter  $\tau/20$ . We then show that

- (i) If  $c_P \leq \frac{\tau}{4}$ , then  $D_H^2(P_D^K, Q_D^K) \leq \sqrt{|K|}\tau$ ,
- (ii) If  $c_P > \frac{\tau}{4}$ , then  $\|P_D^K\|_2 \neq \|Q_D^K\|_2$  (and hence  $P_D^K \neq Q_D^K$ ).

Recall that Lemma 5.5 guarantees that with probability at least  $2/3$ ,  $\|P_D^K\|_2/2 \leq c_P \leq 2\|P_D^K\|_2 + \tau/20$  holds. Assume  $c_P < \tau/4$ , then

$$\tau > \frac{\tau}{10} + 2\frac{\tau}{4} \geq c_Q + 2c_P \geq \frac{1}{\sqrt{|K|}} (\|Q_D^K\|_1 + \|P_D^K\|_1) \geq \frac{1}{\sqrt{|K|}} D_H^2(Q_D^K, P_D^K). \quad (196)$$

Now consider the case when  $c_P > \tau/4$ . Since  $\tau \geq 10c_Q$ , we can say that

$$2\|P_D^K\|_2 + \frac{\tau}{20} > \frac{\tau}{4} = \frac{\tau}{20} + \frac{\tau}{5} \geq \frac{\tau}{20} + 2\|Q_D^K\|_2 \implies \|P_D^K\|_2 > \|Q_D^K\|_2. \quad (197)$$

$\square$

As a corollary of the above lemma, we have the following.

**Corollary 7.13.** *Consider two unknown distributions  $P$  and  $Q$  defined over a set  $[d]$ ,  $S \subseteq [d]$  and  $\varepsilon, \delta, \zeta \in (0, 1)$ . There exists an algorithm `EquivalenceSmall`( $\mathcal{S}_P, S, \zeta, \varepsilon, \delta$ ), which takes a multiset  $\mathcal{S}_P$  of i.i.d. samples from  $P$  with  $\|Q^S\|_2 \leq \zeta$ , and distinguishes with probability at least  $1 - \delta$  whether  $D_H^2(P^S, Q^S) \leq \varepsilon$  or  $P^S \neq Q^S$ , provided  $|\mathcal{S}_P| \geq c_{l2}\sqrt{d}/\eta$  samples, for some instance independent constant  $c_{l2}$ .*



### 7.2.2 Description of the Algorithm

We first state the algorithm testing the large regime, Algorithm 8, which we then explain. The proof of correctness also gives our sample complexity, and we will make use of the tools presented in Section 5.1. There will be three stages. First, we partition  $P_{ABC}^L$  into polylogarithmic many *categories*  $L_{ij}^k$ , illustrated in Figure 5. After simulating a pool of samples from  $Q_{ABC} = P_{AB}P_{C|B}$  using Algorithm 7, we perform equivalence testing individually on each category. Details of the sampling are described in Lemma 7.11, which is built on the fact that  $p_b \geq 1/N_S$  for  $b \in B_L$ . We now define a few quantities we will use in Algorithm 8, which depend on the size of  $L_{ij}^k$ , which are sets we define inside the algorithm, as well as parameters  $k_A := \lceil \log(1/M) \rceil$ ,  $k_B := \lceil \log(1/\nu) \rceil$ ,  $k_C := \lceil \log(1/M) \rceil$ , for  $M$  to be defined, and  $k_{ABC} := \log^3(d_A d_B d_C / \varepsilon)$ . Note that there are  $(k_A + 1)(k_B + 1)(k_C + 1) < k_{ABC}$  categories in total.

$$\gamma(i, j, k) := \begin{cases} \sqrt{\frac{\varepsilon_L e^{k-(i+j+3)}}{e^9 k_{ABC}}} & \text{if } i < k_A, j < k_C, \\ \frac{\varepsilon_L}{4k_{ABC}|L_{ij}^k|^{1/2}} & \text{if either } i = k_A \text{ or } j = k_C, \\ \frac{e^{k+1}\varepsilon_L}{2|B_k|\sqrt{|L_{ij}^k|}} & \text{if } i = k_A, j = k_C, \end{cases} \quad (198)$$

$$b(i, j, k) := e^9 \min \left\{ \sqrt{|L_{ij}^k|} e^{k+4-(i+j)}, \sqrt{e^{k+4-(i+j)}} \right\}, \quad (199)$$

$$\varepsilon_L(i, j, k) := \varepsilon_L / \sqrt{|L_{ij}^k|}. \quad (200)$$

These will follow directly from our derivations. Further, let us define

$$N_L := 10^3 k_{ABC} \log(k_{ABC}) \max\{N_{L,\text{heavy}}, N_{L,\text{mixed}}, N_{L,\text{light}}\}, \quad (201)$$

where

$$N_{L,\text{heavy}} := 10^7 c_{\text{eq}} k_{ABC} \frac{\sqrt{d_A d_B d_C}}{\varepsilon_L}, \quad (\text{Claim 7.14}), \quad (202)$$

$$N_{L,\text{mixed}} := 10^3 c_{\text{eq}} k_{ABC}^2 \min \left\{ \frac{d_A^{3/4} d_B^{3/4} d_C^{1/4}}{\varepsilon_L}, \frac{d_A^{2/3} d_B^{2/3} d_C^{1/3}}{\varepsilon_L^{4/3}} \right\}, \quad (\text{Claim 7.15}), \quad (203)$$

$$N_{L,\text{light}} := 10 \max \left\{ c_{\text{eq}} k_{ABC} \frac{d_A^{1/2} d_B^{3/4} d_C^{1/2}}{\varepsilon_L}, \frac{\log(c_1 k_b)}{\nu} \right\}, \quad (\text{Claim 7.16}). \quad (204)$$

Similar to the independence testing, our overall idea is to partition the distribution such that coordinates of similar weight are grouped together (see Figure 5), and to perform equivalence testing for each of the pieces. To achieve this, we first learn the marginals on  $P_{AB}$ ,  $P_{BC}$ , and  $P_B$  approximately, that is up to a constant factor.

Now we define *categories*  $L_{ij}^k := \cup_{b \in B_k} A_i^b \times b \times C_j^b$ , splitting  $P_{ABC}$  into  $(k_A + 1)(k_B + 1)(k_C + 1) = O(\text{polylog}(d_A d_B d_C / \varepsilon))$  many categories, which we can test individually for precision  $\varepsilon/k_{ABC}$  (pigeonhole principle), at the cost of a polylogarithmic overhead. The partitioning of  $P_{ABC}^L$  is illustrated in Figure 5.

---

**Algorithm 8: CMI Testing Large Regime**

---

**Input:**  $\mathcal{S}$ , multiset of  $N_L$  triplets from  $A \times B \times C$ ,  $\varepsilon_L, \nu \in (0, 1)$   $S_B \subseteq [d_B]$   $\triangleright$  see (201)

**Output:** ‘Yes’ or ‘No’

$\triangleright$  Step 1: Partition  $A \times B \times C$

1  $M \leftarrow \max\{N_{L,\text{heavy}}, N_{L,\text{mixed}}, N_{L,\text{light}}\}$ ,  $\delta \leftarrow \frac{1}{10^3 k_{ABC}}$   $\triangleright$  see (202)-(204)

2  $\mathcal{S}_{AB}, \mathcal{S}_{BC}, \mathcal{S}_B \leftarrow$  remove random disjoint multisets from  $\mathcal{S}$  of size  $8Mk_{ABC}$  each

3  $\forall (a, b) \in A \times B : \hat{p}_{ab} \leftarrow$  empirical frequencies of  $(a, b)$  in  $\mathcal{S}_{AB}$

4  $\forall (b, c) \in B \times C : \hat{p}_{bc} \leftarrow$  empirical frequencies of  $(b, c)$  in  $\mathcal{S}_{BC}$

5  $\forall b \in B : \hat{p}_b \leftarrow$  empirical frequencies of  $b$  in  $\mathcal{S}_B$

6  $\forall i \in [k_A]_0, j \in [k_C]_0, k \in [k_B]_0$ : define  $L_{ij}^k$  according to Figure 5  $\triangleright$  partition  $A \times B \times C$

$\triangleright$  Step 2: Construct Reference Distribution and Simulate Samples

7  $\hat{Q}_{ABC}(a, b, c) \leftarrow \hat{p}_{ab}\hat{p}_{bc}/\hat{p}_b$ ,  $\mathcal{S}_{\text{sim}} \leftarrow$  remove random multiset of size  $N_L/8$  from  $\mathcal{S}$

8  $\mathcal{S}_{CI} \leftarrow \text{SimABC}_{CI}\text{Large}(\mathcal{S}_{\text{sim}})$   $\triangleright$  Algorithm 7, Lemma 7.11,  $|\mathcal{S}_{CI}| = N_L/40$

$\triangleright$  Step 3: Equivalence Testing

9 for all  $i, j, k$  do

10  $\mathcal{S}_P, \mathcal{S}_Q \leftarrow$  remove random multisets from  $\mathcal{S}$  and  $\mathcal{S}_{CI}$ , respectively, of size  $M$  each

11 if  $i < k_A, j < k_C, \|\hat{Q}_{ABC}^{L_{ij}^k}\|_2 < \frac{\varepsilon_L(i, j, k)}{10e^9}$  then

12 if EquivalenceSmall( $\mathcal{S}_P, L_{ij}^k, e^9\|\hat{Q}_{ABC}^{L_{ij}^k}\|_2, \varepsilon_L(i, j, k), \delta$ ) = ‘No’ then

13 return ‘No’  $\triangleright$  Corollary 7.13

14 else

15 if  $i = k_A$  &  $j = k_C$  then

16  $\mathcal{S}_P \leftarrow \{(a, b, c) \in \mathcal{S}_P | b \in B_L\}$ ,  $\mathcal{S}_Q \leftarrow \{(a, b, c) \in \mathcal{S}_Q | b \in B_L\}$   $\triangleright$  Section 7.2.5

17 if EquivalenceL2( $\mathcal{S}_P, \mathcal{S}_Q, L_{ij}^k, b(i, j, k), \gamma(i, j, k), \delta$ ) = ‘Far’ then

18 return ‘No’  $\triangleright$  Corollary 5.7

19 return ‘Yes’

---




Fig. 5: Partition of  $d_A \times d_B \times d_C$  based on  $\hat{P}_{AB}$ ,  $\hat{P}_{BC}$ , and  $\hat{P}_B$ . Indices  $(a, b, c)$  of similar weight  $\hat{p}_{ab}\hat{p}_{bc}/\hat{p}_b$  are grouped together in categories  $L_{ij}^k$ , which are used to perform piecewise equivalence testing with  $P_{ABC}$ . The axes are labeled according to the corresponding category, which is inverse logarithmic to the weight of the probabilities. The color of the categories indicate a different analysis of the sample complexity of testing the categories. The red and orange regimes dominate the sample complexity. The small regime is treated separately.

To perform these equivalence tests, we will need samples from our reference distribution  $Q_{ABC} = P_{AB}P_{C|B}$ . We will sample from a slightly different distribution,  $\tilde{Q}_{ABC}$ , which satisfies  $\tilde{Q}_{ABC}^L = (P_{AB}P_{C|B})^L$  and  $\tilde{Q}_B = Q_B$ . Since we are only interested in ensuring equivalence on  $L$  (or subsets of  $L$ ), we can treat  $Q$  and  $\tilde{Q}$  as equivalent. Our method of sampling is described in Lemma 7.11.

Next, we perform equivalence testing on the individual categories. In total, we have three different types of categories. We will bound the sample complexity of testing these three categories individually. Then, we will combine the results for each category to obtain our final result. For heavy categories where  $\|Q_{ABC}^{L_{ij}^k}\|_2$  is very small, equivalence testing could become very costly, as will become evident from the proof of Claim 7.14. We get around this by noting that if  $\|P_{ABC}^{L_{ij}^k}\|_2$  and  $\|Q_{ABC}^{L_{ij}^k}\|_2$  are small enough, then certainly  $\|P_{ABC}^{L_{ij}^k} - Q_{ABC}^{L_{ij}^k}\|_2 < \varepsilon_L/k_{ABC}$ , without explicitly testing for equivalence. The formal argument is presented in Lemma 7.12, and explains the case distinction in Algorithm 8.

1. **Heavy categories:** This includes all categories  $L_{ij}^k$  where  $i < k_A$  and  $j < k_C$ , such that for all  $a, b, c$ , both  $p_{ab}$  and  $p_{bc}$  are bounded away from zero. Note that some categories will be excluded from the testing, as mentioned above. This also provides us with a lower bound on  $\|Q_{ABC}^{L_{ij}^k}\|_2$ ,  $\|Q_{ABC}^{L_{ij}^k}\|_2 \geq \varepsilon_L(i, j, k)/(10e^{18})$ .

**Claim 7.14.** For any heavy category  $L_{ij}^k$  as defined in Figure 5 with  $\|Q_{ABC}^T\|_2 \geq \Omega(\varepsilon_L(i, j, k))$ ,

distinguishing whether  $P_{ABC}^{L_{ij}^k} = Q_{ABC}^{L_{ij}^k}$  or  $D_H^2(P_{ABC}^{L_{ij}^k}, Q_{ABC}^{L_{ij}^k}) \geq \varepsilon_L/k_{ABC}$  can be done using

$$N_{L,\text{heavy}} := 10^7 c_{eq} k_{ABC} \frac{\sqrt{d_A d_B d_C}}{\varepsilon_L}, \quad (205)$$

samples from  $P_{ABC}$  and  $Q_{ABC}$  each with probability of success at least 99/100.

2. **Mixed categories:** This includes all categories  $L_{ij}^k$  where either  $i = k_A$  or  $j = k_C$ . In the discussion, without loss of generality, we will assume that  $a$  is coming from the last bin.

**Claim 7.15.** For any mixed category  $L_{ij}^k$  as defined in Figure 5, distinguishing whether  $P_{ABC}^{L_{ij}^k} = Q_{ABC}^{L_{ij}^k}$  or  $D_H^2(P_{ABC}^{L_{ij}^k}, Q_{ABC}^{L_{ij}^k}) \geq \varepsilon_L/k_{ABC}$  can be done using

$$N_{L,\text{mixed}} := 10^3 c_{eq} k_{ABC}^2 \min \left\{ \frac{d_A^{3/4} d_B^{3/4} d_C^{1/4}}{\varepsilon_L}, \frac{d_A^{2/3} d_B^{2/3} d_C^{1/3}}{\varepsilon_L^{4/3}} \right\} \quad (206)$$

samples from  $P_{ABC}$  and  $Q_{ABC}$  each (where  $d_A \geq d_C$ ), with probability of success at least 99/100.

3. **Light categories:** This includes all categories  $L_{ij}^k$  where both  $i = k_A$  and  $j = k_C$ .

**Claim 7.16.** For any light category  $L_{ij}^k$  as defined in Figure 5, distinguishing whether  $P_{ABC}^{L_{ij}^k} = Q_{ABC}^{L_{ij}^k}$  or  $D_H^2(P_{ABC}^{L_{ij}^k}, Q_{ABC}^{L_{ij}^k}) \geq \varepsilon_L/k_{ABC}$  can be done using

$$N_{L,\text{light}} := 10 \max \left\{ c_{eq} k_{ABC} \frac{d_A^{1/2} d_B^{3/4} d_C^{1/2}}{\varepsilon_L}, \frac{\log(10^3 k_{ABC})}{\nu} \right\} \quad (207)$$

samples from  $P_{ABC}$  and  $Q_{ABC}$  with probability of success at least 99/100.

We will prove the above three claims in the following. Together, they directly imply Lemma 7.10.

**Lemma 7.17.** With probability at least 2/3, Algorithm 8 correctly performs conditional independence testing in  $D_H^2$  distance in the large regime with the sample complexity specified in Lemma 7.10.

*Proof.* Analogous to the proof of Lemma 5.12, we first argue that our partitioning satisfies certain properties. With probability at least 9/10, it holds that

$$\forall L_{ij}^k, \forall (a, b, c) \in L_{ij}^k : \begin{cases} e^{-(i+2)} \mathbb{1}[i < k_A] \leq q_{ab} \leq e^{-i+1} \\ e^{-(j+2)} \mathbb{1}[j < k_C] \leq q_{bc} \leq e^{-j+1} \\ e^{-(k+2)} \leq q_b \leq e^{-k+1}. \end{cases} \quad (208)$$

The guarantee follows directly from Lemma 5.8 (note that  $8 \log(4d_A d_C / \zeta) / (8M k_{ABC}) \geq 1/(2M)$ , so any element whose probability mass is less than  $1/(2M)$  will not have empirical frequency above  $1/M$ ) and a union bound, analogously for each of  $q_{ab}$ ,  $q_{bc}$  and  $q_b$ . In particular, it holds with high probability that

$$\forall L_{ij}^k, \forall (a, b, c), (a', b', c') \in L_{ij}^k : \begin{cases} \frac{p_{ab}}{p_{a'b'}} \leq e^3 \text{ if } i < k_A, & \text{and } p_{ab} \leq e^{-k_A+1} \text{ otherwise,} \\ \frac{p_{bc}}{p_{b'c'}} \leq e^3 \text{ if } j < k_C, & \text{and } p_{bc} \leq e^{-k_C+1} \text{ otherwise,} \\ \frac{p_b}{p_{b'}} \leq e^3. \end{cases} \quad (209)$$

We will use these inequalities in the following subsections.

To summarize, this implies that no ‘small’  $p_{ab}, p_{bc} \leq 1/(2M)$  is placed in one of the first  $k_A$  buckets, while ‘large’  $p_{ab}$  and  $p_{bc}$  are estimated up to a factor of two, and hence are assigned to either the correct or a neighboring bin.

Within a category  $L_{ij}^k$  for which  $i < k_A$  and  $j < k_C$ , we then have

$$\forall (a, b, c), (a', b', c') \in L_{ij}^k : \frac{p_{ab}p_{bc}/p_b}{p_{a'b'}p_{b'c'}/p_{b'}} \leq e^9. \quad (210)$$

For the following, let

$$s_{ij}^k := \sum_{(a,b,c) \in L_{ij}^k} \left( \sqrt{p_{abc}} - \sqrt{\frac{p_{ab}p_{bc}}{p_b}} \right)^2, \quad \text{such that} \quad D_H^2(P_{ABC}, Q_{ABC}) = \sum_{k,i,j} s_{ij}^k. \quad (211)$$

This implies that if  $D_H^2(P_{ABC}, Q_{ABC}) \geq \varepsilon_L$ , then by the pigeonhole principle, at least one  $s_{ij}^k$  will satisfy  $s_{ij}^k \geq \varepsilon_L/k_{ABC}$ , which we will test this individually for each  $L_{ij}^k$ .

The subroutines testing the individual categories will each succeed with probability at least 99/100. To ensure an *overall* high probability of success, we repeat each test  $6 \log(10^3 k_{ABC})$  times and take a majority vote. The correctness of majority voting follows directly from the Chernoff bound, by choosing  $\delta$  such that  $(1 - \delta)p = 1/2$ , we find the probability that less than 50% of the outcomes indicate the right result is bounded by  $\exp(-Np(1 - 1/2p)^2/2)$ . We have  $k_{ABC}$  instances in total, so performing each one (assuming  $p \geq 99/100$ )

$$\frac{\log(10^3 k_{ABC})}{p(1 - p/2)^2/2} < 6 \log(10^3 k_{ABC}) \quad (212)$$

times ensure an overall probability of success of at least 99/100. We showed that the approximate learning is precise enough to make the distinction into bins correctly with probability at least 99/100, sampling from  $\tilde{Q}_{ABC}$  succeeds with probability at least 99/100 as well, and testing on the individual categories is successful with probability 99/100 too. This implies a probability of success of at least 2/3. The contributions to the sample complexity are

- (i) The sample complexity of Step 1 in Algorithm 8 can be bounded by

$$N_1 \geq 24 \log(c_1 d_A d_B d_C) \max \{N_{L,\text{heavy}}, N_{L,\text{mixed}}, N_{L,\text{light}}\} \quad (213)$$

- (ii) The sample complexity of Step 2 in Algorithm 8 can be upper bounded by

$$N_2 \geq 5 \cdot 6 \log(c_1 k_{ABC}) k_{ABC} \max \left\{ \frac{4}{\varepsilon_L(i, j, k)}, N_{L,\text{heavy}}, N_{L,\text{mixed}}, N_{L,\text{light}} \right\}. \quad (214)$$

The additional factor in the sample complexity is due to simulating samples, as mentioned in Lemma 7.11.

Inserting the sample complexities, we can bound

$$N_1 + N_2 \quad (215)$$

$$\geq c_{\text{eq}} 10^5 k_{ABC}^3 \log(c_1 k_{ABC}) \max \left\{ \min \left\{ \frac{d_A^{3/4} d_B^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_B^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\}, \frac{d_A^{1/2} d_B^{3/4} d_C^{1/2}}{\varepsilon} \right\}. \quad (216)$$

To bound the probability of success, we apply a union bound, over all tasks we perform: we learn  $P_{AB}$ ,  $P_{BC}$  and  $P_B$ , create samples and test the different regimes. The testing is already boosted using majority voting, such that each of the five tasks individually succeeds with probability at least 99/100, giving an overall probability of success at least 9/10.  $\square$

### 7.2.3 Sample Complexity of Heavy Categories

**Claim 7.14.** For any heavy category  $L_{ij}^k$  as defined in Figure 5 with  $\|Q_{ABC}^T\|_2 \geq \Omega(\varepsilon_L(i, j, k))$ , distinguishing whether  $P_{ABC}^{L_{ij}^k} = Q_{ABC}^{L_{ij}^k}$  or  $D_H^2(P_{ABC}^{L_{ij}^k}, Q_{ABC}^{L_{ij}^k}) \geq \varepsilon_L/k_{ABC}$  can be done using

$$N_{L,\text{heavy}} := 10^7 c_{eq} k_{ABC} \frac{\sqrt{d_A d_B d_C}}{\varepsilon_L}, \quad (205)$$

samples from  $P_{ABC}$  and  $Q_{ABC}$  each with probability of success at least 99/100.

*Proof.* Recall that in this case, for all  $a, b, c$ , both  $p_{ab}$  and  $p_{bc}$  are bounded away from zero, and there are  $k_A k_B k_C < k_{ABC}$  many such  $s_{i,j}^k$  for which this is the case. For ease of presentation, we will denote  $L_{ij}^k$  as  $T$  throughout this proof.

We can further assume that all  $p_{ab}$  differ pairwise by at most a factor  $e^3$ , and by the case distinction, that  $\|Q_{ABC}^T\|_2 \geq \varepsilon_L/(8k_{ABC})$ . Consider an arbitrary  $s_{i,j}^k$ . From (211), we have the following:

$$s_{i,j}^k = \sum_{(a,b,c) \in T} \left( \sqrt{p_{abc}} - \sqrt{\frac{p_{ab}p_{bc}}{p_b}} \right)^2 \quad (217)$$

$$= \sum_{(a,b,c) \in T} \frac{\left( p_{abc} - \frac{p_{ab}p_{bc}}{p_b} \right)^2}{\left( \sqrt{p_{abc}} + \sqrt{\frac{p_{ab}p_{bc}}{p_b}} \right)^2} \quad (218)$$

$$\leq \frac{e^9}{\max_{\substack{b \in B_k \\ (a,c) \in S_{ij}^b}} \left\{ \frac{p_{ab}p_{bc}}{p_b} \right\}} \|P_{ABC}^T - Q_{ABC}^T\|_2^2. \quad (219)$$

Thus, following (219), distinguishing whether  $s_{i,j}^k = 0$  or  $s_{i,j}^k \geq \varepsilon_L/k_{ABC}$  is reduced to testing whether:

$$P_{ABC}^T = Q_{ABC}^T \quad \text{or} \quad \|P_{ABC}^T - Q_{ABC}^T\|_2 \geq \eta, \quad (220)$$

where

$$\eta := \sqrt{\frac{\varepsilon_L}{e^9 k_{ABC}} e^{k-(i+j+3)}} \leq \sqrt{\frac{\varepsilon_L}{e^9 k_{ABC}} \max_{(a,b,c) \in T} \left\{ \frac{p_{ab}p_{bc}}{p_b} \right\}}. \quad (221)$$

From Lemma 5.6, we know that equivalence testing with respect to  $\ell_2$  distance with proximity parameter  $\eta$  on a subset  $T$  requires  $O(\|Q_{ABC}^T\|_2/\eta^2 + \sqrt{|T|} + 1/\eta)$  samples. We now analyze the first term,  $\|Q_{ABC}^T\|_2/\eta^2$ . Note that

$$\|Q_{ABC}^T\|_2 = \sqrt{\sum_{(a,b,c) \in T} \left( \frac{p_{ab}p_{bc}}{p_b} \right)^2} \leq \sqrt{|T|} \max_{(a,b,c) \in T} \left\{ \frac{p_{ab}p_{bc}}{p_b} \right\}. \quad (222)$$

Hence, we have that

$$\frac{\|Q_{ABC}^T\|_2}{\eta^2} \leq \frac{e^9 \sqrt{|T|} k_{ABC}}{\varepsilon_L}. \quad (223)$$

This will clearly dominate the term  $\sqrt{|T|}$ , and it remains to consider the term  $1/\eta$ , which is dominant whenever

$$\frac{\|Q_{ABC}^T\|_2}{\eta^2} \leq \frac{1}{\eta} \quad \Leftrightarrow \quad \|Q_{ABC}^T\|_2 \leq \eta. \quad (224)$$

However, since we have a lower bound on  $\|Q_{ABC}^T\|_2$ ,  $\|Q_{ABC}^T\|_2 \geq \varepsilon_L(i, j, k)/(10e^{18})$ , we can easily upper bound  $1/\eta$  by  $1/\eta \leq 10e^{18}/\varepsilon_L(i, j, k)$ , which is less than (223). In order to use Lemma 5.6, we need individual bounds on  $\|Q_{ABC}^T\|_2$  and  $\gamma(i, j, k)$ , which we can bound as follows using the definitions of  $L_{ij}^k$  in Figure 5 and (209)

$$b(i, j, k) := e^9 \min \left\{ \sqrt{|T|} e^{k+4-(i+j)}, \sqrt{e^{k+4-(i+j)}} \right\} \geq \|Q_{ABC}^T\|_2, \quad \gamma(i, j, k) := \eta. \quad (225)$$

Because of these bounds, we get an additional factor of up to  $e^6$  compared to (223).

Due to the additional logarithmic factors, (223) dominates  $1/\eta$ , and the large regime is thus bounded by

$$c_{\text{eq}} \frac{e^{15} \sqrt{|T|} k_{ABC}}{\varepsilon_L} \leq c_{\text{eq}} e^{15} \frac{\sqrt{d_A d_B d_C} k_{ABC}}{\varepsilon_L}. \quad (226)$$

□

## 7.2.4 Sample Complexity of Mixed Categories

**Claim 7.15.** *For any mixed category  $L_{ij}^k$  as defined in Figure 5, distinguishing whether  $P_{ABC}^{L_{ij}^k} = Q_{ABC}^{L_{ij}^k}$  or  $D_H^2(P_{ABC}^{L_{ij}^k}, Q_{ABC}^{L_{ij}^k}) \geq \varepsilon_L/k_{ABC}$  can be done using*

$$N_{L, \text{mixed}} := 10^3 c_{\text{eq}} k_{ABC}^2 \min \left\{ \frac{d_A^{3/4} d_B^{3/4} d_C^{1/4}}{\varepsilon_L}, \frac{d_A^{2/3} d_B^{2/3} d_C^{1/3}}{\varepsilon_L^{4/3}} \right\} \quad (206)$$

*samples from  $P_{ABC}$  and  $Q_{ABC}$  each (where  $d_A \geq d_C$ ), with probability of success at least 99/100.*

*Proof.* For brevity, we will again write  $T$  instead of  $L_{ij}^k$  in the following proof. We also define  $|T_C|$  to be the maximal number of different  $c$ 's appearing in  $T$ ,  $T_C := \{c | \exists a, b : (a, b, c) \in T\}$ . We define  $T_A$  and  $T_B$  analogously. Our proof follows in a similar line as Claim 5.11. We would again like to upper bound  $D_H^2(P_{ABC}^T, Q_{ABC}^T)$  by  $\|P_{ABC}^T - Q_{ABC}^T\|_2$ . However, we can no longer use a similar approach as in (219) as there is no lower bound on  $p_{ab}$ . Let us assume without loss of generality that the unbounded coordinate comes from  $A$ , i.e.  $p_{ab} \leq e/M$ . We denote by  $T_-$  the set of indices in  $T$  for which  $(\sqrt{p_{abc}} - \sqrt{p_{ab}p_{bc}/p_b})^2 \leq x$ , for some  $x \in (0, 1)$  to be determined, and  $T_+ := T \setminus T_-$ . Then

$$D_H^2(P_{ABC}^T, Q_{ABC}^T) \quad (227)$$

$$= \frac{1}{2} \left( \sum_{s \in T_+} \left( \sqrt{P_{ABC}^T} - \sqrt{Q_{ABC}^T} \right)^2 + \sum_{s \in T_-} \left( \sqrt{P_{ABC}^T} - \sqrt{Q_{ABC}^T} \right)^2 \right) \quad (228)$$

$$= \sum_{(a,b,c) \in T_+} \left( \sqrt{p_{abc}} - \sqrt{\frac{p_{ab}p_{bc}}{p_b}} \right)^2 + \sum_{(a,b,c) \in T_-} \left( \sqrt{p_{abc}} - \sqrt{\frac{p_{ab}p_{bc}}{p_b}} \right)^2 \quad (229)$$

$$\leq \sum_{(a,b,c) \in T_+} \frac{1}{x} \left( p_{abc} - \frac{p_{ab}p_{bc}}{p_b} \right)^2 + |(T)_-| x \quad (230)$$

$$\leq \frac{1}{x} \|P_{ABC}^T - Q_{ABC}^T\|_2^2 + |T| x. \quad (231)$$

Since we want to test if  $D_H^2(P_{ABC}^T, Q_{ABC}^T) \geq \varepsilon_L/(2k_{ABC})$ , from (231), we have the following:

$$D_H^2(P_{ABC}^T, Q_{ABC}^T) \geq \frac{\varepsilon_L}{2k_{ABC}} \quad (232)$$

$$\implies \|P_{ABC}^T - Q_{ABC}^T\|_2 \geq \sqrt{\frac{x\varepsilon_L}{2k_{ABC}} - |T|x^2} \quad (233)$$

In order for the first term in the square root to dominate the second, we choose  $x = \varepsilon_L/(4k_{ABC}|T|)$ , and arrive at the following testing problem:

$$\|P_{ABC}^T - Q_{ABC}^T\|_2 \geq \tilde{\varepsilon}_L := \frac{\varepsilon_L}{4k_{ABC}|T|^{1/2}} \quad \text{or} \quad P_{ABC}^T = Q_{ABC}^T. \quad (234)$$

Similar to the proof of the previous claim, we now need to bound  $\|Q_{ABC}^T\|_2$ . Here we will use Fact 3.3, which allows us to bound  $\|Q_{ABC}^T\|_2$  in the following two ways:

$$\|Q_{ABC}^T\|_2 = \sqrt{\sum_{(a,b,c) \in T} \left(\frac{p_{ab}p_{bc}}{p_b}\right)^2} \leq \begin{cases} \sqrt{|T|} \max_{(a,b,c) \in T} \left\{\frac{p_{ab}p_{bc}}{p_b}\right\} & \text{(Case 1)} \\ \sqrt{\max_{(a,b,c) \in T} \left\{\frac{p_{ab}p_{bc}}{p_b}\right\}} & \text{(Case 2)} \end{cases} \quad (235)$$

From the different terms present in the sample complexity of equivalence testing (see Lemma 5.6), we only consider  $\|Q_{ABC}^T\|_2/\tilde{\varepsilon}_L^2$  (we will find that this term dominates over  $1/\tilde{\varepsilon}_L$  and  $\sqrt{|T|}$ ). Depending on which case is more favorable, we obtain:

1. **Case 1:** From (234) and (235), we have:

$$\frac{\|Q_{ABC}^T\|_2}{\tilde{\varepsilon}_L^2} \leq \frac{\sqrt{|T|} \max_{(a,b,c) \in T} \left\{\frac{p_{ab}p_{bc}}{p_b}\right\}}{\tilde{\varepsilon}_L^2} \leq 16k_{ABC}^2 \frac{|T|^{3/2}}{\varepsilon_L^2} \max_{(a,b,c) \in T} \left\{\frac{p_{ab}p_{bc}}{p_b}\right\}. \quad (236)$$

Note that  $\forall (a, b, c) \in T : p_{bc}/p_b \leq e^3/|T_C|$  (since all  $p_{bc}$  for a fixed  $b$  from  $(a, b, c) \in T$  differ by no more than a factor  $e^3$ , and their sum is bounded by  $p_b$ ). Also, by assumption on  $A$ ,  $p_{ab} \leq e^{-(k_A+1)} = e/M$ , and  $|T| \leq |T_A||T_B||T_C|$ , so

$$\frac{\|Q_{ABC}^T\|_2}{\tilde{\varepsilon}_L^2} \leq 4^2 k_{ABC}^2 \frac{e^6(|T_A||T_B|)^{3/2}|T_C|^{1/2}}{\varepsilon_L^2 M} \quad (237)$$

The number of samples,  $N_{L,\text{mixed}}$ , we need to upper bound this. For this, assume that we are in a regime where  $N_{L,\text{mixed}} \geq \max\{N_{L,\text{heavy}}, N_{L,\text{light}}\}$ , such that  $M = N_{L,\text{mixed}}$ . Then

$$4^2 k_{ABC}^2 \frac{e^6(|T_A||T_B|)^{3/2}|T_C|^{1/2}}{\varepsilon_L^2 N_{L,\text{mixed}}} \leq N_{L,\text{mixed}} \implies N_{L,\text{mixed}} \geq \frac{(d_A d_B)^{3/4} d_C^{1/4}}{\varepsilon_L} \cdot 4e^3 k_{ABC} \quad (238)$$

samples are sufficient.

2. **Case 2:** Similarly, using (234) and (235), we can say the following:

$$\frac{\|Q_{ABC}^T\|_2}{\tilde{\varepsilon}_L^2} \leq \frac{\sqrt{\max_{(a,b,c) \in T} \left\{\frac{p_{ab}p_{bc}}{p_b}\right\}}}{\tilde{\varepsilon}_L^2} \leq \sqrt{\frac{e^6}{N_{L,\text{mixed}}|T_C|} \frac{16k_{ABC}^2|T|}{\varepsilon_L^2}} \stackrel{!}{\leq} N_{L,\text{mixed}}, \quad (239)$$



such that we find that

$$N_{L,\text{mixed}} \geq \frac{(d_A d_B)^{2/3} d_C^{1/3}}{\varepsilon_L^{4/3}} \cdot 4^2 e^2 k_{ABC}^2 \quad (240)$$

samples are sufficient as well.

Note that an analogous argument holds for the case where the roles of  $A$  and  $C$  are switched. We explicitly know  $\tilde{\varepsilon}_L$ , and bounded  $\|Q_{ABC}^T\|_2 \leq \min\{\sqrt{|T|}e^{2-k_A}, e^{1-k_A/2}\}$ . Since we can always choose which of the two cases we want to use, the proof of the claim is complete. Also note that the contribution of always dominates the contribution by  $1/\tilde{\varepsilon}_L$  and  $\sqrt{|T|}$ .  $\square$

### 7.2.5 Sample Complexity of Light Categories

In the following we will bound the sample complexity for testing a light category  $T$ . Instead of using Lemma 5.6 for equivalence testing between  $P_{ABC}^T$  and  $Q_{ABC}^T$  directly, we take an intermediate step. Let  $R(T) := \{(a, b, c) | b \in T_B\}$  in the following. First, we take samples from  $P_{ABC}$  and  $Q_{ABC}$  and immediately reject those with  $b \notin T_B$ . This allows us to condition on  $b$  coming from  $T_B$ , and the remaining samples can be seen as coming from the (normalized) distributions  $P_{ABC}^{R(T)}/\ell$  and  $Q_{ABC}^{R(T)}/\ell$ , for  $\ell := \sum_{i \in R(T)} p_i \leq 1$ . We can then apply Lemma 5.6 to test for equivalence between  $P_{ABC}^T/\ell$  and  $Q_{ABC}^T/\ell$ , which we see as ‘subdistributions’ of  $P_{ABC}^{R(T)}/\ell$  and  $Q_{ABC}^{R(T)}/\ell$ .

In our calculation, this results in a tighter bound than using Lemma 5.6 directly to perform the testing in a black box manner. The improvement comes into play in (253), where we use the normalization factor  $\ell$  to bound a term of the form  $\ell/p_b$ . Using Lemma 5.6 without the conditioning would result in a worse factor of  $1/p_b$  instead. Note that  $1/p_b$  does not need to be bounded in the same way in the mixed and heavy categories, such that the approach we use here for the light regime would not provide any advantage in the other regimes.

**Claim 7.16.** *For any light category  $L_{ij}^k$  as defined in Figure 5, distinguishing whether  $P_{ABC}^{L_{ij}^k} = Q_{ABC}^{L_{ij}^k}$  or  $D_H^2(P_{ABC}^{L_{ij}^k}, Q_{ABC}^{L_{ij}^k}) \geq \varepsilon_L/k_{ABC}$  can be done using*

$$N_{L,\text{light}} := 10 \max \left\{ c_{eq} k_{ABC} \frac{d_A^{1/2} d_B^{3/4} d_C^{1/2}}{\varepsilon_L}, \frac{\log(10^3 k_{ABC})}{\nu} \right\} \quad (207)$$

samples from  $P_{ABC}$  and  $Q_{ABC}$  with probability of success at least 99/100.

*Proof.* In this proof, our approach follows a similar line as that of the analysis of mixed categories. First, as mentioned previously, we define  $R(T) := \{(a, b, c) | b \in T_B\}$  and  $\ell := \sum_{i \in R(T)} p_i$ . Since both  $a$  and  $c$  are light here, we can no longer use the fact that  $p_{bc}/p_b \leq O(1/|T_C|)$ . Instead, we use the fact that  $p_{ab}, p_{bc} \leq 1/N_L$ . Note that we define  $T_+$  and  $T_-$  slightly differently in this case. We denote by  $T_-$  the set of indices in  $T$  for which  $(\sqrt{p_{abc}} - \sqrt{p_{ab}p_{bc}/p_b})^2/\ell \leq y$  holds, for some  $y \in (0, 1)$  to be determined. Further, let  $T_+ := T \setminus T_-$ . Note that

$$\left\| \frac{P_{ABC}^{R(T)}}{\ell} \right\|_1 = \sum_{(a,b,c) \in R(T)} \frac{p_{abc}}{\ell} = 1, \quad \left\| \frac{Q_{ABC}^{R(T)}}{\ell} \right\|_1 = \sum_{(a,b,c) \in R(T)} \frac{p_{ab}p_{bc}}{p_b \ell} = 1. \quad (241)$$

We can now write

$$D_H^2(P_{ABC}^T, Q_{ABC}^T) \quad (242)$$

$$= \frac{1}{2} \left( \sum_{t \in T_+} (\sqrt{P_{ABC}[t]} - \sqrt{Q_{ABC}[t]})^2 + \sum_{t \in T_-} (\sqrt{P_{ABC}[t]} - \sqrt{Q_{ABC}[t]})^2 \right) \quad (243)$$

$$= \sum_{(a,b,c) \in T_+} \left( \sqrt{p_{abc}} - \sqrt{\frac{p_{ab}p_{bc}}{p_b}} \right)^2 + \sum_{(a,b,c) \in T_-} \left( \sqrt{p_{abc}} - \sqrt{\frac{p_{ab}p_{bc}}{p_b}} \right)^2 \quad (244)$$

$$\leq \ell \left[ \sum_{(a,b,c) \in T_+} \left( \sqrt{\frac{p_{abc}}{\ell}} - \sqrt{\frac{p_{ab}p_{bc}}{p_b \ell}} \right)^2 + \sum_{(a,b,c) \in T_-} \left( \sqrt{\frac{p_{abc}}{\ell}} - \sqrt{\frac{p_{ab}p_{bc}}{p_b \ell}} \right)^2 \right] \quad (245)$$

$$\leq \ell \left[ \sum_{(a,b,c) \in T_+} \frac{1}{y} \left( \frac{p_{abc}}{\ell} - \frac{p_{ab}p_{bc}}{p_b \ell} \right)^2 + |T_-|y \right] \quad (246)$$

$$\leq \ell \left[ \frac{1}{y} \left\| \frac{P_{ABC}^T}{\ell} - \frac{Q_{ABC}^T}{\ell} \right\|_2^2 + |T|y \right]. \quad (247)$$

If we take  $N_{L,\text{light}}$  samples from both  $P_{ABC}^{R(T)}$ , in expectation we receive  $\ell N_{L,\text{light}}$  samples from  $P_{ABC}^{R(T)}$ , and analogously for  $Q_{ABC}^{R(T)}$ . To ensure that we receive no less than  $\ell N_{L,\text{light}}/2$  samples with high probability, we impose the requirement that  $\ell N_{L,\text{light}} \geq 8 \log(10^3 k_{ABC}) > 8 \log(10^3 k_B)$ . Then, Lemma 5.8 ensures that with probability no less than  $1 - 2/(10^3 k_{ABC})$ , we receive at least  $\ell N_{L,\text{light}}/2$  samples from both  $P_{ABC}^{R(T)}$  and  $Q_{ABC}^{R(T)}$ .

As mentioned previously, the remaining samples can be seen as coming directly from  $R(T)$ , and we can apply Lemma 5.6 to test the equivalence between  $P_{ABC}^{R(T)}/\ell$  and  $Q_{ABC}^{R(T)}/\ell$  on the subset  $T$ . We now determine the precision to which we perform the equivalence testing,

$$\frac{\varepsilon_L}{k_{ABC}} \leq D_H^2(P_{ABC}^T, Q_{ABC}^T) \leq \ell \left[ \frac{1}{y} \left\| \frac{P_{ABC}^T}{\ell} - \frac{Q_{ABC}^T}{\ell} \right\|_2^2 + |T|y \right] \quad (248)$$

$$\implies \left\| \frac{P_{ABC}^T}{\ell} - \frac{Q_{ABC}^T}{\ell} \right\|_2 \geq \sqrt{\frac{\varepsilon_L y}{k_{ABC} \ell} - |T|y^2} =: \tilde{\varepsilon}_L. \quad (249)$$

We choose  $y = \varepsilon/(2|S|\ell k_{ABC})$ , resulting in

$$\tilde{\varepsilon}_L = \frac{\varepsilon_L}{2\ell k_{ABC} \sqrt{|T|}} \geq \frac{e\varepsilon_L}{2k_{ABC} \max_{b \in T_k} \{p_b\} |B_k| \sqrt{|T|}}. \quad (250)$$

In the following, note that  $p_{ab}, p_{bc} \leq e/M$  each. We now use Case 2 of Fact 3.3,

$$\|Q_{ABC}^T/\ell\|_2 = \sqrt{\sum_{a,c \in T_{AC}} \left( \frac{p_{ab}p_{bc}}{p_b \ell} \right)^2} \leq \sqrt{\max \left\{ \frac{p_{ab}p_{bc}}{p_b \ell} \right\}} \leq \sqrt{\max \left\{ \frac{1}{p_b \ell} \right\}} \frac{e}{M}. \quad (251)$$

The sample complexity for equivalence testing using samples from  $Q_{ABC}^{R(T)}$  to precision  $\tilde{\varepsilon}_L$  in  $\ell_2$  is  $O(\max\{\|Q_{ABC}^T/\ell\|_2/\tilde{\varepsilon}_L^2, 1/\tilde{\varepsilon}_L, \sqrt{d_A|T_B|d_C}\})$ , according to Lemma 5.6. We will see that the first term dominates, which can be bounded by

$$c_{\text{eq}} \frac{\|Q_{ABC}^T/\ell\|_2}{\tilde{\varepsilon}_L^2} \leq c_{\text{eq}} \frac{|T|\ell^2 k_{ABC}^2 e}{\sqrt{\min_{b \in T_B} \{p_b\}} \ell \varepsilon^2 M}. \quad (252)$$

This number of samples needs to be upper bounded by the number of samples we see from  $Q_{ABC}^{R(T)}$ , which we argued above to be at least  $\ell N_{L,\text{light}}/2$ , with high probability.

Assuming again the worst case, where  $M = N_{L,\text{light}}$ , such that this case dominates the sample complexity, results in the requirement

$$N_{L,\text{light}} \geq 2e\sqrt{c_{\text{eq}}}\frac{\ell^{1/4}|T|^{1/2}k_{ABC}}{\varepsilon_L(\min_{b \in S_B}\{p_b\})^{1/4}}, \quad (253)$$

which we can ensure by choosing  $N_{L,\text{light}}$  such that

$$N_{L,\text{light}} \geq \frac{|T_B|^{1/4}|T|^{1/2}}{\varepsilon_L} 10\sqrt{c_{\text{eq}}}k_{ABC}. \quad (254)$$

We also required  $\ell N_{L,\text{light}} \geq 8 \log(10^3 k_B)$ . Using  $\ell \geq \min_{b \in B_k}\{p_b\}|T_B| \geq \nu|T_B|$ , such that we find the bound

$$N_{L,\text{light}} \geq 10 \max \left\{ \frac{|T_B|^{1/4}|T|^{1/2}}{\varepsilon_L} \sqrt{c_{\text{eq}}}k_{ABC}, \frac{\log(c_1 k_b)}{\nu|T_B|} \right\}. \quad (255)$$

□

**Remark 7.18.** *In the above calculations, there seem to be alternative options to bound certain terms. Here we will briefly argue why they do not lead to an improvement over our sample complexity reported in Lemma 7.10. First, in (251), we chose to use Case 2 of Fact 3.3. What happens if we use Case 1 instead? We then obtain*

$$\frac{\|Q_{ABC}^T\|_2}{\tilde{\varepsilon}^2} \leq \frac{\sqrt{|T|} \max_{(a,b,c) \in T} \left\{ \frac{p_{ab}p_{bc}}{p_b \ell} \right\}}{\tilde{\varepsilon}^2} \leq \frac{\ell|T|^{3/2}}{p_b \varepsilon^2 N_L^2} \stackrel{!}{\leq} \ell N_L \implies N_L \geq \frac{|T|^{1/2}}{p_b^{1/3} \varepsilon^{2/3}}. \quad (256)$$

We are working in the regime of large  $p_b$ , and we know from Section 7.1 that  $p_b \geq \max\{\varepsilon/d_B^{7/8}, \varepsilon^{8/7}/d_B^{6/7}\}$  (note that including the scaling in  $d_A$  and  $d_C$  would only decrease  $p_b$ ). Thus we obtain  $N_L \geq \max\{|T|^{1/2}d_B^{7/24}/\varepsilon, |T|^{1/2}d_B^{6/21}/\varepsilon^{22/21}\}$ , which is larger than the sample complexity from Case 2,  $|T|^{1/2}d_B^{1/4}/\varepsilon$ , in all settings.

It might also seem that in (253), instead of using  $\ell = \Theta(p_b|T_B|)$  to cancel  $p_b$ , we might try to bound  $\ell$  by 1, and bound  $p_b$  directly by the minimal value of  $p_b$  in the large regime, as guaranteed by Section 7.1. Instead of  $d_B^{1/4}|T|^{1/2}/\varepsilon$ , we would get  $x^{1/4}|T|^{1/2}/\varepsilon$ , where

$$x := \min \left\{ d_B, d_B^{6/7} (d_A d_C)^{2/7} / \varepsilon^{8/7}, d_B^{7/8} (d_A d_C)^{1/4} / \varepsilon \right\}. \quad (257)$$

Suppose  $x_{\min} := d_B^{6/7} (d_A d_C)^{2/7} / \varepsilon^{8/7}$  is the minimum, implying  $d_B \geq (d_A d_C)^2 / \varepsilon^8$ . Now let us check if this would change the sample complexity. This would only happen if the contribution of the third regime, which is also  $x_{\min}$ , was smaller than the original sample complexity:  $(d_A d_C)^{1/2} d_B^{3/4} / \varepsilon > d_B^{6/7} (d_A d_C)^{2/7} / \varepsilon^{8/7}$  implies  $(d_A d_C)^2 \varepsilon^{4/3} \geq d_B$ , a contradiction! In the parameter range where we could get an improvement, another regime is always dominant. Similarly, if  $x_{\min} = d_B^{7/8} (d_A d_C)^{1/4} / \varepsilon$  is the minimum, again  $d_B \geq (d_A d_C)^2 / \varepsilon^8$  holds, and  $x_{\min}$  is the contribution by the small regime. A change in sample complexity would occur if  $(d_A d_C)^{1/2} d_B^{3/4} / \varepsilon > d_B^{7/8} (d_A d_C)^{1/4} / \varepsilon$ , but this implies  $(d_A d_C)^2 \geq d_B$ , again a contradiction.

## 8 Lower Bounds for Conditional Independence Testing

In the following, we derive lower bounds for the sample complexity of conditional mutual information testing.

**Theorem 8.1** (Formalized Result 3). *For the sample complexity of conditional independence testing, Problem 6, where w.l.o.g.  $d_A \geq d_C$ , it holds that*

$$\text{SC}_{\text{CI},H}(\varepsilon, d_A, d_B, d_C) = \Omega(\max\{f_{\text{sym}}(\varepsilon, 1, d_B, 1), f_{\text{asym}}(\varepsilon, d_A, d_B, d_C)\}), \quad (258)$$

where, as defined in (11) and (12),

$$f_{\text{sym}}(\varepsilon, 1, d_B, 1) = \min\left\{\frac{d_B^{7/8}}{\varepsilon}, \frac{d_B^{6/7}}{\varepsilon^{8/7}}\right\}, \quad (259)$$

$$f_{\text{asym}}(\varepsilon, d_A, d_B, d_C) = \min\left\{\frac{d_A^{3/4} d_B^{3/4} d_C^{1/4}}{\varepsilon}, \frac{d_A^{2/3} d_B^{2/3} d_C^{1/3}}{\varepsilon^{4/3}}\right\}. \quad (260)$$

The lower bounds for the first term,

$$\text{SC}_{\text{CI},H}(\varepsilon, d_A, d_B, d_C) = \tilde{\Omega}\left(\min\left\{\frac{d_B^{7/8}}{\varepsilon}, \frac{d_B^{6/7}}{\varepsilon^{8/7}}\right\}\right), \quad (261)$$

follows directly from [CDKS18, Remark A.2]. The lower bounds for the second regime are obtained by a reduction to independence testing. Our result in this section is described below.

**Lemma 8.2.** *Testing for conditional independence with respect to  $D_H^2$ , Problem 6, requires*

$$\text{SC}_{\text{CI},H}(\varepsilon, d_A, d_B, d_C) = \Omega\left(\min\left\{\frac{(d_A d_B)^{3/4} d_C^{1/4}}{\varepsilon}, \frac{(d_A d_B)^{2/3} d_C^{1/3}}{\varepsilon^{4/3}}\right\}\right) \quad (262)$$

samples, assuming w.l.o.g.  $d_A \geq d_C$ .

To prove lower bounds for testing conditional independence, we reduce the problem to the lower bounds we derived for independence testing in Section 6. We will construct two sets of hard distributions, such that the distributions in the first set are conditionally independent, and the distributions in the second set are far from being conditionally independent. Based on the outcome of a fair coin toss  $X$ , we will choose one set and then pick a random distribution  $P$  from that set. Finally, we will obtain samples from the distribution  $P$ . We then argue how a lower bound on the sample complexity of reconstructing  $X$  reliably implies a lower bound on the sample complexity of testing for conditional independence. First let us define a set of distributions over  $A \times B \times C$ , with  $|A| = d_A$ ,  $|B| = d_B$ ,  $|C| = d_C$ , as follows (note that these are the same as in the proof of Corollary 6.2, where  $A \rightarrow AB$ ). Similar to our approach in Section 6, we will take  $\text{Poi}(n)$  samples from the distribution.

*Proof.* We construct a distribution as follows. For all  $(a, b) \neq (0, 0)$ :

- For a given  $(a, b)$ , with probability  $\alpha := \min\{n/(d_A d_B), 1/2\}$ , we set  $p_{abc} = 1/(2nd_C)$  for all  $c$ .

– Otherwise, we set for each  $c$  individually

$$p_{abc} = \begin{cases} \frac{\varepsilon}{d_A d_B d_C} & \text{if } X = 0, \\ \text{uniformly random } \frac{\varepsilon}{2d_A d_B d_C} \text{ or } \frac{3\varepsilon}{2d_A d_B d_C} & \text{if } X = 1. \end{cases} \quad (263)$$

$(0, 0, c)$  is then uniform and carries the rest of the weight.

Where  $n$  is chosen as determined in the lower bound with  $d_A \rightarrow d_A d_B$ . Clearly, Theorem 6.1 about independence testing with respect to  $D_H^2$  directly implies lower bounds (and a value of  $n$ ) of

$$n = \Omega \left( \min \left\{ \frac{(d_A d_B)^{3/4} d_C^{1/4}}{\varepsilon}, \frac{(d_A d_B)^{2/3} d_C^{1/3}}{\varepsilon^{4/3}} \right\} \right). \quad (264)$$

However, we could also use a conditional independence tester as a decoder to decide from which set of distributions we received the samples. If  $X = 0$ , that is, if  $P_{ABC} = P_{AB}P_C$ , then  $P_{BC} = P_B P_C$ , and hence  $P_{AB}P_{BC}/P_B = P_{AB}P_C$ . On the other hand, for  $X = 1$ , we will argue in Lemma 8.3 below that  $D_H^2(P_{ABC}, P_{AB}P_{BC}/P_B) \geq \Omega(\varepsilon)$ . This implies that a possible solution to the decision problem would be to use an algorithm taking samples from  $P_{ABC}$  and deciding whether  $D_H^2(P_{ABC}, P_{AB}P_{BC}/P_B) = 0$  or at least  $\Omega(\varepsilon)$ . The lower bound from Corollary 6.2 implies then a lower bound on how efficient any conditional independence tester can perform the task. This completes our reduction.  $\square$

We will now prove that a distribution constructed as above when  $X = 1$ , is far from being conditionally independent.

**Lemma 8.3.** *With high probability, the distributions constructed above with  $X = 1$  satisfy*

$$D_H^2(P_{ABC}, P_{AB}P_{BC}/P_B) \geq \Omega(\varepsilon). \quad (265)$$

*Proof.* For simplicity, we consider only entries for which  $P_{ABC}(a, b, c) = \varepsilon/(2d_A d_B d_C)$ . Let  $S$  denote the set of all triplets  $(a, b, c)$  for which  $p_{abc} = \varepsilon/(2d_A d_B d_C)$ , excluding elements with  $a = 0$  or  $b = 0$ . By construction of our distributions, we know that with high probability,  $|S|$  is within a constant factor of  $d_A d_B d_C/8$ . Then,

$$D_H^2(P_{ABC}, P_{AB}P_{C|B}) \geq D_H^2(P_{ABC}^S, (P_{AB}P_{C|B})^S) = \sum_{(a,b,c) \in S} \left( \sqrt{p_{abc}} - \sqrt{p_{ab}p_{bc}/p_b} \right)^2. \quad (266)$$

– We have  $p_{ab} = \sum_{c'} p_{abc'}$  and  $p_b = \sum_{a',c'} p_{a'bc'}$ . Recall that, for simplicity, we excluded the case where  $a$  or  $b$  are zero. Let  $X$  indicate for how many  $(a, b)$  we set  $p_{abc}$  according to (263), and by  $Y_{ic}$ ,  $i \in \{0, \dots, X\}$  whether we select  $\varepsilon/(2d_A d_B d_C)$  or  $3\varepsilon/(2d_A d_B d_C)$  for the  $i$ th choice and coordinate  $c$ . Then

$$p_b = (d_A - X)d_C \frac{1}{2nd_C} + \sum_{i=1}^X \left( \frac{d_C}{2} + \sum_{c=1}^{d_C} Y_{ic} \right) \frac{\varepsilon}{d_A d_B d_C}, \quad (267)$$

$$p_{bc} = (d_A - X) \frac{1}{2nd_C} + \sum_{i=1}^X \left( \frac{1}{2} + Y_{ic} \right) \frac{\varepsilon}{d_A d_B d_C}. \quad (268)$$

Since  $1 - \alpha \geq 1/2$ , with high probability,  $X = \Theta(d_A)$ . Using a Chernoff bound, we can thus argue that with high probability,  $\sum_{i=1}^X Y_{ic}$  is close to its expectation value,  $X/2$ . Thus, with high probability,

$$0.9d_C \left[ (d_A - X) \frac{1}{2nd_C} + X \frac{\varepsilon}{d_A d_B d_C} \right] \leq p_b \leq 1.1d_C \left[ (d_A - X) \frac{1}{2nd_C} + X \frac{\varepsilon}{d_A d_B d_C} \right], \quad (269)$$

and analogously for  $p_{bc}$  (without the factor  $d_C$ ) Using a Chernoff bound, we can argue that with high probability,  $p_b$  is close to its expectation value, up to a multiplicative factor close to one.

- For  $p_{ab} = \sum_{c'} p_{abc'}$ , we can argue analogously with a Chernoff bound that with high probability  $p_{ab}$  is close to  $\mathbb{E}[p_{ab}] = \varepsilon/(d_A d_B)$ .

Combined, it holds with high probability that for a constant fraction of the elements in  $S$ ,

$$\frac{p_{ab}p_{bc}}{p_b} \geq \frac{3}{4} \frac{\varepsilon}{d_A d_B d_C}. \quad (270)$$

Using the facts that  $p_{abc} = \varepsilon/(2d_A d_B d_C)$  and  $|S| = \Omega(d_A d_B d_C)$ , we have

$$D_H^2(P_{ABC}, P_{AB}P_{C|B})|_S \geq \Omega \left( d_A d_B d_C \left( \sqrt{\frac{1}{2}} - \sqrt{\frac{3}{4}} \right)^2 \frac{\varepsilon}{d_A d_B d_C} \right) = \Omega(\varepsilon). \quad (271)$$

This completes the proof.  $\square$

## 9 Equivalence Testing of Distributions

In this section, we show that when applied to equivalence testing (Problem 5), our estimator (17) is able to recover the optimal sample complexity for testing in  $\ell_2$  distance, see [CDVV14, Thm. 2] and [DK16, Lemma 2.3] under the guarantee that  $\max\{\|P\|_2, \|Q\|_2\} \leq b$ :

$$\text{SC}_{\text{EQIV}}(\ell_1, \varepsilon, b, n) = O\left(\frac{bn}{\varepsilon^2}\right), \quad (272)$$

$$\text{SC}_{\text{EQIV}}(\ell_2, \varepsilon, b, n) = O\left(\frac{b}{\varepsilon^2}\right). \quad (273)$$

Our equivalence tester (Algorithm 9) takes  $N = 200b/\varepsilon^2$  samples of  $P$  and  $Q$  each to distinguish whether  $P = Q$  or  $\|P - Q\|_2 \geq \varepsilon$ . To test with respect to  $\ell_1$  distance instead of  $\ell_2$  distance, we apply Algorithm 9 with precision  $\varepsilon/\sqrt{n}$  and use  $\|P - Q\|_1 \leq \sqrt{n}\|P - Q\|_2$  (see Fact 3.2).

**Theorem 9.1.** *Consider the problem of equivalence testing, Problem 5. Then the estimator from (17) can be used to solve Problem 5 using*

$$\text{SC}_{\text{EQIV}}(\ell_1, \varepsilon, b, n) = O\left(\frac{bn}{\varepsilon^2}\right), \quad \text{SC}_{\text{EQIV}}(\ell_2, \varepsilon, b, n) = O\left(\frac{b}{\varepsilon^2}\right). \quad (274)$$

*samples. Replacing the condition  $\max\{\|P\|_2, \|Q\|_2\} \leq b$  by  $\|Q\|_2 \leq b$  increases the sample complexity by an additive term in  $O(\sqrt{n})$ .*

---

### Algorithm 9: Equivalence Testing ( $P, Q$ )

---

**Input:** Parameter  $b$ , two multisets  $\mathcal{S}_P$  and  $\mathcal{S}_Q$  with  $N$  elements from  $[n]$  each, parameter  $\varepsilon \in (0, 1)$ .  $\triangleright N = 200b/\varepsilon^2$

**Output:** ‘Yes’ or ‘No’

- 1  $X, X' \leftarrow$  Split  $\mathcal{S}_P$  into two subsets of size at least  $100b/\varepsilon^2$  each
  - 2  $Y, Y' \leftarrow$  split  $\mathcal{S}_Q$  into two subsets of size at least  $100b/\varepsilon^2$  each
  - 3  $\forall i \in [n]: X_i \leftarrow \#i \in X, X'_i \leftarrow \#i \in X', Y_i \leftarrow \#i \in Y, Y'_i \leftarrow \#i \in Y$
  - 4  $Z \leftarrow \sum_{i \in [n]} Z_i, \quad Z_i := X_i X'_i - 2X_i Y_i + Y_i Y'_i$
  - 5 **return** ‘Yes’ if  $Z \leq \frac{\varepsilon^2 N^2}{2}$ , otherwise **return** ‘No’
-

To prove the correctness of Algorithm 9, we will use our results from Section 7. As shown in [DK16, Lemma 3.2], the algorithm can easily be modified to work on the weaker assumption as well, increasing the sample complexity by  $O(\sqrt{n})$ .

*Proof of Theorem 9.1.* We first calculate the expectation value of  $Z$ . Here we denote  $N = O(\frac{b}{\varepsilon^2})$ . Note that  $X_i, X'_i \sim \text{Poi}(Np_i)$ . Similarly,  $Y_i, Y'_i \sim \text{Poi}(Nq_i)$ . Thus

$$\mathbb{E}[Z] = \sum_{i \in [n]} \mathbb{E}[X_i] \mathbb{E}[X'_i] - 2\mathbb{E}[X_i] \mathbb{E}[Y_i] + \mathbb{E}[Y_i] \mathbb{E}[Y'_i] \quad (275)$$

$$= \sum_{i \in [n]} (\mathbb{E}[X_i] - \mathbb{E}[Y_i])^2 \quad (276)$$

$$= \sum_{i \in [n]} (Np_i - Nq_i)^2 \quad (277)$$

$$= N^2 \|P - Q\|_2^2. \quad (278)$$

To bound the variance of  $Z$ , we reuse (377),

$$\text{Var}[Z_i] \leq 2(\mathbb{E}[X_i^2] + \mathbb{E}[Y_i^2])^2 \leq 4N^2(p_i^2 + q_i^2). \quad (279)$$

Because of the Poissonization process, the random variables  $\{Z_1, \dots, Z_n\}$  will be independent of each other, such that

$$\text{Var}[Z] = \sum_{i \in [n]} \text{Var}[Z_i] \leq 4 \sum_{i \in [n]} N^2(p_i^2 + q_i^2) = 4N^2(\|P\|_2^2 + \|Q\|_2^2). \quad (280)$$

Let us first consider the case when  $P = Q$ . This implies that  $p_i = q_i$  for every  $i \in [n]$ . Thus, following (278), we know that  $\mathbb{E}[Z] = 0$  and, since  $\max\{\|P\|_2, \|Q\|_2\} \leq b$ ,  $\text{Var}[Z] \leq 8N^2b^2$  from (280).

Following Chebyshev's inequality with  $t = \varepsilon^2 N^2 / 2$  (our threshold), and using that  $N = 100b/\varepsilon^2$ , we have the following:

$$\Pr \left[ Z \geq \frac{N^2 \varepsilon^2}{2} \right] \leq \frac{4 \text{Var}[Z]}{\varepsilon^4 N^4} \leq \frac{32N^2 b^2}{10^4 N^2 b^2} \leq \frac{1}{3}. \quad (281)$$

Now let us consider the case when  $\|P - Q\|_2 \geq \varepsilon$ . From (278), we can say that  $\mathbb{E}[Z] \geq \varepsilon^2 N^2$ . Using Chebyshev's inequality with again  $t = \varepsilon^2 N^2 / 2$ , we have:

$$\Pr \left[ \frac{N^2 \varepsilon^2}{2} \geq Z \right] \leq \Pr[\mathbb{E}[Z]/2 \geq Z] \quad (282)$$

$$= \Pr[\mathbb{E}[Z] - Z \geq \mathbb{E}[Z]/2] \quad (283)$$

$$\leq \Pr[|Z - \mathbb{E}[Z]| \geq \mathbb{E}[Z]/2] \quad (284)$$

$$\leq \Pr[|Z - \mathbb{E}[Z]| \geq N^2 \varepsilon^2 / 4] \quad (285)$$

$$\leq \frac{16 \text{Var}[Z]}{\varepsilon^4 N^4} \quad (286)$$

$$\leq \frac{16 \cdot 80N^2 b^2}{10^4 N^2 b^2} \leq \frac{1}{3}. \quad (287)$$

If we want to loosen the constraint  $\max\{\|P\|_2, \|Q\|_2\} \leq b$  to  $\|Q\|_2 \leq b$ , we first learn  $c_P$ , an approximation of  $\|P\|_2$  up to a constant multiplicative factor 2 according to Lemma 5.4, which takes  $O(\sqrt{d})$  samples. We reject if this test reveals that  $\|P\|_2$  and  $\|Q\|_2$  cannot be equal,  $\|Q\|_2 \leq b < c/2$ . Otherwise we can proceed as before, since we then know that  $\max\{\|P\|_2, \|Q\|_2\} \leq \Theta(b)$ .  $\square$

## Conclusion

In this work, we introduce the following contributions:

- (i) We design the first sample-optimal (up to polylogarithmic factors) mutual information tester and also prove its optimality.
- (ii) We introduce a novel conditional mutual information tester, and prove that it is optimal in certain regimes. Moreover, we conjecture that our upper bounds are tight.

Along the way, we define a new estimator for equivalence testing of distributions, which works also for the correlated samples generated in our sampling approach. We believe this estimator will be of independent interest to the community. Throughout, our proof techniques provide an intuitive explanation for the complicated sample complexities we encounter in our results.

## Open Questions

There remain several relevant open questions closely related to our work, which we describe below:

- (i) For the problem of conditional mutual information testing, we proved lower bounds in the setting of  $f_{\text{sym}}(\varepsilon, 1, d_B, 1)$ ,  $f_{\text{asym}}(\varepsilon, d_A, d_B, d_C)$  as mentioned in Result 3. For  $f_{\text{asym}}(\varepsilon, d_A, d_B, d_C)$ , this matches with the corresponding upper bounds in Result 2. Proving matching lower bounds for the other three terms in the sample complexity remains an interesting open problem. These regimes also have not yet been resolved for conditional independence testing in variation distance.
- (ii) As we mentioned in the introduction, the authors in [CDKS18] followed an algebraic approach to study the conditional independence testing problem with respect to the  $\ell_1$  distance. In this work, we used a combinatorial approach to the conditional mutual information testing problem. Although our approaches to these two different (but related) problems are different, the sample complexities of both problems are very similar (see Result 2 and Equation (13)). This motivates the question whether these two different approaches could be combined into a more general framework, or whether the approaches could be applied to the other respective problem.
- (iii) In this work, we studied the problems of mutual and conditional mutual information testing in the non-tolerant setting, in the sense that we wanted to distinguish between the classes if the (conditional) mutual information is zero or at least  $\varepsilon$ , for some threshold parameter  $\varepsilon$ . An interesting and more general question is the problem of *tolerant* testing of mutual and conditional mutual information. Here, given two threshold parameters  $\varepsilon_1, \varepsilon_2 \in (0, 1)$  with  $\varepsilon_1 < \varepsilon_2$ , the goal is to distinguish between classes where the (conditional) mutual information is either at most  $\varepsilon_1$ , or at least  $\varepsilon_2$ . Often, tolerant testing requires new techniques. Over the last decade, there has been significant progress in this direction (see [Val11, VV11a, VV11b, WY16, WY19, CJKL22]). However, to the best of our knowledge, the tolerant variant of (conditional) mutual information testing has not been explored yet. It would be very interesting to study this problem in the tolerant setting.
- (iv) Finally, we would like to note that in several instances, for various properties of distributions, there is some inherent tolerance even in the algorithms designed for non-tolerant testing. For example, for the setting when  $\varepsilon_1 \leq \varepsilon_2/(2\sqrt{n})$ , the non-tolerant testing algorithm for identity testing of distributions [BFF<sup>+</sup>01, DGPP19] based on the  $\ell_2$  norm of the unknown distribution



is sufficient, due to the reduction from  $\ell_2$  distance to  $\ell_1$  distance. The equivalence tester in  $\ell_2$  [CDVV14] which we use as a subroutine has inherent robustness as well, implying that some robustness may be extracted directly from our approach.

## Acknowledgements

The authors would like to thank the anonymous reviewers for their suggestions, which improved the presentation of the paper. JS would like to thank Christopher Chubb for discussions on an early version of this problem. The authors thank Josep Lumbreras for proof reading the draft. This project is supported by the National Research Foundation, Singapore through the National Quantum Office, hosted in A\*STAR, under its Centre for Quantum Technologies Funding Initiative (S24Q2d0009).

## References

- [ADK15] Jayadev Acharya, Constantinos Daskalakis, and Gautam Kamath. Optimal testing for properties of distributions. *Advances in Neural Information Processing Systems*, 28, 2015.
- [AJOS14] Jayadev Acharya, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Sublinear algorithms for outlier detection and generalized closeness testing. In *2014 IEEE International Symposium on Information Theory*, pages 3200–3204. IEEE, 2014.
- [BFF<sup>+</sup>01] T. Batu, E. Fischer, L. Fortnow, R. Kumar, R. Rubinfeld, and P. White. Testing random variables for independence and identity. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 442–451. IEEE, 2001.
- [BFR<sup>+</sup>13] Tuğkan Batu, Lance Fortnow, Ronitt Rubinfeld, Warren D Smith, and Patrick White. Testing closeness of discrete distributions. *Journal of the ACM (JACM)*, 60(1):1–25, 2013.
- [BGP<sup>+</sup>23] Arnab Bhattacharyya, Sutanu Gayen, Eric Price, Vincent YF Tan, and NV Vinodchandran. Near-optimal learning of tree-structured distributions by chow and liu. *SIAM Journal on Computing*, 52(3):761–793, 2023.
- [BV15] Bhaswar Bhattacharya and Gregory Valiant. Testing closeness with unequal sized samples. *Advances in Neural Information Processing Systems*, 28, 2015.
- [Can20] Clément L Canonne. A survey on distribution testing: Your data is big. but is it blue? *Theory of Computing*, pages 1–100, 2020.
- [Can22] Clément L Canonne. Topics and techniques in distribution testing: A biased but representative sample. *Foundations and Trends® in Communications and Information Theory*, 19(6):1032–1198, 2022.
- [CDKS18] Clément L Canonne, Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. Testing conditional independence of discrete distributions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 735–748, 2018.
- [CDKS20] Clément L Canonne, Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. Testing bayesian networks. *IEEE Transactions on Information Theory*, 66(5):3132–3170, 2020.
- [CDVV14] Siu-On Chan, Ilias Diakonikolas, Paul Valiant, and Gregory Valiant. Optimal algorithms for testing closeness of discrete distributions. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1193–1203. SIAM, 2014.
- [CJKL22] Clément L. Canonne, Ayush Jain, Gautam Kamath, and Jerry Li. The price of tolerance in distribution testing. In *COLT*, volume 178 of *Proceedings of Machine Learning Research*, pages 573–624. PMLR, 2022.
- [CL68] CKCN Chow and Cong Liu. Approximating discrete probability distributions with dependence trees. *IEEE transactions on Information Theory*, 14(3):462–467, 1968.

- [CT06] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, USA, 2006.
- [CYBC24] Davin Choo, Joy Qiping Yang, Arnab Bhattacharyya, and Clément L. Canonne. Learning bounded-degree polytrees with known skeleton. In Claire Vernade and Daniel Hsu, editors, *International Conference on Algorithmic Learning Theory, 25-28 February 2024, La Jolla, California, USA*, volume 237 of *Proceedings of Machine Learning Research*, pages 402–443. PMLR, 2024.
- [DGK<sup>+</sup>21] Ilias Diakonikolas, Themis Gouleakis, Daniel M Kane, John Peebles, and Eric Price. Optimal testing of discrete distributions with high probability. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 542–555, 2021.
- [DGPP19] Ilias Diakonikolas, Themis Gouleakis, John Peebles, and Eric Price. Collision-based testers are optimal for uniformity and closeness. *Chic. J. Theor. Comput. Sci*, 25:1–21, 2019.
- [DK16] Ilias Diakonikolas and Daniel M Kane. A new approach for testing properties of discrete distributions. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 685–694. IEEE, 2016.
- [DKW18] Constantinos Daskalakis, Gautam Kamath, and John Wright. Which distribution distances are sublinearly testable? In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '18*, page 2747–2764, USA, 2018. Society for Industrial and Applied Mathematics.
- [dMASdBP14] Pablo de Morais Andrade, Julio Michael Stern, and Carlos Alberto de Bragança Pereira. Bayesian test of significance for conditional independence: the multinomial model. *Entropy*, 16(3):1376–1395, 2014.
- [DP09] D.P. Dubhashi and A. Panconesi. Concentration of Measure for the Analysis of Randomized Algorithms. In *Cambridge University Press*, 2009.
- [DP17] Constantinos Daskalakis and Qinxuan Pan. Square hellinger subadditivity for bayesian networks and its applications to identity testing. In *Conference on Learning Theory*, pages 697–703. PMLR, 2017.
- [DP21] Constantinos Daskalakis and Qinxuan Pan. Sample-optimal and efficient learning of tree ising models. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 133–146. ACM, 2021.
- [FO24] Steven T Flammia and Ryan O’Donnell. Quantum chi-squared tomography and mutual information testing. *Quantum*, 8:1381, 2024.
- [Gol17] Oded Goldreich. *Introduction to property testing*. Cambridge University Press, 2017.
- [Gra80] Clive WJ Granger. Testing for causality: A personal viewpoint. *Journal of Economic Dynamics and control*, 2:329–352, 1980.

- [GS02] Alison L. Gibbs and Francis Edward Su. On choosing and bounding probability metrics. *International Statistical Review / Revue Internationale de Statistique*, 70(3):419–435, 2002.
- [KNBW22] Ilmun Kim, Matey Neykov, Sivaraman Balakrishnan, and Larry Wasserman. Local permutation tests for conditional independence. *The Annals of Statistics*, 50(6):3388–3414, 2022.
- [KNBW24] Ilmun Kim, Matey Neykov, Sivaraman Balakrishnan, and Larry Wasserman. Conditional independence testing for discrete distributions: Beyond  $\chi^2$ - and  $g$ -tests. *Electronic Journal of Statistics*, 18(2):4767–4794, 2024.
- [LRR13] Reut Levi, Dana Ron, and Ronitt Rubinfeld. Testing properties of collections of distributions. *Theory of Computing*, 9(1):295–347, 2013.
- [MR95] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- [MV19] Alexander Marx and Jilles Vreeken. Testing conditional independence on discrete data using stochastic complexity. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 496–505. PMLR, 2019.
- [N<sup>+</sup>04] Richard E Neapolitan et al. *Learning bayesian networks*, volume 38. Pearson Prentice Hall Upper Saddle River, 2004.
- [NBW21] Matey Neykov, Sivaraman Balakrishnan, and Larry Wasserman. Minimax optimal conditional independence testing. *The Annals of Statistics*, 49(4):2151–2177, 2021.
- [Pea14] Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Elsevier, 2014.
- [Rub12] Ronitt Rubinfeld. Taming big probability distributions. *XRDS: Crossroads, The ACM Magazine for Students*, 19(1):24–28, 2012.
- [SGS01] Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press, 2001.
- [Sha48] C. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 1948.
- [TH18] Marco Tomamichel and Masahito Hayashi. Operational interpretation of rényi information measures via composite hypothesis testing against product and markov distributions. *IEEE Transactions on Information Theory*, 64:1064–1082, 2 2018.
- [Top07] Flemming Topsoe. Some bounds for the logarithmic function. *Inequality Theory and Applications*, 4, 01 2007.
- [Val11] Paul Valiant. Testing symmetric properties of distributions. *SIAM Journal on Computing*, 40(6):1927–1968, 2011.
- [VV11a] 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, 2011.

- [VV11b] Gregory Valiant and Paul Valiant. The power of linear estimators. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 403–412. IEEE, 2011.
- [WGT<sup>+</sup>24] Yuhao Wang, Ming Gao, Wai Ming Tai, Bryon Aragam, and Arnab Bhattacharyya. Optimal estimation of gaussian (poly)trees. In Sanjoy Dasgupta, Stephan Mandt, and Yingzhen Li, editors, *International Conference on Artificial Intelligence and Statistics, 2-4 May 2024, Palau de Congressos, Valencia, Spain*, volume 238 of *Proceedings of Machine Learning Research*, pages 3619–3627. PMLR, 2024.
- [WH18] Xia Wang and Yongmiao Hong. Characteristic function based testing for conditional independence: A nonparametric regression approach. *Econometric Theory*, 34(4):815–849, 2018.
- [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.
- [WY19] Yihong Wu and Pengkun Yang. Chebyshev polynomials, moment matching, and optimal estimation of the unseen. *The Annals of Statistics*, 47(2):857–883, 2019.
- [Wyn78] A. D. Wyner. A definition of conditional mutual information for arbitrary ensembles. *Information and Control*, 38:51–59, 7 1978.
- [ZPJS11] K Zhang, J Peters, D Janzing, and B Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *27th Conference on Uncertainty in Artificial Intelligence (UAI 2011)*, pages 804–813. AUAI Press, 2011.

## A Calculations

### A.1 Remaining Proof from Section 3

**Lemma 3.9.** *For  $0 \leq x < 1$ , it holds that*

$$1 - x + \frac{x^2}{4} \leq e^{-x} \leq 1 - \frac{x}{2}. \quad (26)$$

Moreover, for any  $x \geq 0$ , we have that

$$1 - x + \frac{x^2}{2} - \frac{x^3}{6} \leq e^{-x} \leq 1 - x + \frac{x^2}{2}. \quad (27)$$

*Proof.* Let us first consider the first inequality. Since  $0 \leq x < 1$ , we know that  $e^{-x} \leq 1 - x/2$  holds. The lower bound holds because  $f(x) := e^{-x}$  and  $g(x) := 1 - x + x^2/4$  satisfy  $f(0) = g(0)$ , and for the first derivative, we have that  $f'(x) = -e^{-x} \geq g'(x) = -1 + x/2$ , which holds since  $e^{-x} \leq 1 - x/2$ .

For the second inequality, since  $0 < x$ ,  $e^{-x} \leq 1 - x + x^2/2$  holds from Taylor's expansion. For the lower bound, we consider  $f(x) := e^{-x}$  and  $h(x) := 1 - x + x^2/2 - x^3/6$  which satisfy  $f(0) = g(0)$ , and for the first derivative, we have that  $f'(x) = -e^{-x} \geq h'(x) = -1 + x - x^2/2$ , which follows from the upper bound.  $\square$

### A.2 Remaining Proof from Section 4

**Lemma 4.1.** *Let  $P_{AC}$  and  $T_{AC}$  be two arbitrary distributions over  $A \times C$  and  $\varepsilon \in (0, 1)$  be a threshold. Then, for any  $\alpha \leq \left(\frac{\varepsilon}{48d_{AdC} \log(d_{AdC}/\varepsilon)}\right)^2$ , it holds that*

$$D(P_{AC} \| P_A P_C) \geq \varepsilon \implies D((1 - \alpha)P_{AC} + \alpha T_{AC} \| (1 - \alpha)P_A P_C + \alpha T_{AC}) \geq \varepsilon/2. \quad (29)$$

*Proof.* We define  $\zeta := \varepsilon/(48 \log(d_{AdC}/\varepsilon)d_{AdC})$  and  $\alpha = \zeta^k$ , for  $k$  to be determined. Note that  $\zeta \log(1/\zeta) \leq \varepsilon/(24d_{AdC})$ . Without loss of generality, we assume that  $\alpha < 1/2$ . Further, we can write

$$D(P_{AC} \| P_A P_C) = \sum_{a,c} p_{ac} \log\left(\frac{p_{ac}}{p_a p_c}\right) = \sum_{a,c} p_{ac} \log(p_{ac}) - \sum_{a,c} p_{ac} \log(p_a p_c). \quad (288)$$

We will choose  $\alpha$  such that each of the  $2d_{AdC}$  terms in (288) changes by at most  $\varepsilon/4d_{AdC}$  when introducing  $T_{AC}$ , which then implies the desired result. Let us fix arbitrary  $a$  and  $c$ . We will show

$$|((1 - \alpha)p_{ac} + \alpha t_{ac}) \log((1 - \alpha)p_{ac} + \alpha t_{ac}) - p_{ac} \log(p_{ac})| \leq \frac{\varepsilon}{4d_{AdC}} \quad (289)$$

$$|((1 - \alpha)p_{ac} + \alpha t_{ac}) \log((1 - \alpha)p_a p_c + \alpha t_{ac}) - p_{ac} \log(p_a p_c)| \leq \frac{\varepsilon}{4d_{AdC}} \quad (290)$$

–  $p_{ac} \log(p_{ac})$ : We will do a case distinction. Let us first assume that  $p_{ac} \geq \zeta$ . Then for  $k \geq 2$ ,

we have

$$|((1 - \alpha)p_{ac} + \alpha t_{ac}) \log((1 - \alpha)p_{ac} + \alpha t_{ac}) - p_{ac} \log(p_{ac})| \quad (291)$$

$$= \left| ((1 - \alpha)p_{ac} + \alpha t_{ac}) \left( \log(1 - \alpha) + \log(p_{ac}) + \log \left( 1 + \frac{\alpha t_{ac}}{(1 - \alpha)p_{ac}} \right) \right) - p_{ac} \log(p_{ac}) \right| \quad (292)$$

$$= \left| \alpha(t_{ac} - p_{ac}) \log((1 - \alpha)p_{ac}) + [(1 - \alpha)p_{ac} + \alpha t_{ac}] \log \left( 1 + \frac{\alpha t_{ac}}{(1 - \alpha)p_{ac}} \right) \right| \quad (293)$$

$$\leq \alpha \log \left( \frac{2}{p_{ac}} \right) + (p_{ac} + \alpha) \frac{2\alpha}{p_{ac}} \quad (294)$$

$$\leq 2 \frac{\alpha}{p_{ac}} + 2\alpha + 2 \frac{\alpha^2}{p_{ac}} \leq 6 \frac{\alpha}{p_{ac}} \leq 6\zeta. \quad (295)$$

Now consider the case when  $p_{ac} \leq \zeta \ll 1$ , then  $|p_{ac} \log(p_{ac})| \leq \zeta \log(1/\zeta)$ . Note that  $(1 - \alpha)p_{ac} + \alpha t_{ac} \leq p_{ac} + \alpha \leq 2\zeta$ , such that we can bound

$$|[(1 - \alpha)p_{ac} + \alpha t_{ac}] \log((1 - \alpha)p_{ac} + \alpha t_{ac})| \leq 2\zeta \log(1/\zeta). \quad (296)$$

Together,

$$|((1 - \alpha)p_{ac} + \alpha t_{ac}) \log((1 - \alpha)p_{ac} + \alpha t_{ac}) - p_{ac} \log(p_{ac})| \leq 3\zeta \log(1/\zeta) \leq \varepsilon/(4d_A d_C). \quad (297)$$

–  $p_{ac} \log(p_a p_c)$ : Similar to the above, we perform a case distinction. Let us first assume that  $p_{ac} \geq \zeta$ . Since  $p_a, p_c \geq p_{ac}$ , this implies  $p_a p_c \geq \zeta^2$ . Then, for  $k \geq 2$ , we have

$$|((1 - \alpha)p_{ac} + \alpha t_{ac}) \log((1 - \alpha)p_a p_c + \alpha t_{ac}) - p_{ac} \log(p_a p_c)| \quad (298)$$

$$= |((1 - \alpha)p_{ac} + \alpha t_{ac}) \left( \log((1 - \alpha)p_a p_c) + \log \left( 1 + \frac{\alpha t_{ac}}{(1 - \alpha)p_a p_c} \right) \right) - p_{ac} \log(p_a p_c)| \quad (299)$$

$$= \left| \alpha(t_{ac} - p_{ac}) \log((1 - \alpha)p_a p_c) + [(1 - \alpha)p_{ac} + \alpha t_{ac}] \log \left( 1 + \frac{\alpha t_{ac}}{(1 - \alpha)p_a p_c} \right) \right| \quad (300)$$

$$\leq \left| 2\alpha(t_{ac} - p_{ac}) \log \left( \sqrt{(1 - \alpha)p_a p_c} \right) \right| + [(1 - \alpha)p_{ac} + \alpha t_{ac}] \log \left( 1 + \frac{\alpha t_{ac}}{(1 - \alpha)p_a p_c} \right) \quad (301)$$

$$\leq \left| 2\alpha \log \left( \frac{2}{\sqrt{p_a p_c}} \right) \right| + (p_{ac} + \alpha) \log \left( 1 + \frac{2\alpha t_{ac}}{p_a p_c} \right) \quad (302)$$

$$\leq \left| 2\alpha \log \left( \frac{1}{\zeta} \right) \right| + (p_{ac} + \alpha) \frac{2\alpha}{p_{ac}^2} \leq 6\zeta. \quad (303)$$

Now let us consider the other case. If  $p_{ac} \leq \zeta$ , then  $|p_{ac} \log(p_a p_c)| \leq |2p_{ac} \log(p_{ac})| \leq 2\zeta \log(1/\zeta)$ . With

$$\log \left( \frac{1}{(1 - \alpha)p_a p_c + \alpha t_{ac}} \right) \leq 2 \log \left( \frac{2}{(1 - \alpha)p_{ac} + \alpha t_{ac}} \right), \quad (304)$$

we find

$$|((1 - \alpha)p_{ac} + \alpha t_{ac}) \log((1 - \alpha)p_a p_c + \alpha t_{ac}) - p_{ac} \log(p_a p_c)| \leq 6\zeta \log(1/\zeta) \leq \varepsilon/(4d_A d_C). \quad (305)$$

This implies that each term in (288) changes by at most  $\varepsilon/(2d_A d_C)$  when introducing  $T_{AC}$ , which shows the desired bound.  $\square$

### A.3 Remaining Proofs from Section 6

**Lemma 6.4.** *With high probability, for a distribution  $P_{AC}$  generated using  $X = 1$ , it holds that*

$$D_H^2(P_{AC}, P_{APC}) \geq \Omega(\varepsilon). \quad (90)$$

*Proof.* When bounding

$$D_H^2(P_{AC}, P_{APC}) = \sum_{(a,c)} (\sqrt{p_{ac}} - \sqrt{p_a p_c})^2, \quad (306)$$

we will only sum over coordinates  $(a, c)$  for which  $q(a, c) = \varepsilon/(2d_A d_C)$ . Let  $S_2^0$  denote the set of all tuples  $(a, c)$  for which  $p_{ac} = \varepsilon/(2d_A d_C)$ . A Chernoff bound guarantees that with high probability, we have  $|S_2^0| \geq \Omega(d_A d_C)$ . Next, we analyze  $p_a$  for all  $a \in S_2$ . For a binomial variable  $X_a \sim \text{Bin}(d_A, 1/2)$ ,  $\mu_a := \mathbb{E}[X_a] = d_A/2$  we can write

$$p_a = \frac{\varepsilon}{2d_C} + X_a \frac{\varepsilon}{d_A d_C}. \quad (307)$$

We now use a Chernoff bound with  $\delta := \sqrt{3} \log(d_A)/\sqrt{\mu_a} = \log(d_A)/\sqrt{d_A/2} < 1$  such that (for  $d_A$  sufficiently large)

$$\Pr[|X - d_A/2| \geq d_A/10] \leq \Pr[|X - d_A/2| \geq \log(d_A)\sqrt{3d_A/2}] \leq e^{-\log(d_A)} = 1/d_A. \quad (308)$$

With a union bound, we conclude that with high probability, simultaneously, all  $p_a$  are within a multiplicative factor 1.2 of their expectation value of  $\varepsilon/d_A d_C$ . Similarly,  $p_c$  is lower bounded by  $(1 - \varepsilon/1.2)/d_C$  since we also sum over  $a = 0$  here.

$$p_{a|a \in S_2^0} p_{c|c \in S_2^0} = \left( \sum_{c' \neq c} p_{ac'} + \frac{\varepsilon}{2d_A d_C} \right) \left( \sum_{a' \neq a} p_{a'c} + \frac{\varepsilon}{2d_A d_C} \right) \quad (309)$$

$$\geq \left( p_a - \frac{\varepsilon}{2d_A d_C} \right) \left( p_c - \frac{\varepsilon}{2d_A d_C} \right). \quad (310)$$

Put together, it holds with high probability that  $\forall (a, c) \in S_2^0 : p_{a|a \in S_2^0} p_{c|c \in S_2^0} > 3\mathbb{E}[p_a p_c]/4 = 3\varepsilon/4d_A d_C$ . We then bound (note that under our constraints,  $p_{ac} < p_a p_c$ )

$$D_H^2(P_{AC}, P_{APC}) \geq \sum_{(a,c) \in S_2^0} (\sqrt{p_{ac}} - \sqrt{p_a p_c})^2 \geq \Omega(d_A d_C) \left( \sqrt{\frac{\varepsilon}{2d_A d_C}} - \sqrt{\frac{3\varepsilon}{4d_A d_C}} \right)^2 \geq \Omega(\varepsilon). \quad (311)$$

□

**Claim 6.11.** *For the hard distribution as defined in Section 6.1, for a fixed  $a \in A$ , the following hold:*

(i) *If there is no occurrence of an index  $a$ , i.e., the occurrence vector  $K_a$  is equal to  $\ell_0 := \vec{0}$ , then we have:*

$$\frac{(\Pr[K_a = \ell_0 | X = 0] - \Pr[K_a = \ell_0 | X = 1])^2}{\Pr[K_a = \ell_0 | X = 0] + \Pr[K_a = \ell_0 | X = 1]} \leq O\left(d_C^2 \left[\frac{\varepsilon n}{d_A d_C}\right]^4\right). \quad (140)$$

(ii) *If a specific index  $a$  appears only once, i.e., the occurrence vector  $K_a$  satisfies  $\|K_a\|_1 = 1$ , we have:*

$$\sum_{\ell, \|\ell\|_1=1} \frac{(\Pr[K_a = \ell | X = 0] - \Pr[K_a = \ell | X = 1])^2}{\Pr[K_a = \ell | X = 0] + \Pr[K_a = \ell | X = 1]} \leq O\left(d_C^3 \left[\frac{\varepsilon n}{d_A d_C}\right]^6 \frac{d_C}{\alpha}\right). \quad (141)$$



*Proof.* Note that we have the following:

$$\Pr[K_a = 0|X = 0] \tag{312}$$

$$= \Pr[K_a = 0|X = 0, a \in S_2] \Pr[a \in S_2] + \Pr[K_a = 0|X = 0, a \in S_1] \Pr[a \in S_1] \tag{313}$$

$$= \Pr[K_a = 0|X = 0, a \in S_2] \underbrace{\Pr[a \in S_2]}_{O(1)} + \underbrace{\Pr[K_a = 0|a \in S_1] \Pr[a \in S_1]}_v. \tag{314}$$

Similarly,

$$\Pr[K_a = 0|X = 1] = \Pr[K_a = 0|X = 1, a \in S_2] \Pr[a \in S_2] + v \tag{315}$$

and the second term cancels in the numerator. Let us start by proving (i).

(i) **Case Without Collisions** (For simplicity, we denote such  $K_a$  by ‘0’), again  $u := \varepsilon n/(d_A d_C)$ ,

$$\Pr[K_a = 0|X = 0] = (e^{-u})^{d_C} + v, \quad \Pr[K_a = 0|X = 1] = (e^{-u/2}/2 + e^{-3u/2}/2)^{d_C} + v \tag{316}$$

Assume  $ud_C$  is upper bounded by some constant smaller than 1. Then

$$\frac{(\Pr[K_a = \ell_0|X = 0] - \Pr[K_a = \ell_0|X = 1])^2}{\Pr[K_a = \ell_0|X = 0] + \Pr[K_a = \ell_0|X = 1]} \tag{317}$$

$$= \frac{((e^{-u})^{d_C} - ([e^{-u/2} + e^{-3u/2}]/2)^{d_C})^2}{(e^{-u})^{d_C} + ([e^{-u/2} + e^{-3u/2}]/2)^{d_C} + 2v} \tag{318}$$

$$= \frac{((1 - u + u^2/2 + O(u^3))^{d_C} - (1 - u + 5u^2/8 + O(u^3))^{d_C})^2}{(1 - u + u^2/2 + O(u^3))^{d_C} + (1 - u + 5u^2/8 + O(u^3))^{d_C} + 2v} \tag{319}$$

$$= \frac{\left(\sum_{j=0}^{d_C} \binom{d_C}{j} [(-u + u^2/2 + O(u^3))^j - (-u + 5u^2/8 + O(u^3))^j]\right)^2}{\Theta(1)} \tag{320}$$

$$= \frac{\left(\sum_{j=0}^{d_C} (-u)^j \binom{d_C}{j} [(1 - uj/2 + O(j^2u^2)) - (1 - 5uj/8 + O(j^2u^2))]\right)^2}{\Theta(1)} \tag{321}$$

$$\leq \frac{\left(\sum_{j=1}^{d_C} \binom{d_C}{j} [u(-u)^j j/8 + O(u^{j+2}j^2)]\right)^2}{\Theta(1)} \tag{322}$$

$$\leq \frac{\left(\sum_{j=1}^{d_C} \binom{d_C}{j} u^{j+1} j\right)^2}{\Theta(1)} \tag{323}$$

$$= \frac{\left(u^2 d_C \sum_{j=1}^{d_C} \binom{d_C-1}{j-1} u^{j-1}\right)^2}{\Theta(1)} \tag{324}$$

$$= \frac{u^4 d_C^2 (1+u)^{2d_C}}{\Theta(1)} = \Theta(u^4 d_C^2). \tag{325}$$

Now let us prove (ii).

(ii) **Single Collision:** Here the procedure is the same, but now one of the  $d_C$  results has to result

in a hit.

$$\sum_{\ell, \|\ell\|_1=1} \frac{(\Pr[K_a = \ell|X = 0] - \Pr[K_a = \ell|X = 1])^2}{\Pr[K_a = \ell|X = 0] + \Pr[K_a = \ell|X = 1]} \quad (326)$$

$$= d_C \frac{(u(e^{-u})^{d_C} - ([e^{-u/2}u + 3e^{-3u/2}u]/4)([e^{-u/2} + e^{-3u/2}]/2)^{d_C-1})^2}{\Theta(u+v)} \quad (327)$$

$$= d_C u \frac{((1-u+u^2/2+O(u^3))^{d_C} - (1-5u/4+O(u^2))(1-u+5u^2/8+O(u^3))^{d_C-1})^2}{\Theta(u+v)} \quad (328)$$

$$= d_C u^2 \frac{\left(\sum_{j=0}^{d_C-1} \binom{d_C-1}{j} [(-u+u^2/2+O(u^3))^j(1+O(u)) - (-u+5u^2/8+O(u^3))^j(1+O(u))]\right)^2}{\Theta(u+v)} \quad (329)$$

$$= d_C u^2 \frac{\left(\sum_{j=0}^{d_C-1} (-u)^j \binom{d_C-1}{j} [(1-u/2+O(u^2))^j(1+O(u)) - (1-5u/8+O(u^2))^j(1+O(u))]\right)^2}{\Theta(u+v)} \quad (330)$$

$$= d_C u^2 \frac{\left(\sum_{j=0}^{d_C-1} (-u)^j \binom{d_C-1}{j} [(1-ju/2+O(j^2u^2))(1+O(u)) - (1-5ju/8+O(j^2u^2))](1+O(u))\right)^2}{\Theta(u+v)} \quad (331)$$

$$= d_C u^2 \frac{\left(\sum_{j=0}^{d_C-1} \binom{d_C-1}{j} [(-u)^j u j + (-u)^j O(u^2 j^2)]\right)^2}{\Theta(u+v)} \quad (332)$$

$$\leq d_C u^2 \frac{\left((d_C-1)u^2 \sum_{j=1}^{d_C-1} u^{j-1} \binom{d_C-2}{j-1}\right)^2}{\Theta(u+v)} = \Theta\left(d_C^3 u^6 (1+u)^{d_C}/v\right). \quad (333)$$

Note that  $v = \alpha/(2d_C)$ .

□

#### A.4 Remaining Proofs from Section 7

**Lemma 7.7.** For  $X_{abc}$  and  $Y_{abc}$  distributed according to (168), and (170), respectively, and  $q_{ac|b} = p_{a|b}p_{c|b}$  as defined before, the following hold:

$$(i) \mathbb{E}[X_b] = \mathbb{E}[Y_b] = \frac{\tilde{N}_S p_b}{2} - \frac{e^{-\tilde{N}_S p_b}}{2} \sinh(\tilde{N}_S p_b), \text{ and } \mathbb{E}[X_b^2] = \mathbb{E}[Y_b^2] = \frac{\tilde{N}_S^2 p_b^2}{4} - \frac{\tilde{N}_S p_b e^{-2\tilde{N}_S p_b}}{4} + \frac{e^{-\tilde{N}_S p_b}}{4} \sinh(\tilde{N}_S p_b).$$

$$(ii) \mathbb{E}[X_{abc}] = p_{ac|b} \mathbb{E}[X_b], \text{ and } \mathbb{E}[X_{abc}^2] = p_{ac|b}(1-p_{ac|b}) \mathbb{E}[X_b] + p_{ac|b}^2 \mathbb{E}[X_b^2].$$

$$\text{Similarly, } \mathbb{E}[Y_{abc}] = q_{ac|b} \mathbb{E}[X_b], \text{ and } \mathbb{E}[Y_{abc}^2] = q_{ac|b}(1-q_{ac|b}) \mathbb{E}[X_b] + q_{ac|b}^2 \mathbb{E}[X_b^2].$$

$$(iii) \text{ For } (a, c) \neq (a', c'), \mathbb{E}[X_{abc} X_{a'bc'}] \leq 12 p_{ac|b} p_{a'c'|b} \tilde{N}_S^3 p_b^3, \text{ and } \mathbb{E}[Y_{abc} Y_{a'bc'}] \leq 12 q_{ac|b} q_{a'c'|b} \tilde{N}_S^3 q_b^3.$$

*Proof.* Since the proof for  $Y_{abc}$  is identical to the proof of  $X_{abc}$ , obtained by replacing  $p_{ac|b}$  with  $q_{ac|b}$ , we will only prove the statements for  $X_{abc}$ . Let us start by bounding the  $\mathbb{E}[X_{abc}]$ .

$$\mathbb{E}[X_{abc}] = \sum_{\ell=0}^{\infty} \Pr[X_b = \ell] \mathbb{E}[X_{abc} | X_b = \ell] \quad (334)$$

$$= \sum_{\ell=0}^{\infty} \Pr[X_b = \ell] \ell p_{ac|b} \quad (335)$$

$$= p_{ac|b} \mathbb{E}[X_b]. \quad (336)$$

Similarly, we can calculate  $\mathbb{E}[X_{abc}^2]$  as well. We have

$$\mathbb{E}[X_{abc}^2] = \sum_{\ell=0}^{\infty} \mathbb{E}[X_{abc}^2 | X_b = \ell] \Pr[X_b = \ell] \quad (337)$$

$$= \sum_{\ell=0}^{\infty} \left[ \sum_{k=0}^{\ell} k^2 \binom{\ell}{k} p_{ac|b}^{\ell} (1 - p_{ac|b})^{k-\ell} \right] \Pr[X_b = \ell] \quad (338)$$

$$= \sum_{\ell=1}^{\infty} \ell p_{ac|b} \left[ \sum_{k=1}^{\ell} k \binom{\ell-1}{k-1} p_{ac|b}^{\ell-1} (1 - p_{ac|b})^{k-\ell} \right] \Pr[X_b = \ell] \quad (339)$$

$$= \sum_{\ell=1}^{\infty} \ell p_{ac|b} \left[ 1 + \sum_{k=1}^{\ell} (k-1) \binom{\ell-1}{k-1} p_{ac|b}^{\ell-1} (1 - p_{ac|b})^{k-\ell} \right] \Pr[X_b = \ell] \quad (340)$$

$$= \sum_{\ell=1}^{\infty} \ell p_{ac|b} [1 + (\ell-1)p_{ac|b}] \Pr[X_b = \ell] \quad (341)$$

$$= p_{ac|b} ((1 - p_{ac|b})\mathbb{E}[X_b] + p_{ac|b}\mathbb{E}[X_b^2]). \quad (342)$$

In the following steps, we omit the indices  $abc$  for simplicity.

$$\sum_{k=0}^{\infty} k \frac{x^{2k} e^{-x}}{(2k)!} = \frac{x}{2} \sum_{k=1}^{\infty} \frac{x^{2k-1} e^{-x}}{(2k-1)!} = \frac{x}{2} \sum_{k=0}^{\infty} \frac{x^{2k+1} e^{-x}}{(2k+1)!} = \frac{x e^{-x}}{2} \sinh(x) \quad (343)$$

$$\sum_{k=0}^{\infty} k \frac{x^{2k+1} e^{-x}}{(2k+1)!} = \frac{1}{2} \sum_{k=0}^{\infty} (2k+1) \frac{x^{2k+1} e^{-x}}{(2k+1)!} - \frac{1}{2} \frac{x^{2k+1} e^{-x}}{(2k+1)!} = \frac{e^{-x}}{2} [x \cosh(x) - \sinh(x)] \quad (344)$$

$$\sum_{k=0}^{\infty} k^2 \frac{x^{2k} e^{-x}}{(2k)!} = \sum_{k=0}^{\infty} \left[ \frac{(2k)(2k-1)}{4} + \frac{k}{2} \right] \frac{x^{2k} e^{-x}}{(2k)!} = \frac{e^{-x}}{4} [x^2 \cosh(x) + x \sinh(x)] \quad (345)$$

$$\sum_{k=0}^{\infty} k^2 \frac{x^{2k+1} e^{-x}}{(2k+1)!} = \sum_{k=0}^{\infty} \left[ \frac{(2k)(2k+1)}{4} - \frac{k}{2} \right] \frac{x^{2k+1} e^{-x}}{(2k+1)!} = \frac{e^{-x}}{4} [x^2 \sinh(x) - x \cosh(x) + \sinh(x)] \quad (346)$$

Let  $x_b := \tilde{N}_{Sp_b}$ . We can now calculate:

$$\mathbb{E}[X_b] = \frac{x_b}{2} - \frac{e^{-x_b}}{2} \sinh(x_b), \quad \mathbb{E}[X_b^2] = \frac{x_b^2}{4} - \frac{x_b e^{-2x_b}}{4} + \frac{e^{-x_b}}{4} \sinh(x_b). \quad (347)$$

Also note that  $\mathbb{E}[X_b] \leq \mathbb{E}[X_b^2]$ . To bound this further, note that writing out  $\sinh$  and using the fact that  $e^x \geq 1 + x$ , we have

$$\mathbb{E}[X_b^2] \leq \frac{2x_b^2 + 1}{8} - \frac{(1 + 2x_b)(1 - 2x_b)}{8} = \frac{3}{4} x_b^2 \quad (348)$$

Moreover, since  $\forall b \in [d_B], \forall k \in \mathbb{N} : \mathbb{E}[X_b^k] = \mathbb{E}[Y_b^k], \mathbb{E}[X_b] = \mathbb{E}[Y_b]$  as well as  $\mathbb{E}[X_b^2] = \mathbb{E}[Y_b^2]$ .

In the following, assume  $(a, c) \neq (a', c')$ , but  $a = a'$  or  $c = c'$  is permitted. Note that for  $0 < x, y$  with  $x + y < 1$ , we have that  $(1 - x - y)^z \leq (1 - x - y + xy)^z = (1 - x)^z(1 - y)^z$ . Then

$$\mathbb{E}[X_{abc}X_{a'bc'}] \tag{349}$$

$$= \sum_{k, \ell, s=0} k\ell \Pr[X_{abc} = k, X_{a'bc'} = \ell | X_b = s] \Pr[X_b = s] \tag{350}$$

$$= \sum_{s=0}^{\infty} \left[ \sum_{k=0}^s \sum_{\ell=0}^{s-k} k\ell \binom{s}{k} \binom{s-k}{\ell} p_{ac|b}^{\ell} p_{a'c'|b}^k (1 - p_{ac|b} - p_{a'c'|b})^{s-k-\ell} \right] \Pr[X_b = s] \tag{351}$$

$$= \sum_{s=2}^{\infty} \left[ \sum_{k=1}^s \sum_{\ell=1}^{s-k} k\ell \binom{s}{k} \binom{s-k}{\ell} p_{ac|b}^{\ell} p_{a'c'|b}^k (1 - p_{ac|b} - p_{a'c'|b})^{s-k-\ell} \right] \Pr[X_b = s] \tag{352}$$

$$\leq \sum_{s=2}^{\infty} \left[ \sum_{k=1}^s k \binom{s}{k} p_{a'c'|b}^k (1 - p_{a'c'|b})^{s-k} \sum_{\ell=1}^{s-k} \ell \binom{s-k}{\ell} p_{ac|b}^{\ell} \frac{(1 - p_{ac|b})^{s-k-\ell}}{(1 - p_{a'c'|b})^{\ell}} \right] \Pr[X_b = s] \tag{353}$$

Due to symmetry, we can assume without loss of generality that  $p_{a'c'|b} < p_{ac|b}$ , which in particular means that  $p_{a'c'|b} < 1/2$ . We now note that

$$\sum_{\ell=1}^{s-k} \frac{\ell}{(1 - p_{a'c'|b})^{\ell}} \binom{s-k}{\ell} p_{ac|b}^{\ell} (1 - p_{ac|b})^{s-k-\ell} \tag{354}$$

$$\leq \sum_{\ell=1}^{s-k} 2^{\ell} \ell \binom{s-k}{\ell} p_{ac|b}^{\ell} (1 - p_{ac|b})^{s-k-\ell} \tag{355}$$

$$\leq 2p_{ac|b}(s-k) \sum_{\ell=1}^{s-k} 2^{\ell-1} \binom{s-k-1}{\ell-1} p_{ac|b}^{\ell-1} (1 - p_{ac|b})^{s-k-\ell} \tag{356}$$

$$= 2p_{ac|b}(s-k) (2p_{ac|b} + (1 - p_{ac|b}))^{s-k-1} \tag{357}$$

$$\leq p_{ac|b}(s-k) 2^{s-k}. \tag{358}$$

In the following, bounding  $s - k \leq s$  also allows us to simplify the sum over  $k$  analogously. In the final steps, we use inequalities from Lemma 3.9.

$$\mathbb{E}[X_{abc}X_{a'bc'}] \leq p_{ac|b} p_{a'c'|b} \sum_{s=2}^{\infty} \left[ \sum_{k=1}^s \binom{s-1}{k-1} s(s-k) 2^{s-k} p_{a'c'|b}^{k-1} (1 - p_{a'c'|b})^{s-k} \right] \Pr[X_b = s] \tag{359}$$

$$\leq p_{ac|b} p_{a'c'|b} \sum_{s=2}^{\infty} s^2 2^s \Pr[X_b = s] \tag{360}$$

$$\leq p_{ac|b} p_{a'c'|b} e^{\sqrt{2}x_b} \sum_{s=2}^{\infty} s^2 \left( \frac{(\sqrt{2}x_b)^{2s} e^{-\sqrt{2}x_b}}{(2s)!} + \frac{(\sqrt{2}x_b)^{2s+1} e^{-\sqrt{2}x_b}}{(2s+1)!} \right) \tag{361}$$

We now let  $k_b := \sqrt{2}x_b = \tilde{N}_S(\sqrt{2}p_b)$ , and  $K_b$  accordingly. Then

$$\mathbb{E}[X_{abc}X_{a'bc'}] = p_{ac|b}p_{a'c'|b}e^{\sqrt{2}x_b}(\mathbb{E}[K_b^2] - (k_b^2e^{-k_b}/2 + k_b^3e^{-k_b}/6)) \quad (362)$$

$$= p_{ac|b}p_{a'c'|b}e^{\sqrt{2}x_b} \left( \frac{2k_b^2 + 1}{8} - \frac{2k_b + 1}{8}e^{-2x_b} - \left( \frac{k_b^2e^{-k_b}}{2} + \frac{k_b^3e^{-k_b}}{6} \right) \right) \quad (363)$$

$$\leq p_{ac|b}p_{a'c'|b}e^{\sqrt{2}x_b} \left( \frac{2k_b^2 + 1}{8} - \frac{2k_b + 1}{8} \left( 1 - 2k_b + 2k_b^2 - \frac{4k_b^3}{3} \right) \right) \quad (364)$$

$$- \left( \frac{k_b^2}{2} + \frac{k_b^3}{6} \right) \left( 1 - k_b + \frac{k_b^2}{4} \right) \quad (365)$$

$$\leq p_{ac|b}p_{a'c'|b}e^{\sqrt{2}x_b} \left( \frac{2k_b^2 + 1}{8} - \frac{1 - 2k_b^2}{8} - \frac{k_b^2}{2} + \frac{k_b^3}{2} \right) \leq 12p_{ac|b}p_{a'c'|b}x_b^3. \quad (366)$$

□

**Lemma 7.9.** For  $i \in [4]$ , let  $\mathcal{S}_i$  be multisets of i.i.d. samples where  $|\mathcal{S}_i|$  is determined by drawing from  $\text{Poi}(\tilde{N}_S)$ . Let  $X = \text{SimABC}(\mathcal{S}_1, B_S)$ ,  $X' = \text{SimABC}(\mathcal{S}_2, B_S)$  (Algorithm 4), and  $Y = \text{SimABC}_{\text{Cl}}(\mathcal{S}_3, B_S)$ ,  $Y' = \text{SimABC}_{\text{Cl}}(\mathcal{S}_4, B_S)$  (Algorithm 5), and  $Z$  constructed from  $X$ ,  $X'$ ,  $Y$ , and  $Y'$  according to (171). Then

$$\text{Var}[Z] \leq 2 \cdot 10^3 (\|P_{ABC}^S\|_2^2 + \|Q_{ABC}^S\|_2^2) \tilde{N}_S^2. \quad (184)$$

*Proof.* Recall that  $Z = \sum_i Z_i$  where  $Z_i := X_iX'_i - 2X_iY_i + Y_iY'_i$  and we would like to prove that

$$\text{Var}[Z] = \sum_{(a,b,c) \in S} \text{Var}[Z_{abc}] + \text{Cov}_{(a,b,c) \neq (a',b,c')} [Z_{abc}, Z_{a'bc'}] \quad (367)$$

$$\leq 2 \cdot 10^3 (\|P_{ABC}^S\|_2^2 + \|Q_{ABC}^S\|_2^2) \tilde{N}_S^2. \quad (368)$$

This result follows directly from the two following claims, Claim A.1 and Claim A.2, which bound the variance and covariance, respectively. □

**Claim A.1.**  $\sum_{(a,b,c) \in S} \text{Var}[Z_{abc}] \leq 8 (\|P_{ABC}^S\|_2^2 + \|Q_{ABC}^S\|_2^2) \tilde{N}_S^2$ .

*Proof.* Let us first compute  $\mathbb{E}[Z_i^2]$  and  $\mathbb{E}[Z_i]^2$ .

$$Z_i^2 = X_i^2(X'_i)^2 + 4X_i^2Y_i^2 + Y_i^2(Y'_i)^2 - 4X_i^2X'_iY_i + 2X_iX'_iY_iY'_i - 4X_iY_i^2Y'_i. \quad (369)$$

So,

$$\mathbb{E}[Z_i^2] = \mathbb{E}[X_i^2]^2 + 4\mathbb{E}[X_i^2]\mathbb{E}[Y_i^2] + \mathbb{E}[Y_i^2]^2 - 4\mathbb{E}[X_i^2]\mathbb{E}[X_i]\mathbb{E}[Y_i] + 2\mathbb{E}[X_i]^2\mathbb{E}[Y_i]^2 - 4\mathbb{E}[X_i]\mathbb{E}[Y_i^2]\mathbb{E}[Y_i] \quad (370)$$

Similarly, we have:

$$\mathbb{E}[Z_i]^2 = \mathbb{E}[X_i]^4 - 4\mathbb{E}[X_i]^3\mathbb{E}[Y_i] + 6\mathbb{E}[X_i]^2\mathbb{E}[Y_i]^2 - 4\mathbb{E}[X_i]\mathbb{E}[Y_i]^3 + \mathbb{E}[Y_i]^4 \quad (371)$$

Now let us compute  $\text{Var}[Z_i]$ .

$$\text{Var}[Z_i] = \mathbb{E}[Z_i^2] - \mathbb{E}[Z_i]^2 \quad (372)$$

$$= \mathbb{E}[X_i^2]^2 + 4\mathbb{E}[X_i^2]\mathbb{E}[Y_i^2] + \mathbb{E}[Y_i^2]^2 - \mathbb{E}[X_i^4] - \mathbb{E}[Y_i^4] \quad (373)$$

$$- 4\mathbb{E}[X_i]\mathbb{E}[Y_i] (\mathbb{E}[X_i]\mathbb{E}[Y_i] + \mathbb{E}[X_i^2] + \mathbb{E}[Y_i^2] - \mathbb{E}[X_i]^2 - \mathbb{E}[Y_i]^2) \quad (374)$$

$$= \mathbb{E}[X_i^2]^2 + 4\mathbb{E}[X_i^2]\mathbb{E}[Y_i^2] + \mathbb{E}[Y_i^2]^2 - \mathbb{E}[X_i^4] - \mathbb{E}[Y_i^4] \quad (375)$$

$$- 4\mathbb{E}[X_i]\mathbb{E}[Y_i] (\mathbb{E}[X_i]\mathbb{E}[Y_i] + \text{Var}[X_i] + \text{Var}[Y_i]) \quad (376)$$

$$\leq 2(\mathbb{E}[X_i^2] + \mathbb{E}[Y_i^2])^2, \quad (377)$$

such that we can bound

$$\sum_{(a,b,c) \in S} \text{Var}[Z_{abc}] \leq 2 \sum_{(a,b,c) \in S} (\mathbb{E}[X_{abc}^2] + \mathbb{E}[Y_{abc}^2])^2 \quad (378)$$

From Lemma 7.7, we know that  $\mathbb{E}[X_{abc}^2] = p_{ac|b}(1 - p_{ac|b})\mathbb{E}[X_b] + p_{ac|b}^2\mathbb{E}[X_b^2]$ , and  $\mathbb{E}[Y_{abc}^2]$  analogously with  $q$  instead of  $p$ . Thus, using  $\mathbb{E}[X_b^2] = \mathbb{E}[Y_b^2] \leq (Np_b)^2$  (see (348)), we have:

$$\sum_{(a,b,c) \in S} \text{Var}[Z_{abc}] \leq 2 \sum_{(a,b,c) \in S} (\mathbb{E}[X_{abc}^2] + \mathbb{E}[Y_{abc}^2])^2 \quad (379)$$

$$\leq 2 \sum_{(a,b,c) \in S} (p_{ac|b}\mathbb{E}[X_b] + p_{ac|b}^2\mathbb{E}[X_b^2] + q_{ac|b}\mathbb{E}[Y_b] + q_{ac|b}^2\mathbb{E}[Y_b^2])^2 \quad (380)$$

$$\leq 8 \sum_{(a,b,c) \in S} (p_{ac|b}\mathbb{E}[X_b^2] + q_{ac|b}\mathbb{E}[Y_b^2])^2 \quad (381)$$

$$\leq 8 \sum_{(a,b,c) \in S} (p_{ac|b} + q_{ac|b})^2 \tilde{N}_S^4 p_b^4 \quad (382)$$

$$\leq 8(\|P_{ABC}^S\|_2^2 + \|Q_{ABC}^S\|_2^2) \tilde{N}_S^2. \quad (383)$$

□

**Claim A.2.**  $\text{Cov}_{(a,b,c) \neq (a',b,c')}[Z_{abc}, Z_{a'bc'}] \leq 12^3 (\|P_{ABC}^S\|_2^2 + \|Q_{ABC}^S\|_2^2) \tilde{N}_S^2$ .

*Proof.* Let us now compute the covariance. From the definition, we have the following:

$$\text{Cov}_{(a,b,c) \neq (a',b,c')}[Z_{abc}, Z_{a'bc'}] = \sum_{(a,c)} \sum_{(a',c') \neq (a,c)} \mathbb{E}[(Z_{abc} - \mathbb{E}[Z_{abc}])(Z_{a'bc'} - \mathbb{E}[Z_{a'bc'}])] \quad (384)$$

$$= \sum_{(a,c)} \sum_{(a',c') \neq (a,c)} \mathbb{E}[Z_{abc}Z_{a'bc'}] - \mathbb{E}[Z_{abc}]\mathbb{E}[Z_{a'bc'}] \quad (385)$$

$$\leq \sum_{(a,c)} \sum_{(a',c') \neq (a,c)} \mathbb{E}[Z_{abc}Z_{a'bc'}] \quad (386)$$

For ease of reading, let us denote  $Z_{abc}$  as  $Z_i$  and  $Z_{a'bc'}$  as  $Z'_i$ . We use similar notations for  $X_i, X'_i$  and  $Y_i, Y'_i$  as well.

Then

$$Z_i Z'_i = (X_i X'_i - 2X_i Y_i + Y_i Y'_i)(X_j X'_j - 2X_j Y_j + Y_j Y'_j) \quad (387)$$

$$= X_i X'_i X_j X'_j - 2X_i X'_i X_j Y_j + X_i X'_i Y_j Y'_j - 2X_i Y_i X_j X'_j + 4X_i Y_i X_j Y_j - 2X_i Y_i Y_j Y'_j \quad (388)$$

$$+ Y_i Y'_i X_j X'_j - 2Y_i Y'_i X_j Y_j + Y_i Y'_i Y_j Y'_j \quad (389)$$

So, we have:

$$\mathbb{E}[Z_i Z'_i] = \mathbb{E}[(X_i X'_i - 2X_i Y_i + Y_i Y'_i)(X_j X'_j - 2X_j Y_j + Y_j Y'_j)] \quad (390)$$

$$= (\mathbb{E}[X_i X_j])^2 - 2\mathbb{E}[X_i X_j]\mathbb{E}[Y_j]\mathbb{E}[X_i] + (\mathbb{E}[X_i Y_j])^2 - 2\mathbb{E}[X_i X_j]\mathbb{E}[Y_i]\mathbb{E}[X_j] \quad (391)$$

$$+ 4\mathbb{E}[X_i X_j]\mathbb{E}[Y_i Y_j] - 2\mathbb{E}[Y_i Y_j]\mathbb{E}[X_i]\mathbb{E}[Y_j] \quad (392)$$

$$+ (\mathbb{E}[Y_i X_j])^2 - 2\mathbb{E}[Y_i Y_j]\mathbb{E}[X_j]\mathbb{E}[Y_i] + (\mathbb{E}[Y_i Y_j])^2 \quad (393)$$

and summing up gives

$$\sum_{i,j} (\mathbb{E}[X_i X_j])^2 - 4(\mathbb{E}[X_i X_j] + \mathbb{E}[Y_i Y_j])\mathbb{E}[Y_j]\mathbb{E}[X_i] \quad (394)$$

$$+ 2\mathbb{E}[X_i]^2\mathbb{E}[Y_j]^2 + 4\mathbb{E}[X_i X_j]\mathbb{E}[Y_i Y_j] + (\mathbb{E}[Y_i Y_j])^2 \quad (395)$$

$$\leq 2 \sum_{i,j} ((\mathbb{E}[X_i X_j] + \mathbb{E}[Y_i Y_j])^2 + 2\mathbb{E}[X_i]^2\mathbb{E}[Y_j]^2) \quad (396)$$

$$\leq 4 \sum_{(a,b,c)} \sum_{(a',b,c') \neq (a,b,c)} (\mathbb{E}[X_{abc} X_{a'bc'}]^2 + \mathbb{E}[Y_{abc} Y_{a'bc'}]^2 + 2\mathbb{E}[X_{abc}]^2\mathbb{E}[Y_{a'bc'}]^2). \quad (397)$$

Now let us first bound the term  $\sum_{(a,b,c)} \sum_{(a',b,c') \neq (a,b,c)} \mathbb{E}[X_{abc} X_{a'bc'}]^2$  using Lemma 7.7.

$$\sum_b \sum_{(a,c)} \sum_{(a',c') \neq (a,c)} \mathbb{E}[X_{abc} X_{a'bc'}]^2 \leq 12^2 \sum_b \sum_{(a,c)} \sum_{(a',c')} p_{ac|b}^2 p_{a'c'|b}^2 p_b^6 \tilde{N}_S^6 \quad (398)$$

$$\leq 12^2 \sum_b \sum_{(a,c)} p_{abc}^2 p_b^4 \tilde{N}_S^6 \quad (399)$$

$$\leq 12^2 \|P_{ABC}^S\|_2^2 \tilde{N}_S^2, \quad (400)$$

Similarly, we can bound  $\sum_{(a,b,c)} \sum_{(a',b,c') \neq (a,b,c)} \mathbb{E}[Y_{abc} Y_{a'bc'}]^2 \leq 12^2 \|Q_{ABC}^S\|_2^2 \tilde{N}_S^2$ .

Now let us bound

$$\sum_b \sum_{(a,c)} \sum_{(a',c') \neq (a,c)} \mathbb{E}[X_{abc}]^2 \mathbb{E}[Y_{a'bc'}]^2 \leq \sum_b \sum_{(a,c)} \sum_{(a',c')} p_{ac|b}^2 (p_{a'|b} p_{c'|b})^2 \tilde{N}_S^8 p_b^8 \quad (401)$$

$$\leq \sum_b \sum_{(a,c)} p_{ac|b}^2 \tilde{N}_S^8 p_b^8 \leq \|P_{ABC}^S\|_2^2 \tilde{N}_S^2. \quad (402)$$

Combining all the above, we can say that

$$\text{Cov}_{(a,b,c) \neq (a',b,c')} [Z_{abc}, Z_{a'bc'}] \leq 12^3 (\|P_{ABC}^S\|_2^2 + \|Q_{ABC}^S\|_2^2) \tilde{N}_S^2. \quad (403)$$

This completes the proof.  $\square$