

# On Sample-Based Testers\*

Oded Goldreich<sup>†</sup>      Dana Ron<sup>‡</sup>

September 22, 2015

## Abstract

The standard definition of property testing endows the tester with the ability to make arbitrary queries to “elements” of the tested object. In contrast, sample-based testers only obtain independently distributed elements (a.k.a. labeled samples) of the tested object. While sample-based testers were defined by Goldreich, Goldwasser, and Ron (*JACM* 1998), with few exceptions, most research in property testing has focused on query-based testers.

In this work, we advance the study of sample-based property testers by providing several general positive results as well as by revealing relations between variants of this testing model. In particular:

- We show that certain types of query-based testers yield sample-based testers of sublinear sample complexity. For example, this holds for a natural class of proximity oblivious testers.
- We study the relation between distribution-free sample-based testers and one-sided error sample-based testers w.r.t. the uniform distribution.

While most of this work ignores the time complexity of testing, one part of it does focus on this aspect. The main result in this part is a sublinear-*time* sample-based tester, in the dense graphs model, for  $k$ -Colorability, for any  $k \geq 2$ .

**Keywords:** Property Testing, Graph Properties, One-Sided Error

---

\*An extended abstract of this paper appeared in the proceedings of the *6th Innovations in Theoretical Computer Science* (ITCS), 2015. Partially supported by the Israel Science Foundation (grant No. 671/13).

<sup>†</sup>Department of Computer Science, Weizmann Institute of Science, Rehovot, ISRAEL. E-mail: [oded.goldreich@weizmann.ac.il](mailto:oded.goldreich@weizmann.ac.il)

<sup>‡</sup>Department of EE-Systems, Tel-Aviv University, Ramat-Aviv, ISRAEL. E-mail: [danar@eng.tau.ac.il](mailto:danar@eng.tau.ac.il)

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Sample-based testers of sublinear sample complexity . . . . .	1
1.2	Perspectives on other testing models . . . . .	3
1.2.1	Testing symmetric properties and testing distributions . . . . .	3
1.2.2	Distribution-Free Testing and One-Sided Error . . . . .	3
1.3	Related work . . . . .	4
1.4	The computational complexity aspect . . . . .	5
1.5	Some of the ideas that underly our proofs . . . . .	5
1.5.1	Deriving sample-based testers from fair POTs . . . . .	5
1.5.2	Distribution-free versus one-sided error testing . . . . .	6
1.5.3	Computationally efficient tester for $k$ -colorability . . . . .	7
1.6	Organization . . . . .	7
<b>2</b>	<b>Preliminaries</b>	<b>7</b>
2.1	The standard definitions of query-based testers . . . . .	8
2.2	Sample-based testers . . . . .	9
<b>3</b>	<b>POTs and sample-based testers</b>	<b>10</b>
3.1	The main result . . . . .	10
3.2	On the fairness condition . . . . .	13
3.3	A generalization (of Theorem 3.1) . . . . .	17
<b>4</b>	<b>Sample-based testers in the dense graphs model</b>	<b>19</b>
4.1	General results . . . . .	19
4.2	Graph partition problems . . . . .	21
<b>5</b>	<b>Sample-based testers for linearity</b>	<b>31</b>
<b>6</b>	<b>Invariance under the symmetric group</b>	<b>33</b>
6.1	Queries versus samples . . . . .	33
6.2	The role of the range size . . . . .	34
6.3	Relation to testing distributions . . . . .	35
<b>7</b>	<b>Distribution-Free Testing</b>	<b>38</b>
7.1	A general result . . . . .	39
7.2	Specific results . . . . .	41
	<b>Acknowledgements</b>	<b>49</b>
	<b>References</b>	<b>49</b>
	<b>Appendix: In Passing</b>	<b>53</b>
A.1	A relaxed notion of POT . . . . .	53
A.2	Relating various notions of sampling . . . . .	54
A.3	On the sample complexity of testing monomials . . . . .	55

# 1 Introduction

In the last couple of decades, the area of property testing has attracted much attention (see, e.g., [Gol10, Ron08, Ron10]). Loosely speaking, property testing typically refers to sub-linear complexity probabilistic algorithms for deciding whether a given object has a predetermined property or is far from any object having this property. Such algorithms, called testers, obtain local views of the object by performing queries; that is, the object is seen as a function and the testers get oracle access to this function (and thus may be expected to work in time that is sub-linear in the length of the object).

The standard definition of property testing, which is alluded to above, endows the tester with the ability to make queries. This is the definition presented in [RS96], and it is also the main definition studied in [GGR98] as well as in most research on property testing. Nevertheless, a weaker model, where the tester is only provided with uniformly distributed labeled samples (or, equivalently, is only allowed to make uniformly and *independently distributed* queries)<sup>1</sup> was also presented (and briefly studied) in [GGR98]. We call such testers *sample-based*. (For precise definitions, see Section 2.)

While sample-based testers were studied in [GGR98]<sup>2</sup> and in some subsequent works (see, e.g., [GGL<sup>+</sup>00]), these studies tend to be negative in nature: Their focus is on showing why the query ability utilized in the main positive results is necessary for efficiently testing the property that is being considered. This amounts to presenting high lower bounds on the sample complexity of sample-based testers for that property. A few positive results on sample-based testers have been obtained in previous work [GGL<sup>+</sup>00, KR00, BBY12], and they are further discussed in Section 1.3. We also discuss the relation to testing properties of *distributions* (in Section 1.2.1).

In this work we aim at a broader study of sample-based property testing and we focus on sample-based testers of sublinear sample complexity. We obtain several general positive results and reveal relations between variants of this testing model.

We believe that sample-based testers are of practical potential. Although sample-based testers are typically much less efficient than general (query-based) testers, in many applications, random labeled samples are easier to obtain than a full query capacity. Hence, sample-based testers of sublinear (sample) complexity may be preferable to query-based testers in such applications.

## 1.1 Sample-based testers of sublinear sample complexity

The complexity of testers (of various types) is measured as a function of the *size of the object*, denoted  $n$ , and the *proximity parameter*, denoted  $\epsilon$ . (Both these parameters are given as input to standard testers.) When we say that a complexity measure is *sublinear*, we mean that it is sublinear as a function of  $n$ , when fixing  $\epsilon$ . (Typically, the query and sample complexities are at least linear in  $1/\epsilon$ .) We stress that when the tested object is a function, its size is the size of the function's domain, not the length of elements in that domain; for example, the size of  $f: \{0, 1\}^\ell \rightarrow \{0, 1\}$  is  $n = 2^\ell$ . With one exception (i.e., Section 5), all our positive results assert sublinear sample complexity for testing properties that are impossible to learn within such complexity (under the uniform distribution).

---

<sup>1</sup>We stress that the canonical testers of graph properties, which are used in many works and are studied explicitly in [GT03, Section 4], do *not* make uniformly and independently distributed queries. Indeed, they select at random a uniformly distributed set of vertices, but their queries (which correspond to all vertex pairs) are dependent, although each is uniformly distributed. In general, many testers make uniformly distributed queries, but these queries are typically not fully independent of one another. See further discussion following Theorem 1.1.

<sup>2</sup>In fact, the main definition in [GGR98] refers to (distribution-free) sample-based testers (cf. [GGR98, Definition 2.1]).

Indeed, we view any tester of sublinear sample complexity as a positive result, even if its complexity has the form  $O_\epsilon(n^{0.9})$ . This stands in contrast to some earlier work in which complexity lower bounds of the form  $\Omega(n^{0.1})$  were viewed as negative results. This change in perspective is related to viewing the size of the object (i.e.,  $n$ ) as the point of reference rather than viewing the length of the queries to the object (i.e.,  $\log n$ ) as the point of reference.

**The POT connection.** One source of positive results is the model of Proximity Oblivious Testers (POTs). Specifically, we transform certain POTs into sample-based testers of sublinear sample complexity, where the “level of sublinearity” depends on the (constant) query complexity of the POT. Loosely speaking, Proximity Oblivious Testers (POTs) are “basic testers” that do not get a proximity parameter as input, and make only a constant number of queries to the tested object [GR11]. They are required to reject objects with probability that is lower bounded by a function of the object’s distance to the property (see Definition 2.2).<sup>3</sup> This function is called the *detection probability function*.

**Theorem 1.1** (POTs with uniformly distributed queries imply sample-based testers of sublinear sample complexity): *Suppose that  $\Pi$  has a  $q$ -query Proximity Oblivious Tester that makes uniformly distributed queries, and let  $\varrho$  denote its detection probability function. Then,  $\Pi$  has a sample-based tester with sample complexity  $s(n, \epsilon) = O(n^{1-(1/q)}/\varrho(\epsilon)^{2+(3/q)})$ . Furthermore, if the POT has one-sided error, then so does the sample-based tester.*

We stress that the premise of Theorem 1.1 is merely that each of the queries made by the POT is uniformly distributed, although the queries may be dependent on one another and may even be chosen adaptively. A typical example is the celebrated **linearity** tester of [BLR93], which is based on a three-query POT. This three-query POT selects  $x$  and  $y$  uniformly and independently in the function’s domain (which is a group), and makes the queries  $x, y$  and  $x + y$ .

Theorem 1.1 implies that **linearity** has a sample-based tester with sample complexity  $O(n^{2/3}\text{poly}(1/\epsilon))$ , but a better (and tight) result can be obtained directly (see Theorem 5.1). In general, POTs that satisfy the “uniform query distribution” condition are quite common; in fact, in the dense graph model, any POT can be converted to one that makes uniformly distributed queries, while at most squaring the number of queries (see [GR11, Section 4.2], following [GT03, Section 4]).<sup>4</sup> Actually, we can get stronger results for the dense graph model.

**Theorem 1.2** (on sample-based testers in the dense graph model): *Let  $n$  denote the size of the adjacency matrix of a graph, that is, the number of vertex-pairs.*

1. (From quasi-canonical testers to sample-based testers): *Let  $\Pi$  be a graph property and consider the task of testing  $\Pi$  in the dense graph model. If  $\Pi$  has a POT with detection probability function  $\varrho$  that inspects the subgraph induced by a random set of  $\nu$  vertices, then  $\Pi$  has a sample-based tester with sample complexity  $s(n, \epsilon) = O(n^{1-(1/(\nu-1))}/\varrho(\epsilon)^2)$ . Furthermore, if the POT has one-sided error, then so does the sample-based tester.*
2. *In the dense graph model, for every  $k \geq 2$ ,  $k$ -Colorability has a sample-based tester with sample complexity  $s(n, \epsilon) = O(1/\epsilon) \cdot \sqrt{n} \log k$ . Furthermore, this tester has one-sided error.*

---

<sup>3</sup>In contrast, standard testers do get a proximity parameter, denoted  $\epsilon$ , as input, make a number of queries that depend on  $\epsilon$ , and are required to reject any object that is  $\epsilon$ -far from the property with probability at least  $2/3$ .

<sup>4</sup>The resulting POT inspects the subgraph induced by a random set of  $O(q)$  vertices, where  $q$  is the query complexity of the original POT.

Note that Bipartiteness (2-Colorability) does not have a POT (cf. [GR11]). Hence, while Item 1 merely presents a quantitative improvement over Theorem 1.1, Item 2 goes beyond the scope of Theorem 1.1. Two more comments are in order:

1. The square-root dependence on  $n$  of the sample complexity in the second item of the theorem is tight for (sample-based) testing of bipartiteness.
2. The tester referred to in the second item of the theorem is not computationally efficient. We address this issue in Section 1.4.

## 1.2 Perspectives on other testing models

Sample-based testing offers interesting perspectives on several notions.

### 1.2.1 Testing symmetric properties and testing distributions

As argued in [KS08] (see also [GK11]), natural families of properties are defined in terms of invariances, where each property in the family is invariant under a group of permutations acting on the functions' domain. (Indeed, testing graph properties in the dense graph model (as in Theorem 1.2) is one famous example, and linearly invariant properties [KS08] is another.) At the very extreme, one may consider *symmetric properties*, which are properties that are invariant under the symmetric group acting on the functions' domain. We observe that *when testing symmetric properties, samples are essentially as good as queries* (see Theorem 6.1).

We also articulate Sudan's observation that *testing distributions* (for any property of distributions, cf. [BFR<sup>+</sup>00]) *can be tightly reduced to testing symmetric properties* [Sud, Section 2.1]. In particular, Theorem 6.4 formally relates the model of testing distributions<sup>5</sup> to the standard model of property testing (i.e., testing properties of functions w.r.t. the uniform distribution over their domains). While establishing this relation is technically simple, to the best of our knowledge, this is the first time that the model of testing distributions has been rigorously related to the standard model of property testing, which relates to testing functions (w.r.t. the uniform distribution over their domains).

### 1.2.2 Distribution-Free Testing and One-Sided Error

The focus on sample-based testers brings us closer to the area of computational learning theory (cf. [Val84, KV94]). Within this mind-frame, it is natural to consider distribution-free testing, which was also defined in [GGR98] but received relatively little attention so far. In distribution-free testing, the sample that the testing algorithm receives is distributed according to an arbitrary and unknown distribution, and distance between functions (and hence distance to having the property) is defined according to the same distribution. Our initial feeling was that there may be a relation between *one-sided error* sample-based testing (under the uniform distribution) and *distribution-free* (sample-based) testing. This feeling was partially confirmed by a general upper bound on the sample complexity of the former in terms of the sample complexity of the latter, but beyond this upper bound the two complexity measures may exhibit different relationships. Our results

---

<sup>5</sup>In this model the tested object is a distribution  $D$  over some domain  $X$  and the algorithm receives points  $x \in X$  that are distributed according to  $D$ . The goal is to decide whether  $D$  has a particular property or is far from any distribution that has the property (where the distance measure between distributions is usually the variation distance).

are summarized in the following theorem, where  $\text{OSE}(\Pi)$  denotes the sample complexity of one-sided error sample-based testing  $\Pi$  (under the uniform distribution), and  $\text{DF}(\Pi)$  denotes the sample complexity of distribution-free (sample-based) testing  $\Pi$ .

**Theorem 1.3** (distribution-free sample-based testers versus one-sided error sample-based testers under the uniform distribution):

1. For every property  $\Pi$ , it holds that  $\text{OSE}(\Pi) = \tilde{O}(\text{DF}(\Pi)^2)$ .
2. There exists a property  $\Pi$  such that  $\text{OSE}(\Pi) = \Omega(n)$  but  $\text{DF}(\Pi) \leq \text{poly}(1/\epsilon) \cdot \frac{n}{\log n}$ .
3. There exists a property  $\Pi$  such that  $\text{OSE}(\Pi) = \Theta(\text{DF}(\Pi))$ .
4. There exists a property  $\Pi$  such that  $\text{OSE}(\Pi) = \tilde{\Theta}(1/\epsilon) \cdot \log \text{DF}(\Pi)$ .
5. There exists a property  $\Pi$  such that  $\text{OSE}(\Pi) = \tilde{O}(1/\epsilon)$  but  $\text{DF}(\Pi) = \Omega(n)$ .

Hence, there exists a general upper bound of  $\text{OSE}$  in terms of  $\text{DF}$  (i.e., Item 1), but in specific cases we may see quite different relations ranging from  $\text{OSE} = \omega(\text{DF})$ , to  $\text{OSE} = \Theta(\text{DF})$ , to  $\text{OSE} = o(\text{DF})$  (cf., Items 2-5). We mention that the properties used in Items 2-4 are natural ones (see Section 1.5.2 for a description of these properties).

Regarding the general upper bound of  $\text{OSE}$  in terms of  $\text{DF}$ , note that Item 1 asserts  $\text{OSE} = \tilde{O}(\text{DF}^2)$  whereas Item 2 rules out a result better than  $\text{OSE} = \tilde{O}(\text{DF})$ . This leaves a considerable gap in our understanding, and addressing it is a natural open problem.

### 1.3 Related work

As mentioned upfront, sample-based testers were considered in several prior works, starting with the work of Goldreich, Goldwasser, and Ron [GGR98]. Ironically, the main definition in their work (i.e. [GGR98, Definition 2.1]), refers to (distribution-free) sample-based testers, whereas the now-standard definition of query-based testers (with respect to the uniform distribution) is presented there as a variant (see Item 3 in [GGR98, Section 2]). However, the bulk of [GGR98] is devoted to the study of query-based testers, and this notion became the standard in the area.

Nevertheless, sample-based testers were considered also in subsequent works, which focused on query-based testers. Typically, the perspective is negative; that is, the focus is on lower bounds on sample complexity, which are presented as a justification for the use of queries in the main positive results. For example, the first study of testing monotonicity [GGL<sup>+</sup>00] focuses on query-based testers, but also provides a lower bound on the sample complexity of sample-based testers, which is also shown to be tight. Indeed, the latter lower bound (i.e., [GGL<sup>+</sup>00, Theorem 5]) is used as justification for the use of queries in the main result (i.e., [GGL<sup>+</sup>00, Theorem 1]), and the sample-based tester (of [GGL<sup>+</sup>00, Theorem 6]) is viewed as indicating that this lower bound is tight.

Two notable exceptions appear in [KR00, BBBY12]: Kearns and Ron [KR00] consider sample-based testing (under the uniform distribution) for decision trees of a bounded size  $s$  over  $[0, 1]^d$  (for constant  $d$ ) and for a special class of neural networks with  $s$  hidden units. They design testers whose sample complexity is significantly lower than that required for learning the corresponding class of functions. However, their testers are only required to reject functions that are far from a superclass of the tested class (determined by a larger size parameter  $s'$ ). For the special case of interval functions ( $d = 1$ ), Balcan *et al.* [BBBY12] showed how this relaxation of the rejection requirement can be removed, and they obtained optimal (in terms of the dependence on  $s$ ) sample-based testers.

Balcan *et al.* [BBBY12] also present sample-based testers (which out-perform learning algorithms) for linear threshold functions under the Gaussian distribution.

Actually, the work of Balcan *et al.* [BBBY12] puts forward a more refined notion, called **active testers**. These may be viewed as a generalization of sample-based testers, where the testers are provided with *unlabeled* samples, and may query the function only at points that appear in the given sample. Balcan *et al.* [BBBY12] consider both the sample-complexity and the query-complexity of these testers, where the latter is typically smaller (since the tester does not query the function on all the sample points). They view sample-based testers as a special case (which they call *passive testing*) in which the testers query the function on all points in the sample (and so in this case their notion of query-complexity equals the notion of sample-complexity).

## 1.4 The computational complexity aspect

Throughout most of this work, we ignore the computational complexity aspect (i.e., the running-time of the various testers). This choice seems crucial to some of our results; examples include Proposition 2.4, Theorem 3.1 (see also Theorems 1.1 and 3.5), Theorem 4.5 (see Item 2 in Theorem 1.2), and Theorem 7.2. In contrast, the following result (i.e., Theorem 1.4) is of interest only because it addresses the computational complexity aspect, which was ignored in Item 2 of Theorem 1.2.

**Theorem 1.4** (on computationally efficient sample-based testers in the dense graph model):

1. *Bipartiteness* has a (one-sided error) *sample-based tester* of time complexity  $O(1/\epsilon) \cdot \sqrt{n}$ .
2. For every  $k \geq 3$ , *k-Colorability* has a (one-sided error) *sample-based tester* of time complexity  $f_k(\epsilon) \cdot n^{1-(1/2k)}$ , where  $f_k(\epsilon) \stackrel{\text{def}}{=} \exp(\exp(\tilde{O}(k/\epsilon)))$ .<sup>6</sup>

The computationally efficient Bipartiteness tester (of Theorem 1.4) matches the sample complexity asserted in Item 2 of Theorem 1.2, but for  $k \geq 3$  the sublinear time tester for *k-Colorability* uses a sublinear sample complexity that is larger than the bound asserted in Item 2 of Theorem 1.2. A begging open problem is *whether, for  $k \geq 3$ , there exists a sample-based tester of time complexity  $\text{poly}(1/\epsilon) \cdot \sqrt{n}$ . for *k-Colorability*.*

## 1.5 Some of the ideas that underly our proofs

In this section we attempt to give a flavour of the ideas underlying the proofs of our results. In some cases we provide proof outlines and in other only some of the ideas.

### 1.5.1 Deriving sample-based testers from fair POTs

Let us say that a  $q$ -query POT is *fair* if, when given query access to a function  $f$ , each of its (possibly adaptive)  $q$  queries is uniformly distributed in  $f$ 's domain  $[n]$ . We stress that these  $q$  queries may depend on one another and may even depend on the answers to prior queries (since the POT may be adaptive).

Fixing a fair POT  $T$ , for each setting  $\omega$  of its random bits, we obtain a deterministic  $q$ -query tester  $T_\omega$  (which may be adaptive). Consider a *labeled* sample  $S = ((x_1, f(x_1)), \dots, (x_s, f(x_s)))$  of size  $s = qt$  where  $t$  is roughly  $n^{1-(1/q)}$ , and let us view it as a sequence of  $q$  samples,  $S_1, \dots, S_q$ , each of size  $t$ . We shall say that  $S$  is *good* for  $T_\omega$ , if for each  $j \in [q]$ , the sample  $S_j$  includes the

---

<sup>6</sup>Here and elsewhere, the notation  $\exp(x)$  means  $c^x$  for some constant  $c > 1$ .

$j^{\text{th}}$  query of  $T_\omega$ . The sample-based tester takes a sample  $S$  of size  $s = qt$  and uniformly selects a setting of  $\omega$  such that  $S$  is good for  $T_\omega$  (the probability that there is no such  $\omega$  is small). It then computes the output of  $T_\omega$  based on the answers to its queries as appearing in the sample.

In order to analyze this sample based tester, observe first that for any fixed  $\omega$ , the probability that  $S$  is good for  $T_\omega$  equals  $\mu = (1 - (1 - 1/n)^t)^q \approx (t/n)^q$ . Furthermore, using the fairness of  $T$  we show that, with high probability over the choice of  $S$ , the fraction of settings of  $\omega$  for which  $S$  is good, is very close to  $\mu$ . This is then used to show that the sample-based tester approximately emulates “true” executions of the POT  $T$ .

### 1.5.2 Distribution-free versus one-sided error testing

Recall that Theorem 1.3 relates the sample complexity of distribution-free testers (which may have two-sided error) and one-sided error testers (under the uniform distribution). We first show that the latter is not too big in terms of the former (i.e., Item 1).

Suppose that  $\Pi$  has a distribution-free tester  $T_{\text{df}}$  with error probability  $1/6$  and sample complexity  $s \ll \sqrt{n}$ . Consider the following one-sided error tester  $T_{\text{ose}}$  that takes a sample of size  $s' = \Theta(s^2)$  labeled by a function  $f$ , and accepts if and only if the sample is consistent with some function in  $\Pi$ . We claim that  $T_{\text{ose}}$  must reject, with probability at least  $2/3$ , any function  $f$  that is far (under the uniform distribution) from  $\Pi$ .

Assume, contrary to the claim, that  $T_{\text{ose}}$  accepts some function  $f$  that is far from  $\Pi$  with probability greater than  $1/3$ . That is, more than a third of the samples of size  $s'$  labeled by  $f$  are each consistent with some function in  $\Pi$ . Then, for each of these samples, we consider the uniform distribution over the corresponding (multi-)set of points. Since the distribution-free tester  $T_{\text{df}}$  must work for any distribution, for each of these samples,  $T_{\text{df}}$  accepts with probability at least  $5/6$  when provided with a (sub-)sample of  $s$  elements. This can be shown to imply that  $T_{\text{df}}$  accepts a sample of size  $s$  that is uniformly selected from the whole domain and labeled by  $f$  with probability at least  $\frac{1}{3} \cdot \frac{5}{6} - 0.1 > \frac{1}{3}$ , in contradiction to the hypothesis regarding  $T_{\text{df}}$ .

In order to show that one-sided error testing (under the uniform distribution) may require more samples than distribution-free testing (i.e., Item 2), we consider the set of *Boolean functions* (over  $[n]$ ) *having a majority of 1-values*. Clearly, any one-sided error tester must use more than  $n/2$  samples (or else it is bound to always accept and hence err on the all zero function). Our distribution-free tester first approximates the average value of the function (under the unknown distribution  $\mu$ ). If the estimate is above  $1 - (\epsilon/2)$ , then it accepts (since this is likely to happen only when the function is  $\epsilon$ -close (under the distribution  $\mu$ ) to being the all-1 function, which is in the property). Otherwise, we may assume that the expected value is at most  $1 - (\epsilon/4)$ , so that  $\mu(f^{-1}(0)) \geq \epsilon/4$ . We now consider the residual distribution of  $x$  selected from the unknown distribution conditioned on  $f(x) = 0$  (i.e.,  $x \in f^{-1}(0)$  is selected with probability  $\mu(x)/\mu(f^{-1}(0))$ ). Since  $\mu(f^{-1}(0)) \geq \epsilon/4$ , we can efficiently sample from the residual distribution and estimate the support size of the residual distribution using the algorithm of Valiant and Valiant [VV11]. This algorithm uses  $O(n/\log n)\text{poly}(1/\epsilon)$  samples, and Item 2 follows.

In order to show that the sample complexity of one-sided error testing (under the uniform distribution) may equal the sample complexity of distribution-free testing (i.e., Item 3), we consider the set of *monotone non-decreasing functions from  $[n]$  to  $[n]$* . It can be shown that (even under the uniform distribution) sample-based testers (of two-sided error) require  $\Omega(\sqrt{n/\epsilon})$  samples, whereas  $O(\sqrt{n/\epsilon})$  samples suffice for a tester of one-sided error (which actually works in the distribution-free setting). Item 3 follows.

In order to show that one-sided error testing (under the uniform distribution) may be signifi-

cantly easier than distribution-free testing (i.e., Item 4), we consider the set of *Boolean functions* from  $\{0,1\}^\ell$  to  $\{0,1\}$  that may be expressed as *monomials* (of the  $\ell$  variables). We show that distribution-free testing requires  $\Omega(\sqrt{\ell}/\epsilon)$  samples, but one-sided error testing (under the uniform distribution) can be performed using  $\tilde{O}(\log \ell/\epsilon)$  samples. Item 4 follows.

### 1.5.3 Computationally efficient tester for $k$ -colorability

Item 2 of Theorem 1.2 is easily proved by presenting an inefficient tester that considers all possible  $k$ -way partitions of the graph. In contrast, Theorem 1.4 asserts efficient testers for  $k$ -Colorability. The first item is proved by noting that for 2-colorability all 2-partitions can be considered implicitly; that is, the tester just checks whether the subgraph revealed in the labeled sequence of samples is bipartite. For  $k > 2$ , we use a much more sophisticated argument, which builds on a technique introduced by Alon and Krivelevich [AK02].

Specifically, Alon and Krivelevich showed that if a graph is  $\epsilon$ -far from being  $k$ -colorable, then, with high constant probability, a subgraph induced by  $\Theta(k \log k/\epsilon^2)$  uniformly selected vertices will not be  $k$ -colorable. To this end they introduced (as a mental experiment) a process in which new vertices that are added to the sample *restrict* the legal colorings of previously added vertices.<sup>7</sup> However, their analysis uses the fact that with each newly added vertex we get all edges to previously selected vertices (as part of the induced subgraphs). In contrast, with the sample size that we use, we cannot expect to obtain so much information, and this is one aspect in which we depart from their analysis. The second aspect is that we need to turn the mental experiment into an efficient algorithm (which finds a small subgraph that is not  $k$ -colorable), let alone that the algorithm lacks some of the knowledge that the mental experiment has.

## 1.6 Organization

In Section 2.1 we recall the standard definitions of (query-based) testers and proximity oblivious testers. In Section 2.2 we recall and briefly discuss the definition of sample-based testers.

In Section 3 we show that certain POTs imply sample-based testers of sublinear sample complexity. In particular, we present a generalization of Theorem 1.1, which derives sample-based testers of sublinear complexity from any POT that refrains from “heavy queries” (i.e., queries that assign much weight to a specific location). Sample-based testers in the dense graphs model are studied in Section 4, and the proofs of Theorems 1.2 and 1.4 appear there. Section 5 presents an optimal (one-sided error) sample-based tester of linearity. In Section 6 we study symmetric properties.

In Section 7 we study the relation between one-sided error sample-based testing (under the uniform distribution) and distribution-free (sample-based) testing: The results of this study are summarized in Theorem 1.3.

In Appendix A.1, we consider a relaxation of the definition of POTs. The only relation of this section to the rest of the paper is that such relaxed POTs can be used in Section 3.

## 2 Preliminaries

Property testing is a relaxation of decision problems and it focuses on algorithms that can only read parts of the input. Thus, the input is represented as a function (to which the tester has oracle

---

<sup>7</sup>An alternative process was introduced previously in [GGR98], but this alternative did not lend itself to our purposes.

access) and the tester is required to accept functions that have some predetermined property (i.e., reside in some predetermined set) and reject any function that is “far” from the set of functions having the property. Distances between functions are defined as the fraction of the domain on which the functions disagree, and the threshold determining what is considered far is presented as a proximity parameter, which is explicitly given to the tester.

An asymptotic analysis is enabled by considering an infinite sequence of domains, functions, and properties. The domains and properties (in the infinite sequence) are described by a finite sequence of parameters, which include the size of the domain, denoted  $n$ . For simplicity, we shall present our results for sequences that are determined by this size parameter. That is, for any  $n$ , we consider functions from  $D_n$  to  $R_n$ , where  $|D_n| = n$ . (Often, one just assumes that  $D_n = [n] \stackrel{\text{def}}{=} \{1, 2, \dots, n\}$ .) Thus, in addition to the input oracle, representing a function  $f: D_n \rightarrow R_n$ , the tester is explicitly given two parameters: a size parameter, denoted  $n$ , and a proximity parameter, denoted  $\epsilon$ .

## 2.1 The standard definitions of query-based testers

In this section we recall the standard definition of property testing as well as the definition of proximity oblivious testers. Both definitions refer to oracle machines that are given oracle access to a function  $f: D_n \rightarrow R_n$  (as well as free access to some relevant parameters such as  $n$ ). We denote by  $M^f(p)$  the output of oracle machine  $M$  on input parameter  $p$  when given oracle access to  $f$ .

**Definition 2.1** (property tester, following [RS96, GGR98]): *Let  $\Pi = \bigcup_{n \in \mathbb{N}} \Pi_n$ , where  $\Pi_n$  contains functions defined over the domain  $D_n$  and ranging over  $R_n$ . A tester for a property  $\Pi$  is a probabilistic oracle machine  $T$  that satisfies the following two conditions:*

1. *The tester accepts each  $f \in \Pi$  with probability at least  $2/3$ ; that is, for every  $n \in \mathbb{N}$  and  $f \in \Pi_n$  (and every  $\epsilon > 0$ ), it holds that  $\Pr[T^f(n, \epsilon) = 1] \geq 2/3$ .*
2. *Given proximity parameter  $\epsilon > 0$  and oracle access to any  $f$  that is  $\epsilon$ -far from  $\Pi$ , the tester rejects with probability at least  $2/3$ ; that is, for every  $n \in \mathbb{N}$  and  $\epsilon > 0$ , if  $f: D_n \rightarrow R_n$  is  $\epsilon$ -far from  $\Pi_n$ , then  $\Pr[T^f(n, \epsilon) = 0] \geq 2/3$ , where  $f$  is  $\epsilon$ -far from  $\Pi_n$  if, for every  $g \in \Pi_n$ , it holds that  $\delta(f, g) \stackrel{\text{def}}{=} |\{e \in D_n : f(e) \neq g(e)\}|/n > \epsilon$ .*

*Indeed, the distance of  $f$  from  $\Pi$ , denoted  $\delta_\Pi(f)$ , equals  $\inf_{g \in \Pi_n} \{\delta(f, g)\}$ .*

*If the tester accepts every function in  $\Pi$  with probability 1, then we say that it has one-sided error; that is,  $T$  has one-sided error if for every  $f \in \Pi_n$  and every  $\epsilon > 0$ , it holds that  $\Pr[T^f(n, \epsilon) = 1] = 1$ . A tester is called **non-adaptive** if it determines all its queries based solely on its internal coin tosses (and the parameters  $n$  and  $\epsilon$ ); otherwise it is called **adaptive**.*

Definition 2.1 does not specify the query complexity of the tester, and indeed an oracle machine that queries the entire domain of the function qualifies as a tester (with zero error probability...). Needless to say, we are interested in testers that have significantly lower query complexity.

Some testers (e.g., the celebrated linearity tester of [BLR93]) operate by repeating some basic tests for a number of times that depends on the proximity parameter, whereas the basic test is oblivious of the proximity parameter. Such basic tests are captured by the following definition.

**Definition 2.2** (Proximity Oblivious Tester (POT), following [GR11, GS12]): *Let  $\Pi$  be as in Definition 2.1 and let  $\varrho: (0, 1] \rightarrow (0, 1]$  be monotone non-decreasing. A POT with detection probability  $\varrho$  for  $\Pi$  is a probabilistic oracle machine  $T$  that makes a constant number of queries and satisfies the following two conditions with respect to some constant  $c \in (0, 1]$ :*

1. For every  $n \in \mathbb{N}$  and  $f \in \Pi_n$ , it holds that  $\Pr[T^f(n)=1] \geq c$ .
2. For every  $n \in \mathbb{N}$  and  $f: D_n \rightarrow R_n$  not in  $\Pi_n$ , it holds that  $\Pr[T^f(n)=1] \leq c - \varrho(\delta_\Pi(f))$ ,

The constant  $c$  is called the **threshold probability**. A POT is said to have **one-sided error** if  $c = 1$ .

We stress that, in contrast to a standard tester, a POT only gets one explicit parameter (i.e., the size parameter,  $n$ ). Standard testers are obtained by invoking the POT for an adequate number of times. Specifically, for one-sided error POTs, we invoke the POT for  $O(1/\varrho(\epsilon))$  times and accept if and only if all invocations returned 1. For general POTs with threshold probability  $c$ , we invoke the POT for  $O(1/\varrho(\epsilon)^2)$  times and accept if and only if at least  $c - (\varrho(\epsilon)/2)$  of the invocations returned 1.

## 2.2 Sample-based testers

As discussed in the introduction, a sample-based tester does not have query access to the function  $f: D_n \rightarrow R_n$  but rather obtains a sequence of “labeled sample points” of the function values. That is, it obtains a sequence of points that are independently and uniformly distributed in  $D_n$  such that these points are coupled with the corresponding values of the function. The length of this sequence is predetermined based on  $n$  and  $\epsilon$ .

**Definition 2.3** (sample-based tester, following [GGR98]): *Let  $\Pi$  be as in Definition 2.1 and  $s: \mathbb{N} \times (0, 1] \rightarrow \mathbb{N}$ . A (sample-based) tester of sample complexity  $s$  for a property  $\Pi$  is a probabilistic algorithm  $T$  that satisfies the following two conditions for every  $n \in \mathbb{N}$ :*

1. For every  $f \in \Pi_n$  (and every  $\epsilon > 0$ ), it holds that

$$\Pr_{r_1, \dots, r_s \in D_n}[T(n, \epsilon, (r_1, f(r_1)), \dots, (r_s, f(r_s))) = 1] \geq \frac{2}{3},$$

where  $s = s(n, \epsilon)$  and  $(r_1, \dots, r_s)$  is uniformly distributed in  $D_n^s$ .

2. For every  $\epsilon > 0$ , and every  $f: D_n \rightarrow R_n$  that is  $\epsilon$ -far from  $\Pi$ ,

$$\Pr_{r_1, \dots, r_s \in D_n}[T(n, \epsilon, (r_1, f(r_1)), \dots, (r_s, f(r_s))) = 0] \geq \frac{2}{3},$$

where again  $s = s(n, \epsilon)$  and  $(r_1, \dots, r_s)$  is uniformly distributed in  $D_n^s$ .

The sequence  $(r_1, f(r_1)), \dots, (r_s, f(r_s))$  is called a **sample** labeled by  $f$ . As in Definition 2.1, if the tester accepts every function in  $\Pi$  with probability 1, then we say that it has **one-sided error**.

An alternative definition views such a tester as one that may make queries to the function *as long as its queries are distributed uniformly in the domain, independently of one another*. That is, each query of such a tester is distributed uniformly in  $D_n$ , independently of all prior queries (and of the answers provided to these queries). The equivalence of this definition to Definition 2.3 is based on our disregard of the computational complexity of the tester (i.e., the complexity of generating queries). This is clarified in the proof of the following proposition.

**Proposition 2.4** (equivalence of the two formulations): *A property  $\Pi$  has a tester of query complexity  $q$  that uses queries that are uniformly and independently distributed in the domain if and only if it has a sample-based tester of sample complexity  $q$ .*

**Proof:** Given a sample-based tester  $T_{\text{sb}}$  (as defined in Definition 2.3), we easily obtain a corresponding querying tester  $T_{\text{q}}$  by having  $T_{\text{q}}$  emulate  $T_{\text{sb}}$  in a straightforward manner. That is, on input  $n$  and  $\epsilon$ , the tester  $T_{\text{q}}$  selects uniformly and independently an adequate (i.e.,  $q(n, \epsilon)$ ) number of queries, obtains the corresponding answers by querying its oracle, and feeds  $T_{\text{sb}}$  with the corresponding sequence of query and answer pairs.<sup>8</sup>

The other direction is a bit more tricky. The issue is that the final decision of the (query based) tester may depend on the coins that it has used for producing its queries, and so merely replacing the queries by a uniformly distributed (labeled) sample will not do. We should augment this (labeled) sample with a sequence of random coins that would have led the original tester to make the corresponding queries. Details follow.

We are given a query-based tester  $T_{\text{q}}$  with a guarantee on the distribution of its queries (to  $f$ ), and wish to construct a sample-based tester  $T_{\text{sb}}$ . The sample-based tester  $T_{\text{sb}}$  is given a sequence of pairs, denoted  $(r_1, v_1), \dots, (r_s, v_s)$ , where  $v_i = f(r_i)$  for each  $i \in [s]$ , and it operates by selecting at random coins  $\omega$  such that using coins  $\omega$  and having oracle access to  $f$ , the oracle machine  $T_{\text{q}}$  makes the queries  $r_1, \dots, r_s$ . By the hypothesis, the set of such  $\omega$ 's constitutes an  $n^{-s}$  fraction of the set of all possible coin tosses (of  $T_{\text{q}}$ ). Furthermore,  $T_{\text{sb}}$  can reconstruct the former set without making any queries to  $f$  (by emulating the execution of  $T_{\text{q}}$  using coins  $\omega$  and using  $v_i$  as the answer to the  $i^{\text{th}}$  query, regardless of the identity of this query).<sup>9</sup> ■

**On the operation of sample-based one-sided error testers.** When ignoring the computational complexity aspect, the operation of any sample-based *one-sided error* tester is totally determined by its sample: Being sample-based, this tester has no control over its access to the function, and having one-sided error it has no real control on its decision; that is, without loss of generality, the tester accepts if and only if the labeled sample that it has obtained is consistent with some function that has the property. This is the case because in case of consistency it must accept, whereas in case of inconsistency it better reject (since this may only improve its performance). Note that the consistency check may be computationally infeasible, and so a computationally efficient tester may need to utilize a more conservative check.

### 3 POTs and sample-based testers

The fact that *certain* POTs yield sample-based testers of sublinear sample complexity was proven implicitly in [Gol, Appendix A.2]. The context of that result is of testing graph properties in the bounded-degree model (introduced in [GR02]). Here we present much more general results.

#### 3.1 The main result

In this subsection we prove Theorem 1.1, or more precisely, a slightly stronger version of it. Specifically, we consider POTs that make queries that are *each* almost uniformly distributed in the domain, although their joint distribution may be very dependent. Stated formally, for  $\alpha \in (0, 1]$ , a tester (or a POT) is called  $\alpha$ -fair if, *for each  $i$  and  $j \in D_n$  (and for every  $f: D_n \rightarrow R_n$ ), the  $i^{\text{th}}$  query of the tester (when it accesses the oracle  $f$ ) equals  $j$  with probability at most  $1/\alpha n$* . We stress that here we consider the marginal distribution of the  $i^{\text{th}}$  query. Note that by [GT03, Section 4], we may assume, without loss of generality, that any constant-query POT in the dense graphs model is

<sup>8</sup>This argument ignores the complexity of computing  $q(n, \epsilon)$ .

<sup>9</sup>Indeed, we ignore the computational complexity of selecting  $\omega$  uniformly among all sequences that are consistent with the query-answer sequence.

1-fair.<sup>10</sup> Lastly, we mention that the notion of a fair tester is reminiscent of the notion of a “smooth decoder” (which was studied in [KT00]).

**Theorem 3.1** (from POTs to sample-based testers – basic version): *Suppose that  $\Pi$  has a  $q$ -query POT with detection probability function  $\rho$  that is  $\Omega(1)$ -fair. Then,  $\Pi$  has a sample-based tester with sample complexity  $s(n, \epsilon) = O(n^{1-(1/q)}/\rho(\epsilon)^{2+(3/q)})$ . Furthermore, if the POT has one-sided error, then so does the sample-based tester (and the sample complexity can be reduced to  $O(n^{1-(1/q)}/\rho(\epsilon)^{1+(3/q)})$ ).*

Note that the POT in the premise of Theorem 3.1 implies a standard (query-based) tester of query complexity  $O(1/\rho(\epsilon)^b)$ , where  $b = 2$  in the general case and  $b = 1$  in the one-sided error case. The fairness condition is essential to Theorem 3.1; however, it is satisfied by almost all known POTs. See further discussion in Section 3.2.

**Proof:** As in [Gol, Appendix A.2], the basic idea is that a random sample of the said size is very likely to contain a sequence of queries that are made by the POT on *some* setting of its random coins. Furthermore, this setting is almost uniformly distributed among all possible settings of the random coins. The proof is devoted to actually establishing that the above assertions hold, in particular, given that the algorithm may be adaptive. We note that the situation here is more complex than in [Gol, Appendix A.2], since the property is not necessarily closed under any non-trivial invariance.

Throughout the analysis we fix the POT, denoted  $T$ , and *fix the function, denoted  $f$ , tested by  $T$* . By the hypothesis,  $T$  is  $\alpha$ -fair, for some constant  $\alpha > 0$ . Denoting the randomness complexity of  $T$  by  $r$ , it follows that  $2^r \geq \alpha n$ . We consider all possible random strings  $\omega \in \{0, 1\}^r$ . A key notion is that of the *sequence of queries generated by  $\omega$* , which is defined as *the sequence of queries that  $T$  makes when using randomness  $\omega$  and having access to  $f$* .

We start by presenting a sample-based algorithm that emulates the POT up to an  $O(\delta)$ -deviation, where  $\delta$  is a parameter to be determined later.<sup>11</sup> For  $t = \Theta(n^{1-(1/q)}/\delta^{3/q})$ , the sample-based tester that we wish to construct is given a sequence of  $qt$  labeled examples  $(s_1, f(s_1)), \dots, (s_{qt}, f(s_{qt}))$ , where  $\overline{S} = (s_1, \dots, s_{qt})$  is uniformly distributed in  $[n]^{qt}$ . For each  $i \in [q]$ , we let  $\overline{S}_i$  denote the subsequence  $(s_{(i-1)t+1}, \dots, s_{it})$ .

We say that a pair  $(\omega, \overline{S}) \in \{0, 1\}^r \times [n]^{qt}$  is *good* if the sequence generated by  $\omega$  is in  $(\overline{S}_1, \dots, \overline{S}_q)$  (i.e., for each  $i \in [q]$ , when given oracle access to  $f$  and using coins  $\omega$ , the  $i^{\text{th}}$  query of  $T$  is in  $\overline{S}_i$ ). In such a case we say that  $\overline{S}$  is *good for  $\omega$* , and that  $\omega$  is *good for  $\overline{S}$* . For starters, note that for every fixed  $\omega \in \{0, 1\}^r$ , the probability that a uniformly distributed  $\overline{S} \in [n]^{qt}$  is good for  $\omega$  equals  $(1 - (1 - (1/n))^t)^q \approx (t/n)^q$ . We shall show that, with high probability, a random  $\overline{S} \in [n]^{qt}$  is good for a  $(1 \pm o(1)) \cdot (t/n)^q$  fraction of the settings of  $\omega \in \{0, 1\}^r$ , where we say that a value is  $v \pm \delta$  if it is in the interval  $[v - \delta, v + \delta]$ . This will imply that if the sample-based tester selects at random an  $\omega$  that is good for the sample that it receives, and emulates  $T$  using  $\omega$  as its random coins (while answering  $T$ 's queries by using the corresponding labels), then this tester emulates the POT quite well. We start with the first claim.

**Claim 3.1.1** *Let  $\delta > 0$  and  $\mu = (1 - (1 - (1/n))^t)^q$ . If  $\overline{S}$  is uniformly distributed in  $[n]^{qt}$ , then with probability at least  $1 - \frac{q}{\alpha \delta^2 \mu n}$  the sequence  $\overline{S}$  is good for a  $(1 \pm \delta) \cdot \mu$  fraction of the settings of  $\omega \in \{0, 1\}^r$ .*

<sup>10</sup>As mentioned in the introduction, applying this claim may square the query complexity. Stronger results are obtained (see Theorems 1.2 and 4.3) by considering a more refined parameter (i.e., the number of vertices in the inspected subgraph).

<sup>11</sup>This tester is reminiscent of the notion of a relaxed POT, as presented in Definition A.1, but we shall proceed without any definition regarding this matter.

**Proof:** For every  $\omega \in \{0,1\}^r$ , we denote by  $\zeta_\omega = \zeta_\omega(\bar{S})$  the indicator random variable that is 1 if and only if  $(\omega, \bar{S})$  is good. Recall that  $E[\zeta_\omega] = \Pr_{\bar{S} \in [n]^{qt}}[\zeta_\omega(\bar{S}) = 1] = \mu$ . Letting  $N \stackrel{\text{def}}{=} 2^r$  and  $\bar{\zeta}_\omega \stackrel{\text{def}}{=} \zeta_\omega - \mu$ , the claim asserts that

$$\Pr_{\bar{S} \in [n]^{qt}} \left[ \left| \sum_{\omega \in \{0,1\}^r} \bar{\zeta}_\omega \right| > \delta \cdot \mu N \right] < \frac{q}{\alpha \delta^2 \mu n}. \quad (1)$$

We prove Equation (1) by using Chebyshev's Inequality:

$$\begin{aligned} \Pr_{\bar{S} \in [n]^{qt}} \left[ \left| \sum_{\omega \in \{0,1\}^r} \bar{\zeta}_\omega \right| > \delta \cdot \mu N \right] &< \frac{E \left[ \left( \sum_{\omega \in \{0,1\}^r} \bar{\zeta}_\omega \right)^2 \right]}{(\delta \mu N)^2} \\ &= \frac{1}{\delta^2 \mu^2 N^2} \cdot \sum_{\omega_1, \omega_2 \in \{0,1\}^r} E[\bar{\zeta}_{\omega_1} \bar{\zeta}_{\omega_2}]. \end{aligned} \quad (2)$$

We partition the sum in Equation (2) into two sums separating pairs that generate intersecting sequences of queries from pairs that generate non-intersecting sequences of queries, where the sequences  $(u_1, \dots, u_q)$  and  $(v_1, \dots, v_q)$  intersect if there exists an  $i \in [q]$  such that  $u_i = v_i$ .

We first note that the fraction of intersecting pairs is at most  $q/\alpha n$ , since for any fixed sequence  $\bar{u} = (u_1, \dots, u_q) \in [n]^q$  the sequence of queries generated by a uniformly distributed  $\omega \in \{0,1\}^r$  intersects  $\bar{u}$  with probability at most  $q \cdot (1/\alpha n)$ . On the other hand, for any pair  $(\omega_1, \omega_2)$ , it holds that  $E[\bar{\zeta}_{\omega_1} \bar{\zeta}_{\omega_2}] < E[\zeta_{\omega_1} \zeta_{\omega_2}] \leq E[\zeta_{\omega_1}] = \mu$ . This bound is far from being tight. In particular, as shown next,  $E[\bar{\zeta}_{\omega_1} \bar{\zeta}_{\omega_2}]$  is negative for any *non-intersecting* pair  $(\omega_1, \omega_2)$ . Considering such a non-intersecting pair, note that  $E[\bar{\zeta}_{\omega_1} \bar{\zeta}_{\omega_2}] = \Pr[\zeta_{\omega_1} = \zeta_{\omega_2} = 1] - \mu^2$ , whereas  $\Pr[\zeta_{\omega_1} = \zeta_{\omega_2} = 1] = \mu \cdot \Pr[\zeta_{\omega_2} = 1 | \zeta_{\omega_1} = 1]$  and  $\Pr[\zeta_{\omega_2} = 1 | \zeta_{\omega_1} = 1] = (1 - (1 - (1/n))^{t-1})^q < (1 - (1 - (1/n))^t)^q = \mu$ . Combining all the above facts, we upper bound Equation (2) by

$$\frac{1}{\delta^2 \mu^2 N^2} \cdot \left( \frac{q}{\alpha n} \cdot N^2 \right) \cdot \mu = \frac{q}{\delta^2 \mu \alpha n}$$

and the claim follows. ■

We set  $\delta$  so as to equate the two errors in Claim 3.1.1; that is, we set  $\delta = \frac{q}{\delta^2 \mu \alpha n}$ , which implies  $\mu n = 1/\alpha \delta^3 = \Theta(1/\delta^3)$ . Indeed, avoiding triviality requires that  $\mu n > 1$ . Since  $\mu = (1 - (1 - (1/n))^t)^q < (t/n)^q$ , this implies that  $t > n^{(q-1)/q}$ . Using Claim 3.1.1, the analysis of the sample-based tester outlined above reduces to the following claim.

**Claim 3.1.2** *Let  $G = ((X, Y), E)$  be a bipartite graph such that each vertex in  $X$  has degree  $|E|/|X|$  and at least a  $1 - \delta$  fraction of the vertices in  $Y$  have degree  $(1 \pm \delta) \cdot |E|/|Y|$ , where  $\delta \in [0, 0.5)$ . Then, uniformly selecting a vertex  $y$  in  $Y$ , and uniformly selecting a neighbor of  $y$ , yields a distribution that is  $2\delta$ -close to the uniform distribution on  $X$ .*

(We shall apply Claim 3.1.2 with  $X = \{0,1\}^r$  and  $Y = [n]^{qt}$ , where there is an edge between  $x \in X$  and  $y \in Y$  if  $x$  is good for  $y$ .)

**Proof:** Since each vertex  $x \in X$  has degree  $|E|/|X|$ , if we uniformly select an edge in  $E$  and take its endpoint that belongs to  $X$ , then we get a uniformly selected vertex in  $X$ . In general, any distribution  $D$  over  $E$  induces a distribution  $D_X$  over  $X$  (defined by selecting an edge according

to  $D$ , and taking its  $X$ -endpoint). In particular, we consider the distribution  $D$  over  $E$  described in the claim. That is,  $D$  is defined by uniformly selecting a vertex  $y$  in  $Y$  and then uniformly selecting an edge incident to  $y$ . For every two random variables  $\zeta$  and  $\zeta'$ , and every function  $f$ , the variation distance between  $f(\zeta)$  and  $f(\zeta')$  is upper-bounded by the variation distance between  $\zeta$  and  $\zeta'$ . Therefore, in order to prove the claim, it suffices to upper-bound the variation distance between  $D$  and the uniform distribution over  $E$ , as we do next.

The probability that an edge in  $E$  is selected according to  $D$  is  $\frac{1}{|Y|} \cdot \frac{1}{d_y}$ , where  $d_y$  is the degree of its endpoint  $y \in Y$ . Let  $Y'$  be the set of vertices in  $Y$  that have degree  $(1 \pm \delta) \cdot |E|/|Y|$ , recall that  $|Y'| \geq (1 - \delta) \cdot |Y|$ , and let  $E'$  be the set of edges with an endpoint in  $Y'$ . Then, each edge in  $E'$  is selected with probability  $\frac{1}{|Y|} \cdot \frac{1}{(1 \pm \delta) \cdot |E|/|Y|} = \frac{1}{(1 \pm \delta) \cdot |E|}$ , and  $|E'| \geq (1 - \delta) \cdot |Y| \cdot (1 - \delta) \cdot |E|/|Y| > (1 - 2\delta) \cdot |E|$ . We next use the fact that, for every two random variables,  $\zeta$  and  $\zeta'$ , over a set  $Z$ , and every subset  $S$  of  $Z$ , the variation distance between  $\zeta$  and  $\zeta'$  is upper-bounded by  $\frac{1}{2} \cdot (\sum_{e \in S} |\Pr[\zeta = e] - \Pr[\zeta' = e]| + \max(\Pr[\zeta \notin S], \Pr[\zeta' \notin S]))$ . Therefore, the variation distance between  $D$  and the uniform distribution over  $E$  is upper bounded by  $\frac{1}{2} \cdot (\sum_{e \in E'} \left| \frac{1}{(1 \pm \delta) \cdot |E|} - \frac{1}{|E|} \right| + \max\left(\frac{|Y \setminus Y'|}{|Y|}, \frac{|E \setminus E'|}{|E|}\right))$ , which is at most  $\frac{1}{2} \cdot (|(1 - \delta)^{-1} - 1| + 2\delta) < 2\delta$ , as required. ■

**The sample-based tester (formalized):** On input parameters  $n, \epsilon$  and a sequence of labeled samples  $((s_1, v_1), \dots, (s_{qt}, v_{qt}))$ , where  $t = \Theta(n^{1-(1/q)}/\rho(\epsilon)^{3/q})$ , the tester proceeds as follows, with the intention of emulating a single execution of  $T$  with statistical deviation of at most  $\rho(\epsilon)/3$ .

1. The tester selects uniformly  $\omega \in \{0, 1\}^r$  such that  $\omega$  is good for  $\bar{S}$ , where  $\bar{S} = (s_1, \dots, s_{qt})$ .

If no such  $\omega$  exists, then the tester halts and outputs 1 with probability  $c$  (and outputs 0 otherwise), where  $c$  is the threshold probability of  $T$ . (This bad event occurs with probability at most  $\rho(\epsilon)/10$ .)

(The next step is executed only if  $\omega$  was selected in Step 1.)

2. Let  $(i_1, \dots, i_q) \in [qt]$  be such that  $(s_{i_1}, \dots, s_{i_q})$  is the sequence of queries generated by  $\omega$ .

The tester emulates  $T$  using randomness  $\omega$ , while using  $v_{i_j}$  as the answer to the  $j^{\text{th}}$  query, and outputs whatever  $T$  does.

Note that for our choice of  $t = \Theta(n^{1-(1/q)}/\rho(\epsilon)^{3/q})$ , it holds that  $\mu n \approx (t/n)^q \cdot n = \Theta(1/\rho(\epsilon)^3)$ , and so (by Claims 3.1.1 and 3.1.2) this emulation of  $T$  deviates from a perfect one by at most  $2\delta = \sqrt[3]{q/\alpha\mu n}$ , which is smaller than  $\rho(\epsilon)/3$  (by an adequate choice of  $t$ , which depends on the constant  $q/\alpha$ ). The final sample-based tester is obtained by running  $O(1/\rho(\epsilon)^b)$  copies of the above emulation, where  $b = 2$  for the general case and  $b = 1$  for the one-sided error case. Details follow.

On input  $n, \epsilon$  and  $((s_1, v_1), \dots, (s_{qt'}, v_{qt'}))$ , where  $t' = \Theta(t/\rho(\epsilon)^b) = \Theta(n^{1-(1/q)}/\rho(\epsilon)^{b+(3/q)})$ , the tester runs  $O(1/\rho(\epsilon)^b)$  copies of the above emulation, when using the subsequence of labeled samples  $((s_{(i-1)t+1}, v_{(i-1)t+1}), \dots, (s_{it}, v_{it}))$ , in the  $i^{\text{th}}$  emulation. In the general case the tester outputs 1 if and only if at least  $c - \rho(\epsilon)/2$  of the emulations returned 1, whereas in the one-sided error case the tester outputs 1 if and only if all the emulations returned 1. ■

### 3.2 On the fairness condition

The fairness condition is essential to the foregoing proof of Theorem 3.1. Indeed, if there exists  $i \in [n]$  such that the given POT always makes the query  $i$ , then we have no chance to emulate it while using a random sample of size  $o(n)$ . One might conjecture that *any property that has a  $q$ -query POT, also has a  $q$ -query POT that is  $\Omega(1)$ -fair*; but as shown next this conjecture is false.

**Proposition 3.2** (POTs do not necessarily imply fair POTs): *There exists a property  $\Pi$  that has a (one-sided error) three-query POT with linear detection probability, but has no constant-query POT that is  $\Omega(1)$ -fair.*

Proposition 3.2 does not mean that  $\Pi$  has no sample-based tester of sublinear complexity; in fact, the property used in the following proof does have such a tester. We shall show later (see Proposition 3.3) that there exists a property  $\Pi$  that has a two-query POT but does not have a sample-based tester of sublinear complexity.<sup>12</sup>

**Proof:** The proof is by a reduction to results regarding MA-POTs, a notion introduced by Gur and Rothblum [GR13]. Loosely speaking, an MA-POT is a POT in which the tester obtains an explicit (short) proof (or witness)  $w$  in addition to oracle access to the input function  $f$ . It is required that for every  $f \in \Pi$  there exists a witness  $w_f$  (of the designated length) such that on input  $w_f$  and oracle access to  $f$  the tester always accepts, whereas if  $f$  is not in  $\Pi$ , then, no matter which  $w$  is provided as an alleged witness, the tester rejects  $f$  with probability at least  $\varrho(\delta_\Pi(f))$ , for a monotonically non-decreasing (detection probability) function  $\varrho$  (as in Definition 2.2 for the case  $c = 1$ ). We stress that the tester has free access to the witness, and the query complexity only accounts for the access to  $f$ .

Gur and Rothblum [GR13] showed a  $q$ -query MA-POT that utilizes proofs of logarithmic length for a class of Boolean functions having no (standard) property tester of query complexity  $n^{1-\Omega(1/\sqrt{q})}$ .<sup>13</sup> For our purposes, we shall use a simpler two-query MA-POT that utilizes proofs of logarithmic length for a class of Boolean functions having no (standard) property tester of query complexity  $\Omega(\sqrt{n})$ . Such a result is mentioned in [FGL13],<sup>14</sup> but we shall present an alternative one. In any case, let us detail the reduction first.

**Claim 3.2.1** *Let  $\Pi = \bigcup_{n \in \mathbb{N}} \Pi_n$ , where  $\Pi_n$  contains functions from  $[n]$  to  $\{0, 1\}^{\ell(n)}$ . Suppose that  $\Pi$  has a  $q$ -query MA-POT with proof complexity  $\ell(n)$ , but has no standard tester of query complexity that only depends on the proximity parameter. Then, there exists a property  $\Pi'$  that has a  $(q + 1)$ -query (one-sided error) POT, but no constant-query POT that makes  $\Omega(1)$ -fair queries (not even one with two-sided error). Furthermore,  $\Pi' = \bigcup_{n \in \mathbb{N}} \Pi'_{n+1}$  such that  $\Pi'_{n+1}$  contains functions from  $\{0, 1, \dots, n\}$  to  $\{0, 1\}^{\ell(n)}$ , and the detection probability of the POT is  $\varrho'(\delta) \geq \min(\varrho(\delta/2), \text{poly}(\delta))$ , where  $\varrho$  is the detection probability of the MA-POT.*

**Proof:** We define  $\Pi'_{n+1}$  such that  $f' \in \Pi'_{n+1}$  if there exists  $f \in \Pi_n$  such that  $f'(i) = f(i)$  for every  $i \in [n]$  and  $f'(0) = w$  for some witness  $w$  that makes the guaranteed MA-POT accept  $f$  with probability 1. Note that at least one such  $w$  exists.

We first observe that  $\Pi'$  has no constant-query POT that makes  $\Omega(1)$ -fair queries (not even one with two-sided error). The reason is that an  $\Omega(1)$ -fair POT may only query the function  $f'$  at 0 with probability  $O(1/n)$ , whereas, for some monotonically non-decreasing function  $F$ , it must reject functions that are  $\epsilon$ -far from  $\Pi'$  with probability that is  $F(\epsilon)$  smaller than the probability that it accepts functions in  $\Pi'$ . Hence, if such a constant-query POT existed, it could be modified to a POT that never queries  $f'$  at 0 and still maintains the foregoing behavior; that is, for any constant  $\epsilon > 0$ , this POT distinguishes between functions that are  $\epsilon$ -far  $\Pi$  and functions in  $\Pi$  with

<sup>12</sup>Indeed, Proposition 3.3 implies Proposition 3.2, alas it uses a more contrived property (which consists of functions having a huge range).

<sup>13</sup>We refer to [GR13, Theorem 3.1] and note that its proof implicitly provides a constant-query POT of detection probability  $\varrho(\epsilon) = \exp(-1/\epsilon)$ . (Actually, the detection probability is  $\Omega(\epsilon)$  if  $\epsilon > 1/\log n$  (and at least  $1/n$  otherwise).)

<sup>14</sup>We refer to their observation that the property  $\Pi = \{uu^Rvv^R : u, v \in \{0, 1\}^*\}$  has a constant-query MA-POT of logarithmic proof complexity, whereas any standard tester for  $\Pi$  requires  $\Omega(\sqrt{n})$  queries [AKNS01].

a gap of at least  $F(\epsilon) - O(1/n) > F(\epsilon)/2$ . But this would have yielded a standard tester of query complexity  $O(1/F(\epsilon))$ , in contradiction to the hypothesis.

We now present a  $(q+1)$ -query POT for  $\Pi'$ . This POT always queries  $f'$  at 0, and invokes the MA-POT while providing it with oracle access to the corresponding  $f = f'_{[n]}$  (i.e.,  $f'$  restricted to  $[n]$ ) and a witness that equals  $f'(0)$ . Note that this POT accepts each  $f' \in \Pi'$  with probability 1 (by the definition of  $f'(0)$  and the one-sided error of the MA-POT). If  $f' \notin \Pi'$ , then there are two cases with respect to  $f = f'_{[n]}$ :

1. If  $f \notin \Pi$ , then, regardless of the value of  $f'(0)$ , the MA-POT (and hence the constructed POT) rejects with probability at least  $\varrho(\delta_{\Pi}(f)) \geq \varrho(\delta_{\Pi'}(f')/2)$ , since  $(n+1) \cdot \delta_{\Pi'}(f') \leq n \cdot \delta_{\Pi}(f) + 1 \leq 2n \cdot \delta_{\Pi}(f)$ .
2. If  $f \in \Pi$ , then  $w = f'(0)$  is not a witness that makes the MA-POT accept  $f$  with probability 1, but rather with probability strictly smaller than 1. Letting  $\rho$  denote the randomness complexity of the MA-POT, this probability is at most  $1 - 2^{-\rho(n)} \leq 1 - (1/\text{poly}(n))$ , since we may assume without loss of generality that  $\rho(n) = O(\log n)$  (cf., e.g., [GR13, Lemma 4.6]). Observe that in this case  $\delta_{\Pi'}(f') = \frac{1}{n+1}$ , and so the detection probability is polynomial in the distance.

By combining the two cases, we get that the rejection probability function  $\varrho'$  of the resulting POT satisfies  $\varrho'(\epsilon) \geq \min(\varrho(\epsilon/2), \text{poly}(\epsilon))$ . ■

**Claim 3.2.2** *Let  $\Pi = \bigcup_{n \in \mathbb{N}} \Pi_n$ , where  $\Pi_n$  contains functions from  $[n]$  to  $[n]$  such that  $f \in \Pi_n$  if and only if there exists  $i \in [\lfloor n/2 \rfloor]$  such that for every  $j \in [\lfloor n/2 \rfloor]$  it holds that  $f(\lfloor n/2 \rfloor + j) = f(\text{sh}_i(j))$ , where  $\text{sh}_i(j) \stackrel{\text{def}}{=} ((j+i) \bmod \lfloor n/2 \rfloor) + 1$ ; that is,  $(f(\lfloor n/2 \rfloor + 1), \dots, f(2\lfloor n/2 \rfloor))$  is a cyclically shifted version of  $(f(1), \dots, f(\lfloor n/2 \rfloor))$ . Then,  $\Pi$  has a two-query MA-POT with proof complexity  $\log_2 n$ , but any standard tester for  $\Pi$  has query complexity  $\Omega(\sqrt{n})$ .<sup>15</sup>*

**Proof:** Let  $m \stackrel{\text{def}}{=} \lfloor n/2 \rfloor$ . The two-query MA-POT gets  $i \in [m]$  as a witness, selects  $j$  uniformly in  $[m]$  and compares  $f(m+j)$  to  $f(\text{sh}_i(j))$ , by making the corresponding queries. This MA-POT has linear detection probability. Turning to the lower bound, we show that a machine of query complexity  $o(\sqrt{n})$  cannot distinguish a function uniformly selected in  $\Pi_n$  from a totally random function. Towards this end, we consider a matching of  $[m+1, 2 \cdot m]$  to  $[1, m]$  such that  $m+j$  is matched to  $\text{sh}_i(j)$ . The point is that in the former case (i.e., of  $f \in \Pi$ ), as long as the queried locations contain no matched pair, the answers are uniformly distributed in  $[n]$  (just as in the case of a totally random function). ■

Combining the two claims, the proposition follows. Note that the property that we obtained (i.e., the one satisfying the proposition) is a property of functions with a range that equals their domain. ■

**Proposition 3.3** (POTs do not necessarily imply testers with sublinear sample complexity): *There exists a property  $\Pi$  that has a one-sided error two-query POT with linear detection probability, but has no sample-based tester of sublinear sample complexity.*

<sup>15</sup>This lower bound is tight, since  $\Pi$  has a standard tester of query complexity  $O(\sqrt{n}/\epsilon^2)$ . The key observation is that  $2\sqrt{n}$  queries suffice for determining  $i$  such that  $(f(\lfloor n/2 \rfloor + 1), \dots, f(2\lfloor n/2 \rfloor))$  is an  $i$ -shifted version of  $(f(1), \dots, f(\lfloor n/2 \rfloor))$ , and within the stated complexity we can detect that for each  $i \in [\lfloor n/2 \rfloor]$  the sequence  $(f(\lfloor n/2 \rfloor + 1), \dots, f(2\lfloor n/2 \rfloor))$  is  $\epsilon$ -far from the  $i$ -shifted version of  $(f(1), \dots, f(\lfloor n/2 \rfloor))$ . (In fact, a random sample of  $\tilde{O}(\sqrt{n}/\epsilon^2)$  points can be used towards the same end, yielding a sample-based tester.)

Proposition 3.3 is proved by using a class of functions with a huge range (i.e.,  $|R_n| = \exp(n)$ ). In a preliminary version of this work, we wondered whether the result holds also for a class of functions with significantly smaller range (e.g.,  $|R_n| = n$  or maybe even  $|R_n| = 2$ ). Subsequently, Fischer *et al.* [FGL14, Corollary 1.9] proved that the result of Proposition 3.3 does not hold for Boolean functions (i.e., for  $|R_n| = 2$ ). Specifically, they showed that any property of Boolean functions that has a one-sided error two-query POT also has a sample-based tester of sample complexity  $s(n, \epsilon) = O(n^{2/3}/\epsilon)$ .

**Proof:** Starting from any hard to test property  $\Pi'$  (cf., e.g., [GGR98, Section 4.1]), we consider the property  $\Pi$  such that  $f \in \Pi_{n+1}$  if and only if there exists  $f' \in \Pi'_n$  such that  $f(i) = f'(i)$  for every  $i \in [n]$  and  $f(0) = \langle f' \rangle$ , where  $\langle f' \rangle$  denotes the full description of  $f': [n] \rightarrow R'_n$ . Indeed, we use  $D_{n+1} = \{0, 1, \dots, n\} = \{0\} \cup D'_n$  and  $R_{n+1} = R'_n \cup (R'_n)^n$ . (We could have presented the proof by using the connection to MA-POTs, as done in the proof of Proposition 3.2, but prefer to present a direct argument.)<sup>16</sup>

We first show that  $\Pi$  has a two-query POT with one-sided error and detection probability  $\varrho(\epsilon) \geq \epsilon/2$ . Given access to a function  $f: D_{n+1} \rightarrow R_{n+1}$ , we query  $f$  at 0 and at a uniformly distributed  $i \in [n]$ , obtaining the values  $v$  and  $u$ , which may be assumed to reside in  $(R'_n)^n$  and  $R'_n$  respectively (or else we reject). We accept if and only if  $v$ , viewed as a function from  $[n]$  to  $R'_n$ , is in  $\Pi'_n$  and  $u$  equals the  $i^{\text{th}}$  element in the  $n$ -long sequence  $v$ .<sup>17</sup> This machine accepts any  $f \in \Pi_{n+1}$  with probability 1. On the other hand, suppose that  $f$  is  $\epsilon$ -far from  $\Pi_{n+1}$ , and let  $f'$  denote the restriction of  $f$  to  $[n]$ . We first discard the case that  $\epsilon = 1/(n+1)$ , noting that in this case the machine rejects with probability at least  $\epsilon$ . Now, using  $\epsilon \geq 2/(n+1)$ , we infer that  $f'$  is  $\epsilon'$ -far from  $\Pi'_n$ , where  $\epsilon' = ((n+1)\epsilon - 1)/n > \epsilon/2$ . We may assume, w.l.o.g., that  $f(0) \in \Pi'_n$ , which implies that there are at least  $\epsilon n/2$  indices  $i \in [n]$  such that  $f(i) = f'(i)$  differs from the  $i^{\text{th}}$  element in the  $n$ -long sequence  $f(0)$ . In this case, the machine rejects with probability at least  $\epsilon/2$ .

Finally, we claim that  $\Pi$  has no sample-based tester of sublinear sample complexity. The reason is that a sublinear sized sample is unlikely to yield the value of  $f(0)$ , whereas without this value testing is hard (i.e., requires a linear number of queries, let alone samples). ■

**A sufficient condition for fairness.** While Propositions 3.2 and 3.3 assert that not every POT can be transformed into a fair one, we next state a sufficient condition for such a transformation. We warn, however, that this condition is not necessary. In fact, any POT in the bounded-degree graph model can be transformed into a fair one (cf. [GR11, Claim 5.5.2]), although the condition stated next does not hold.<sup>18</sup>

**Proposition 3.4** (invariance under 1-transitivity implies 1-fairness): *Suppose that the property  $\Pi$  is invariant under a 1-transitive permutation group that acts on its domain; that is, for every  $n \in \mathbb{N}$  and  $i, j \in D_n$ , there exists a permutation  $\pi: D_n \rightarrow D_n$  such that  $\pi(i) = j$  and for every  $f: D_n \rightarrow R_n$  it holds that  $f \in \Pi_n$  if and only if  $f \circ \pi \in \Pi_n$ , where  $(f \circ \pi)(x) = f(\pi(x))$ . If  $\Pi$  has a nonadaptive  $q$ -query POT, then it has one that is 1-fair, and if the former has one-sided error then so does the latter.*

<sup>16</sup>Such an alternative presentation would proceed by showing that  $\Pi$  has a one-query MA-POT of proof complexity  $n \log |R'_n|$ .

<sup>17</sup>This machine corresponds to an MA-POT for  $\Pi'$ , which is given  $\langle f' \rangle$  as a witness, and makes a single query obtaining the value  $f'(i)$  (and acting accordingly).

<sup>18</sup>Graph properties in this model are not invariant under any 1-transitive permutation group that acts on their domain, and (in general) the corresponding POTs are adaptive [RS06].

For constant-size range (i.e.,  $|R_n| = O(1)$ ), non-adaptivity can be obtained at the cost of decreasing the detection probability by a factor that is exponential in  $q$ . (This can be done by guessing the answers at random beforehand, determining the corresponding queries and making them non-adaptively, outputting the verdict of the original POT if all guesses turn out to be correct, and outputting 1 with probability  $c$  otherwise (i.e., if any of these guesses turns out to be wrong).)

**Proof:** Consider the group  $G_n$  generated by the permutations guaranteed in the hypothesis. Given a  $q$ -query POT  $T$  for  $\Pi$ , consider the machine  $T'$  that, when given oracle access to a function  $f: D_n \rightarrow R_n$ , selects  $\pi \in G_n$  uniformly at random, and invokes  $T$  when providing it with oracle access to  $f \circ \pi$ ; that is, query  $j$  is answered by querying  $f$  at  $\pi(j)$ . Then,  $T'$  is a POT for  $\Pi$ , since  $\delta_\Pi(f \circ \pi) = \delta_\Pi(f)$  by the invariance of  $\Pi$  under the permutations in  $G_n$ . On the other hand, by the 1-transitivity of  $G_n$ , a random  $\pi \in G_n$  maps each  $i \in D_n$  uniformly over  $D_n$  (i.e.,  $\Pr_{\pi \in G_n}[\pi(i) = j] = 1/n$  for every  $j \in D_n$ ). Hence, each of the queries of  $T'$  is uniformly distributed in  $D_n$ , where here we rely on the non-adaptivity of  $T$ . ■

We note that non-adaptivity is essential to the proof of Proposition 3.4. Consider, for example, the property  $\Pi_n = \{f_i: i \in [n]\}$  such that  $f_i: \mathbb{Z}_n \rightarrow \mathbb{Z}_n$  satisfies  $f_i(j) = j + i \pmod n$ . This property is 1-transitive, but is only invariant under cyclic permutations (i.e., permutations of the form  $\pi_s(i) = i + s$ ). Now consider a three-query adaptive POT  $T$  that, uniformly selects  $r \in \mathbb{Z}_n$ , queries its oracle  $f$  on  $0, r$ , and  $v = r - f(r)$ , and accepts if and only if  $f(0) = f(r) - r$ . (Indeed,  $T$  ignores the third answer, and its analysis reduces to observing that it accepts  $f$  on coins  $r$  if and only if  $f(r) = r + f(0)$ .) Then  $T'$ , as defined in the foregoing proof, select  $s \in \mathbb{Z}_n$  uniformly at random, and queries  $f = f_i$  on  $\pi_s(0), \pi_s(r)$  and  $\pi_s(v)$ , whereas  $\pi_s(v) = \pi_s(r - f_i(\pi_s(r))) = (r - ((r+s)+i)) + s = i$ , which strongly violates the fairness condition (since the third query is totally determined by  $f_i$ ). Of course this does not mean that  $\Pi_n$  has no POT that is non-adaptive; in fact, this  $\Pi_n$  has a sample-based POT that uses two samples (i.e., given  $(r_1, f(r_1))$  and  $(r_2, f(r_2))$ , the POT accepts if and only if  $f(r_1) - f(r_2) = r_1 - r_2$ ).<sup>19</sup>

### 3.3 A generalization (of Theorem 3.1)

For any  $k \in [q]$ , we consider POTs that make  $q$  queries such that any  $k$  of these queries are (almost) uniformly distributed in  $D_n^k$ . Indeed, the case of  $k = 1$  is handled in Theorem 3.1 and the case of  $k = q$  is straightforward, but here we are interested in the intermediate cases. For example, the basic test that underlies the Linearity Tester of [BLR93] selects three elements in the domain such that any pair are uniformly and independently distributed. In general, for  $\alpha \in (0, 1]$ , we say that a POT is  $(k, \alpha)$ -fair if, for every  $1 \leq i_1 < i_2 < \dots < i_k \leq q$  and  $e_1, \dots, e_k \in D_n$  (and for every  $f: D_n \rightarrow R_n$ ), with probability at most  $1/\alpha n^k$  it holds that for every  $j \in [k]$  the  $i_j^{\text{th}}$  query of the tester equals  $e_j$  (when it accesses the oracle  $f$ ).

**Theorem 3.5** (Theorem 3.1, generalized): *Suppose that  $\Pi$  has a  $q$ -query POT with detection probability  $\varrho$  that is  $(k, \Omega(1))$ -fair. Then,  $\Pi$  has a sample-based tester of sample complexity  $s(n, \epsilon) = \max(O(n^{1-(k/q)}/\varrho(\epsilon)^{b+(3/q)}), O(\varrho(\epsilon)^{-(3+b)}))$ , where  $b = 2$  in the general case and  $b = 1$  if POT has one-sided error. Furthermore, if the POT has one-sided error, then so does the sample-based tester.*

Applying Theorem 3.5 to the 3-query POT for Linearity (cf. [BLR93]), we obtain a (one-sided error) sample-based tester of sample complexity  $s(n, \epsilon) = \text{poly}(1/\epsilon) \cdot n^{1/3}$ . Note that this sample

---

<sup>19</sup>This POT has detection probability  $\varrho(\delta) = \delta$ . This is shown by noting that if this POT accepts with probability  $p$ , then there is a choice of  $r_1$  for which it accepts with probability at least  $p$ , and it follows that  $f$  is  $(1-p)$ -close to  $\Pi_n$  (just as in the analysis of the foregoing  $T$ ).

complexity is optimal as an emulation of this POT, since a sample of  $o(n^{1/3})$  group elements is unlikely to contain a triplet that sums-up to zero, but more efficient sample-based tester can be obtained directly (see Section 5).

**Proof:** We proceed as in the proof of Theorem 3.1, while focusing on strengthening Claim 3.1.1 under the stronger fairness hypothesis made here. Specifically, we partition the sum in Equation (2) into  $k + 1$  sums separating pairs according to the number of intersections that occur in these pairs. More specifically, we say that the sequences  $(u_1, \dots, u_q)$  and  $(v_1, \dots, v_q)$  are  $j$ -intersecting if  $|\{i \in [q] : u_i = v_i\}| = j$ . (Indeed, the proof of Claim 3.1.1 only distinguishes 0-intersecting pairs from intersecting pairs (i.e., pairs that  $j$ -intersect for some  $j \in [q]$ .) Letting  $I_j$  denote the set of  $j$ -intersecting pairs, we have:

$$\begin{aligned} & \sum_{\omega_1, \omega_2 \in \{0,1\}^r} \mathbb{E}[\bar{\zeta}_{\omega_1} \bar{\zeta}_{\omega_2}] \\ &= \sum_{(\omega_1, \omega_2) \in I_0} \mathbb{E}[\bar{\zeta}_{\omega_1} \bar{\zeta}_{\omega_2}] + \sum_{j \in [k]} \sum_{(\omega_1, \omega_2) \in I_j} \mathbb{E}[\bar{\zeta}_{\omega_1} \bar{\zeta}_{\omega_2}] + \sum_{(\omega_1, \omega_2) \in \bigcup_{j>k} I_j} \mathbb{E}[\bar{\zeta}_{\omega_1} \bar{\zeta}_{\omega_2}]. \end{aligned} \quad (3)$$

Recall that the contribution of the 0-intersecting pairs is negative (see the proof of Claim 3.1.1). Now, for every  $j \in [k]$ , the fraction of  $j$ -intersecting pairs is at most  $\binom{q}{j}/\alpha n^j$ , whereas each  $j$ -intersecting pair  $(\omega_1, \omega_2)$  contributes  $\mu \cdot (1 - (1 - (1/n))^{t-1})^{q-j} < \mu^2 \cdot (1 - (1 - (1/n))^t)^{-j}$ , since  $\mathbb{E}[\bar{\zeta}_{\omega_1} \bar{\zeta}_{\omega_2}] < \Pr[\zeta_{\omega_1} = \zeta_{\omega_2} = 1]$  and  $\Pr[\zeta_{\omega_2} = 1 | \zeta_{\omega_1} = 1] = (1 - (1 - (1/n))^{t-1})^{q-j}$ . Similarly, the fraction of pairs in  $\bigcup_{j>k} I_j$  is at most  $\binom{q}{k}/\alpha n^k$ , but we can only upper bound the contribution of each such pair by  $\mu$ . Since  $t = s(n, \epsilon)/q = o(n)$ , we have  $(1 - (1 - (1/n))^{t-1})^{-j} < 2 \cdot (t/n)^{-j}$ , and so Equation (3) is upper bounded by

$$\begin{aligned} & \sum_{j \in [k]} \frac{\binom{q}{j} N^2}{\alpha n^j} \cdot 2\mu^2 \cdot (t/n)^{-j} + \frac{\binom{q}{k} N^2}{\alpha n^k} \cdot \mu \\ & < 2^{q+1} \cdot N^2 \cdot \left( \frac{\mu^2}{\alpha t} + \frac{\mu}{\alpha n^k} \right). \end{aligned} \quad (4)$$

where  $N = 2^r$  (as in the proof of Claim 3.1.1). Combining Equation (2) and Equation (4), it follows that

$$\begin{aligned} \Pr_{\bar{\omega} \in [n]^{qt}} \left[ \left| \sum_{\omega \in \{0,1\}^r} \bar{\zeta}_{\omega} \right| > \delta \cdot \mu N \right] &< \frac{2^{q+1}}{\delta^2 \mu^2} \cdot \left( \frac{\mu^2}{\alpha t} + \frac{\mu}{\alpha n^k} \right) \\ &= \frac{2^{q+1}}{\delta^2 \alpha t} + \frac{2^{q+1}}{\delta^2 \mu \alpha n^k}. \end{aligned} \quad (5)$$

Wishing to upper bound each of the terms in Equation (5) by  $\delta/2$ , we get  $t \geq 2^{q+2}/\alpha \delta^3$  and  $\mu n^k \geq 2^{q+2}/\alpha \delta^3$ . By setting  $t = \max(O(n^{1-(k/q)}/\delta^{3/q}), O(1/\delta^3))$ , both inequalities are satisfied. Continuing as in the proof of Theorem 3.1, the current theorem follows.  $\blacksquare$

**Remark 3.6** (applicability to standard testers): *The ideas underlying the proofs of Theorem 3.1 and 3.5 can be applied also to standard testers. In this case the query complexity  $q$  is not constant, but rather depends on the proximity parameter  $\epsilon$  and possibly also on the size parameter  $n$ . On the other hand, the detection probability (represented by the function  $\varrho$ ) is a constant (i.e.,  $1/3$ ). The most appealing version refers to the case that the query complexity only depends on  $\epsilon$ , in which case we get:*

Suppose that  $\Pi$  has a  $\Omega(1)$ -fair tester of query complexity  $q: [0, 1] \rightarrow \mathbb{N}$  that only depends on the proximity parameter. Then,  $\Pi$  has a sample-based tester of sample complexity  $s(n, \epsilon) = O(n^{1-1/q(\epsilon)})$ . Furthermore, if the original tester has one-sided error, then so does the sample-based tester.

Hence, sample-based testers of sublinear sample complexity exist for many natural properties.

## 4 Sample-based testers in the dense graphs model

In this section, we consider testing graph properties in the adjacency matrix model (a.k.a. the dense graphs model). Our aim here is to present results that are either quantitatively or qualitatively stronger than the general result presented in the previous section (specifically, in Section 3.1). The quantitative improvements are due to either referring to an alternative complexity measure or to a restriction on a specific class of problems (see Sections 4.1 and 4.2, respectively). The qualitative improvements, presented in Section 4.2, refer to the computational complexity of the testers, which is ignored in most of the current work.

In the adjacency matrix model a graph  $G = (V, E)$  is represented by the Boolean function  $g: \binom{V}{2} \rightarrow \{0, 1\}$  such that  $g(u, v) = 1$  if and only if  $u$  and  $v$  are adjacent in  $G$  (i.e.,  $\{u, v\} \in E$ ). Thus, the domain size  $n$  equals  $\binom{|V|}{2}$ , and the distance between graphs is measured in terms of their aforementioned representation (i.e., as the fraction of (the number of) unordered vertex pairs on which they differ over  $n$ ). Here we only consider *graph properties*, which are sets of graphs that are closed under isomorphism; that is,  $\Pi$  is a **graph property** if for every graph  $G = (V, E)$  and every permutation  $\pi$  of  $V$  it holds that  $G \in \Pi$  if and only if  $\pi(G) \in \Pi$ , where  $\pi(G) \stackrel{\text{def}}{=} (V, \{\{\pi(u), \pi(v)\} : \{u, v\} \in E\})$ .

We shall sometimes view  $E$  as a set of ordered pairs (rather than a set of unordered ones). Ditto regarding the sequence of samples (i.e., vertex pairs).

### 4.1 General results

The starting point of our general study of sample-based testers in the dense graphs model is a transformation of standard testers (in this model) into “canonical” ones [GT03, Theorem 2]. In fact, we shall use part of this transformation, which yields “quasi-canonical” testers as defined next. (This definition introduces an alternative complexity measure – the number of vertices in the inspected subgraph – and enables quantitative improvements by reference to that measure.)

**Definition 4.1** (quasi-canonical testers): *A quasi-canonical tester of vertex complexity  $\nu: \mathbb{N} \times [0, 1] \rightarrow \mathbb{N}$  for a graph property  $\Pi$  is a tester of  $\Pi$  (in the dense graphs model) that operates as follows:*

1. *On input parameters  $n = \binom{|V|}{2}$  and  $\epsilon$ , the tester selects uniformly a set of  $\nu(n, \epsilon)$  vertices, denoted  $S$ , and queries all the vertex pairs in it (i.e., all pairs in  $\binom{S}{2}$ ).*
2. *Let  $H$  be the unlabeled induced graph that corresponds to the answers obtained in Step 1. The tester makes its decision based solely on  $n, \epsilon$  and  $H$ , and possibly on fresh coin tosses (but not on the coin tosses used to select the sample  $S$ ).*

Definition 4.1 differs from the definition of canonical testers in [GT03] only in that Step 2 is allowed to be randomized.

**Theorem 4.2** (Combining Lemma 4.1 and Claim 4.3 in [GT03]): *Let  $\Pi$  be a graph property and suppose that  $\Pi$  has a tester with query complexity  $q: \mathbb{N} \times [0, 1] \rightarrow \mathbb{N}$  (in the dense graphs model). Then  $\Pi$  has a quasi-canonical tester with vertex complexity  $2q$ .*

(The canonical tester of [GT03] is obtained by transforming any quasi-canonical tester into a canonical one, while increasing the vertex complexity by a constant factor.<sup>20</sup> We avoid this additional transformation, which is presented in [GT03, Lemma 4.4], because we do not need it.)

We show that quasi-canonical tester yield sample-based testers of sample complexity that is lower than what we obtained for general (fair) testers of the corresponding query complexity (which is quadratic in the vertex complexity).

**Theorem 4.3** (from quasi-canonical testers to sample-based testers): *Let  $\Pi$  be a graph property and suppose that  $\Pi$  has a quasi-canonical tester of vertex complexity  $\nu: \mathbb{N} \times [0, 1] \rightarrow \mathbb{N}$ . Then,  $\Pi$  has a sample-based tester of sample complexity  $s(n, \epsilon) = O(n^{1-1/(\nu(n, \epsilon)-1)})$ .*

In particular, if  $\nu(n, \epsilon) = \nu'(\epsilon)$  does not depend on  $n$ , then for every constant  $\epsilon$  we obtain sublinear sample complexity.

**Proof:** A random sample of size  $s = s(n, \epsilon)$  corresponds to selecting each sample point (i.e., a vertex pair) with probability  $p = s/n$ , independently of all other choices.<sup>21</sup> Recall that  $n = \binom{|V|}{2} \approx |V|^2/2$ . Call a set of vertices  $S$  good for the sample if the sample contains all vertex pairs that correspond to this set  $S$  (i.e., for every two distinct vertices  $u, v \in S$ , it holds that the pair  $\{u, v\}$  is in the sample). Then, the expected number of vertex sets of size  $\nu = \nu(n, \epsilon)$  that are good for a random sample (selected as above), denoted  $\mu$ , is given by

$$\begin{aligned} \binom{|V|}{\nu} \cdot p^{\binom{\nu}{2}} &> (\sqrt{2n}/\nu)^\nu \cdot p^{\nu \cdot (\nu-1)/2} \\ &= ((2n/\nu^2) \cdot (s/n)^{\nu-1})^{\nu/2}, \end{aligned}$$

which is greater than 1 if  $s > 2 \cdot n^{1-1/(\nu-1)}$ . As shown next, with probability at least 0.99, a random sample of  $O(n^{1-1/(\nu-1)})$  vertex pairs contains a good set of  $\nu$  vertices, and each such good set is uniformly distributed among all vertex sets of size  $\nu$ .

For every set of  $\nu$  vertices,  $S$ , let  $\zeta_S$  denote the indicator random variable representing whether  $S$  is good for a random sample, and let  $\bar{\zeta}_S = \zeta_S - \mathbb{E}[\zeta_S]$ . Letting  $N \stackrel{\text{def}}{=} \binom{|V|}{\nu}$  and  $\mu = \mathbb{E}[\zeta_S] = p^{\nu \cdot (\nu-1)/2}$ , the claim follows by applying Chebyshev's Inequality: Starting with

$$\Pr \left[ \sum_{S \in \binom{V}{\nu}} \zeta_S > 0 \right] < \frac{\mathbb{E} \left[ \left( \sum_{S \in \binom{V}{\nu}} \bar{\zeta}_S \right)^2 \right]}{(\mu N)^2}$$

we use the fact that  $\mathbb{E}[\bar{\zeta}_S \bar{\zeta}_T] = 0$  if  $S \cap T = \emptyset$  (and in fact also if  $|S \cap T| = 1$ ). The key observation is that for  $|S \cap T| \geq 1$ , it holds that  $\mathbb{E}[\bar{\zeta}_S \bar{\zeta}_T] < \mathbb{E}[\zeta_S \zeta_T] = p^{\binom{|S|}{2} + \binom{|T|}{2} - \binom{|S \cap T|}{2}}$ , since  $\zeta_S \zeta_T = 1$  if and

<sup>20</sup>This transformation can be avoided in case the original tester had one-sided error (see [GT03, Claim 4.3]).

<sup>21</sup>Actually, the correspondence is not exact, but it suffices for our purposes. See Appendix A.2. The issue at hand here is the difference between picking  $m$  elements uniformly in  $[n]$  and picking each element in  $[n]$  with probability  $m/n$ .

only if the sample contains  $\binom{S}{2} \cup \binom{T}{2}$ . Hence,

$$\begin{aligned}
\sum_{S,T \in \binom{V}{\nu}} \mathbb{E}[\bar{\zeta}_S \bar{\zeta}_T] &= \sum_{i \in [\nu]} \sum_{S,T \in \binom{V}{\nu}: |S \cap T|=i} \mathbb{E}[\bar{\zeta}_S \bar{\zeta}_T] \\
&< \sum_{i \in [\nu]} \binom{|V|}{2\nu-i} \cdot \binom{2\nu-i}{\nu} \cdot \binom{\nu}{i} \cdot p^{2\binom{\nu}{2}-\binom{i}{2}} \\
&< \sum_{i \in [\nu]} O(|V|/\nu)^{2\nu-i} \cdot p^{(2\nu-i) \cdot (\nu-1)/2} \\
&= \exp(\nu) \cdot \sum_{i \in [\nu]} \left( (|V|/\nu) \cdot p^{(\nu-1)/2} \right)^{2\nu-i},
\end{aligned}$$

which is significantly smaller than  $(N\mu)^2 > ((|V|/\nu) \cdot p^{(\nu-1)/2})^{2\nu}$ , since  $|V| \cdot p^{(\nu-1)/2} > \exp(\nu)$ .

Having shown that, with probability at least 0.99, a random sample of  $O(n^{1-1/(\nu-1)})$  vertex pairs contains a good set of  $\nu$  vertices, we note that the subgraph induced by such a set is visible in the sample. Noting that each such good set is uniformly distributed among all vertex sets of size  $\nu$ , it begs to emulate the quasi-canonical tester using this set. Indeed, our sample-based tester selects at random a good set, uniformly among all  $\nu$ -vertex sets that are good for the given sample, and emulates the quasi-canonical tester using the induced subgraph (which is visible in the sample). ■

REV: (29)  
asks about exp notation.  
D: added earlier that  $c^x$  for some constant  $c$

**Quasi-canonical proximity oblivious testers.** Using the same ideas, we can obtain analogous results for proximity oblivious testers (in the dense graphs model). We first need an analogue definition of quasi-canonical POTs. Such a definition is obtained by combining Definitions 2.2 and 4.1. Specifically, a quasi-canonical POT of vertex complexity  $\nu$  is a POT that operates as in Definition 4.1, which means that  $\nu$  is a constant. Analogously to Theorem 4.2, it follows that every  $q$ -query POT yields a quasi-canonical POT of vertex complexity  $2q$  that maintains the performance guarantees of the original POT, where this fact is only implicit in [GT03, Section 4]. Using the ideas underlying the proof of Theorem 4.3, we get the first item in Theorem 1.2, restated next.

**Theorem 4.4** (from quasi-canonical POTs to sample-based testers): *Suppose that  $\Pi$  has a quasi-canonical POT of vertex complexity  $\nu$  and detection probability  $\varrho$ . Then,  $\Pi$  has a sample-based tester of sample complexity  $s(n, \epsilon) = O(n^{1-1/(\nu-1)}/\varrho(\epsilon)^2)$ . Furthermore, if the POT has one-sided error, then so does the sample-based tester (and the sample complexity can be reduced to  $O(n^{1-1/(\nu-1)}/\varrho(\epsilon))$ ).*

(The  $\varrho(\epsilon)^{-b}$  factor, for  $b = 1, 2$ , arises from the need to repeat the resulting “sample-based POT” for an adequate number of times so as to derive a tester.) Note that Theorem 4.4 improves upon Theorem 3.1 whenever we are given a quasi-canonical POT, since the query complexity of such a POT is quadratic in its vertex complexity. In general, we get an improvement whenever we are given a  $q$ -query POT such that its  $q$  queries always refer to at most  $\nu \leq q$  vertices, since in this case the transformation in [GT03, Section 4] yields a quasi-canonical POT of vertex complexity  $\nu$ .

## 4.2 Graph partition problems

Here we refer to the general framework of graph partition problems presented in [GGR98, Section 9]. Each problem in this framework is specified by  $k + k + \binom{k}{2}$  pairs of parameters in  $[0, 1]$ , which specify

lower and upper bounds on the densities of the parts in a  $k$ -partition of the graph as well as on the number of edges within each part and between the parts. For example,  $k$ -colorability is specified by trivial lower and upper bounds for all densities, except for upper bounds (of zero) on the number of edges within each part, which mandate that each part is an independent set. Another example is  $\rho$ -clique, which can be captured by requiring that the first part has density exactly  $\rho$  and the number of edges inside it is  $\approx \rho^2 \cdot n$ . We start with the following generic result, of which the second item in Theorem 1.2 is a special case.

**Theorem 4.5** (sample complexity linear in the number of vertices): *Every graph partition property has a sample-based tester of sample complexity  $s(n, \epsilon, k) = O_k(\sqrt{n}/\epsilon^2) = \tilde{O}(k^4) \cdot \sqrt{n}/\epsilon^2$ . Furthermore, for  $k$ -colorability the bound is  $O(\log k) \cdot \sqrt{n}/\epsilon$ .*

Unfortunately, the tester that we use for proving Theorem 4.5 runs in exponential-time (i.e., in time that is exponential in its sample complexity). We thus view Theorem 4.5 as a starting point for a search for efficient sample-based testers (i.e., testers that run in time that is linear in their sample complexity).

**Proof:** The tester considers all possible  $k$ -partitions of the vertex set of the graph that match the vertex density bounds. (Indeed, there are at most  $k^{|V|}$  such partitions.) For each such partition, the tester tries to find evidence for violation of the bounds regarding edge density. By using  $O(tk^4/\epsilon^2)$  labeled samples, it can estimate the edge density between pairs of vertex sets up to an additive error of  $\epsilon/2k^2$  with error probability at most  $\exp(-t)$ . Setting  $t = O(|V| \log k)$  and applying a union bound, we may assume that all estimates obtained by the tester are within this deviation. Hence, the tester rejects if and only if evidence of violation was found for each of the foregoing  $k$ -partitions of  $V$ . The improved bound for  $k$ -Colorability follows by observing that in this case the violation events refer to existence of certain edges (rather than to deviation from a given edge density).<sup>22</sup> ■

As stated above, Theorem 4.5 is the starting point of our search for sample-based testers of sublinear time complexity. We first restate and prove the first item in Theorem 1.4 as well as show that result is tight.

**Theorem 4.6** (on the complexity of sample-based testers for Bipartiteness):

1. *Bipartiteness has a (one-sided error) sample-based tester of time complexity  $O(\sqrt{n}/\epsilon)$ . In particular, this tester uses  $O(\sqrt{n}/\epsilon)$  samples.*
2. *Any sample-based tester for Bipartiteness has sample complexity  $s(n, \epsilon) = \tilde{\Omega}(\sqrt{n}/\epsilon)$ , provided  $\epsilon > 1/\sqrt{n}$ .*

**Proof:** The tester accepts if and only if the subgraph of the input graph that is revealed by the samples labeled 1 (which indicate the existence of an edge) is bipartite. That is, the tester uses the sample in order to construct a subgraph of the input graph, and accepts if and only if this subgraph is bipartite.

This algorithm always accepts graphs that are bipartite. In analyzing what happens when the input graph,  $G$ , is  $\epsilon$ -far from being bipartite, we follow the idea that underlies the proof of Theorem 4.5, but use it as a mental experiment (rather than an actual algorithm, where the actual algorithm is as stated above). Specifically, considering all possible 2-partitions of the graph

---

<sup>22</sup>Indeed, the existence of  $\epsilon n$  edges that violate a specific  $k$ -partition (i.e., reside inside the part) can be detected, with probability  $1 - \exp(-t)$ , by using  $O(t/\epsilon)$  labeled samples.

$G = (V, E)$  (which is  $\epsilon$ -far from being bipartite), we note that, with very high probability, a sample of size  $O(|V|/\epsilon)$  contains at least one violating edge for each of these 2-partitions (where a violating edge is one that connects two vertices that are on the same side of the 2-partition). Hence, the subgraph of revealed edges (considered by the algorithm) is not bipartite, and Item 1 follows.

Turning to Item 2, we first put aside the dependence on  $\epsilon$  and consider the following two distributions over  $n'$ -vertex graphs, where  $n' = \lceil \sqrt{2n} \rceil$  and each distribution is obtained by a random labeling of the corresponding unlabeled graph: The two graphs are

1. A 6-cycle of equal-sized independent sets, denoted  $V_0, \dots, V_5$ , such that each pair  $(V_i, V_{i+1 \bmod 6})$  is connected by a complete bipartite graph. In other words, this is a  $n'/6$ -factor blow-up of a 6-vertex cycle.
2. Two disjoint 3-cycles of (equal-sized) independent sets such that the three sets in each 3-cycle are connected by bipartite graphs (as in the first graph).

(A random relabeling is used so as to deem the vertex labels uninformative.) Note that all graphs in the first distribution are bipartite, while all graphs in the second distribution are  $1/18$ -far from any bipartite graph. These two distributions are indistinguishable by samples that reveal no cycles, whereas the probability of revealing a cycle of length  $i$  when each edge is selected with probability  $p$  is at most  $n' \cdot 2^{i-1} \cdot (n'/6)^{i-1} \cdot p^i$ . Hence, distinguishing the two distributions requires  $\sum_{i \geq 3} n' \cdot 2^{i-1} \cdot (n'/6)^{i-1} \cdot p^i = \Omega(1)$ , and  $n'p = \Omega(1)$  follows. Noting that  $p$  corresponds to  $s(n, 1/18)/n$ , we infer that  $s(n, 1/18)/n = p = \Omega(1/n')$ ,<sup>23</sup> which implies  $s(n, 1/18) = \Omega(\sqrt{n})$ . By applying the above construction to a  $\sqrt{\epsilon}$  fraction of the vertices, we get  $s(n, \epsilon) = \Omega(s(\lceil \sqrt{\epsilon n'} \rceil, 1/18)/\epsilon) = \Omega(\sqrt{n/\epsilon})$ , which establishes a weaker lower bound than the one in Item 2.

To establish the claimed lower bound, we generalize the construction as follows (for any  $\epsilon > 1/\sqrt{n}$ ).<sup>24</sup> We consider the following two distributions on  $n'$ -vertex graphs, where  $n' = \sqrt{2n}$ ,  $\ell \in \{\lfloor \log_2(1/\epsilon) \rfloor - 1, \lfloor \log_2(1/\epsilon) \rfloor\}$  is odd, and  $m = 2\ell^2/\epsilon$ .

1. A collection of  $2\ell m$  equal-sized independent sets, arranged in  $m$  disjoint cycle of length  $2\ell$ .
2. A collection of  $2\ell m$  equal-sized independent sets, arranged in  $2m$  disjoint cycle of length  $\ell$ .

Note that each of the graphs in the second distribution is  $\epsilon$ -far from any bipartite graph (since  $\epsilon = 1/2\ell^2 m$ ). On the other hand, these two distributions are indistinguishable by samples that reveal no cycle of length at least  $\ell$ , whereas the probability of revealing a cycle of length  $i$  when each edge is revealed with probability  $p$  is at most  $n' \cdot 2^{i-1} \cdot (n'/2\ell m)^{i-1} \cdot p^i$ . Hence,  $\sum_{i \geq \ell} n' \cdot 2^{i-1} \cdot (n'/2\ell m)^{i-1} \cdot p^i = \Omega(1)$  must hold, and  $n'p/\ell m = \Omega(1)$  follows. Noting that  $p$  corresponds to  $s(n, \epsilon)/n$ , we infer that  $s(n, \epsilon)/n = p = \Omega(\ell m/n')$ , which implies  $s(n, \epsilon) = \Omega(\sqrt{n}/(\epsilon \log(1/\epsilon)))$ . ■

**Digest.** As clarified in the proof of Item 1 of Theorem 4.6, the proof of Theorem 4.5 actually asserts that, for any graph partition property  $\Pi$ , with very high probability, the subgraph of  $G = (V, E)$  revealed by  $s = O(|V|/\epsilon^2)$  samples of its vertex pairs indicates whether  $G$  is in  $\Pi$  or is  $\epsilon$ -far from  $\Pi$ . Specifically, for every graph partition problem  $\Pi_n$ , consider the corresponding problem  $\Pi'_s$  in which edge densities are taken as a fraction of  $s$  (rather than as a fraction of  $n$ ). Then, with high probability, the distance of  $G$  from  $\Pi_n$  is approximated by the distance of the

<sup>23</sup> Again, the correspondence is not accurate, but it suffices for our purposes: See Appendix A.2.

<sup>24</sup> In the foregoing paragraph, we used  $\ell = 3$  and  $m = 1$ , which corresponds to  $\epsilon = 1/18$ .

revealed graph from  $\Pi'_s$ . Furthermore, in the case of  $k$ -coloring, we may let  $\Pi' = \Pi$  (and  $s$  can be reduced to  $O(|V|/\epsilon)$  as in the case of Bipartiteness). Thus, the computational complexity of the sample-based tester for  $\Pi_n$  is at most the computational complexity of  $\Pi'_s$ .

One natural question is whether we can obtain more efficient (w.r.t. computational complexity) sample-based testers for graph partition properties, while maintaining the sample complexity of  $s(n, \epsilon) = s'(\epsilon) \cdot \sqrt{n}$ , where  $s': (0, 1] \rightarrow \mathbb{N}$  is an arbitrary function. Recall that using a sample of size  $\tilde{O}(n)$ , which allows to fully reconstruct the graph, we can test  $\Pi_n$  in  $(\tilde{O}(n) + \exp(\text{poly}(1/\epsilon)))$ -time (by reconstructing a random induced graph of size  $\text{poly}(1/\epsilon)$ , and applying brute force on it; cf. [GGR98, Section 9]). More generally, we obtain the following trade-off.

**Theorem 4.7** (sample vs time trade-off): *For every  $t = t(n, \epsilon) \in [\text{poly}(1/\epsilon), \sqrt{n}]$ , every graph partition property (with a constant number of parts) has a sample-based tester of sample complexity  $s(n, \epsilon) = O(n/(t\epsilon^2))$  and time complexity  $\exp(t) + s(n, \epsilon)$ .*

Theorem 4.5 corresponds to  $t = \sqrt{n}$ , whereas the prior comment corresponds to  $t = \text{poly}(1/\epsilon)$ .

**Proof:** When obtaining  $s(n, \epsilon)$  random (vertex pair) samples labeled by the graph  $G = (V, E)$ , where  $n = \binom{|V|}{2}$ , the tester selects a random set of  $t$  vertices, denoted  $U$ , and (w.v.h.p) derives a residual sample of  $\Omega(s(n, \epsilon) \cdot (t/|V|)^2)$  pairs in  $U \times U$ . Hence, the size of the residual sample is  $\Omega(t/\epsilon^2)$ , and it is labeled by  $G_U$  (the subgraph of  $G$  induced by  $U$ ). Using this residual sample, the (sample-based) tester emulates the (sample-based) tester of Theorem 4.5 (for testing  $G_U$  with respect to a relaxed version of  $\Pi$ ). Hence, our tester has time complexity  $\exp(t)$ , and its analysis is based on the fact that a random induced subgraph (of that size) maintains the distance of  $G$  to  $\Pi$ . (The latter fact originates in [GGR98, Section 5.2], see also [GGR98, Corollary 7.2] and [GGR98, Corollary 8.9].) ■

**An improved result for  $k$ -Colorability.** While not resolving the general question raised above, the following result obtains a relatively efficient (w.r.t. computational complexity) sample-based testers of sublinear sample complexity for the special case of  $k$ -Colorability (as asserted in the second item of Theorem 1.4). Specifically:

**Theorem 4.8** (on the complexity of sample-based testers for  $k$ -Colorability): *For every  $k \geq 3$ , there exists a (one-sided error) sample-based tester of time complexity  $f_k(\epsilon) \cdot n^{1-(1/2k)}$  for  $k$ -Colorability, where  $f_k(\epsilon) \stackrel{\text{def}}{=} \exp(\exp(\tilde{O}(k/\epsilon)))$ . Furthermore, this tester uses  $s_k(n, \epsilon) = (k/\epsilon^2) \cdot n^{1-(1/2k)}$  samples.*

**Proof:** We first prove a weaker result and later modify it to obtain the claimed result. The weaker result asserts a sample-based tester of sample complexity  $s_k(n, \epsilon)$  that has time complexity  $f_k(\epsilon) \cdot n^{\log f_k(\epsilon)}$ , and is based upon the following tester:

On input a random labeled sample of size  $s_k(n, \epsilon)$ , the tester constructs the subgraph  $R$  of the input graph that is revealed by the samples labeled 1 (which indicate the existence of an edge), and accepts if and only if for every set  $U$  of  $k^{3k/\epsilon}$  vertices the subgraph of  $R$  induced by  $U$  is  $k$ -colorable.

That is, letting  $G = (V, E)$  denote the input graph and  $S$  denote the set of samples (i.e.  $S \in (V \times V)^{s_k(n, \epsilon)}$ ), the tester first constructs the subgraph  $R = (V, E \cap S)$ . Then, for every  $U \subset V$  of size  $k^{3k/\epsilon}$ , the tester checks whether the induced subgraph  $R_U = (U, (E \cap S) \cap (U \times U))$  is  $k$ -colorable. The tester accepts if and only if all checks are positive.

Thus, the foregoing tester always accepts  $k$ -colorable graphs, and has running time  $\binom{|V|}{k^{3k/\epsilon}} \cdot \exp(k^{3k/\epsilon})$ . The rest of the analysis is aimed at showing that *if  $G$  is  $\epsilon$ -far from  $k$ -colorable, then, with high probability over the choice of the sample, there exists a set  $U$  of  $k^{3k/\epsilon}$  vertices such that the subgraph of  $R$  induced by  $U$  is not  $k$ -colorable.*

Our analysis adapts the analysis of the (query-based)  $k$ -colorability tester of [AK02], which in turn builds upon the analysis of [GGR98]. We use [AK02] (rather than [GGR98]) as our starting point not because of its improved query complexity but rather because its “adaptive” nature lends itself well to our current needs. Specifically, the analysis of [AK02] offers a flexible analysis of an adequate  $k$ -ary tree of “color traces” (defined next). In all our definitions, we refer to a fixed graph  $G = (V, E)$  (and to a fixed value of the proximity parameter  $\epsilon$ ).

**Definition 4.8.1** (color traces): *A color trace is a sequence of  $t \geq 0$  pairs over  $V \times [k]$ ; that is,  $((v_1, c_1), \dots, (v_t, c_t))$  is a color trace if for every  $i \in [t]$  it holds that  $v_i \in V$  and  $c_i \in [k]$ .*

Note that the empty sequence is a color trace. The  $k$ -ary tree alluded to above consists of nodes holding vertices of  $G$  (i.e., the  $v_i$ 's)<sup>25</sup> and edges labeled by colors in  $[k]$  (i.e.,  $c_i$ 's) such that color traces are read by taking a path from the root of this tree and reading both the labels of the nodes and of the edges. In other words, we shall maintain a set of color traces that is closed under taking prefixes; that is, if  $((v_1, c_1), \dots, (v_t, c_t))$  is in the set, then  $((v_1, c_1), \dots, (v_{t-1}, c_{t-1}))$  is also in the set. Furthermore, the  $k$ -ary requirement corresponds to saying that if  $((v_1, c_1), \dots, (v_t, c_t))$  is in the set then for every  $c \in [k]$  the color trace  $((v_1, c_1), \dots, (v_{t-1}, c_{t-1}), (v_t, c))$  is also in the set. Color traces are also viewed as partial assignments (of colors to the vertices), where the color trace  $((v_1, c_1), \dots, (v_t, c_t))$  corresponds to assigning vertex  $v_i$  the color  $c_i$ , for every  $i \in [t]$ . We shall consider the constraints imposed by such color traces on the (rest of the) vertices in  $G$ .

**Definition 4.8.2** (forbidden colors): *The set of forbidden colors of vertex  $v$  w.r.t. the color trace  $\overline{\tau} = ((v_1, c_1), \dots, (v_t, c_t))$ , denoted  $F_{\overline{\tau}}(v)$ , is the set of colors assigned by  $\overline{\tau}$  to vertices that reside on  $\overline{\tau}$  and neighbor  $v$ ; that is,  $F_{\overline{\tau}}(v) \stackrel{\text{def}}{=} \{c_i : i \in [t] \wedge \{v, v_i\} \in E\}$ .*

That is, the forbidden set of  $v$  is the set of colors that are not allowed for  $v$  in any  $k$ -coloring that extends the coloring represented in the trace. Indeed, the color trace imposes natural restrictions on itself, and violating these restrictions yield an illegal color trace (as defined next).

**Definition 4.8.3** (illegal traces): *A color trace  $\overline{\tau} = ((v_1, c_1), \dots, (v_t, c_t))$  is illegal if there exists  $i \in [t]$  such that  $c_i$  is a forbidden color of vertex  $v_i$  w.r.t. the color trace  $\overline{\tau}' = ((v_1, c_1), \dots, (v_{i-1}, c_{i-1}))$ ; that is, if  $c_i \in F_{\overline{\tau}'}(v_i)$ . Otherwise, the trace is called legal.*

Indeed, illegal traces must have length at least two. Turning back to the forbidden sets of all vertices in  $G$ , we consider the effect on these sets that is due to augmenting a specific color trace with a specific vertex-color pair.

**Definition 4.8.4** (restricting pairs and vertices): *We say that  $(v, c) \in V \times [k]$  is restricting w.r.t.  $\overline{\tau}$  if  $|R_{\overline{\tau}}(v, c)| \geq \epsilon|V|/3$ , where*

$$R_{\overline{\tau}}(v, c) \stackrel{\text{def}}{=} \{w \in V : F_{\overline{\tau}, (v, c)}(w) \supset F_{\overline{\tau}}(w)\}. \quad (6)$$

*We say that  $v \in V$  is restricting w.r.t.  $\overline{\tau}$  if for every  $c \in [k]$  either  $c \in F_{\overline{\tau}}(v)$  or  $(v, c)$  is restricting w.r.t.  $\overline{\tau}$ .*

---

<sup>25</sup>With the exception of leaving the leaves of the tree empty.

Indeed, for every  $\bar{\tau}, (v, c)$  and  $w$ , it holds that  $F_{\bar{\tau},(v,c)}(w)$  equals either  $F_{\bar{\tau}}(w)$  or  $F_{\bar{\tau}}(w) \cup \{c\}$ , and  $F_{\bar{\tau},(v,c)}(w) \supset F_{\bar{\tau}}(w)$  holds if and only if  $c \notin F_{\bar{\tau}}(w)$  and  $\{v, w\} \in E$ . Also, every  $v$  such that  $F_{\bar{\tau}}(v) = [k]$  is restricting w.r.t.  $\bar{\tau}$ . It turns out that if  $G$  is  $\epsilon$ -far from being  $k$ -colorable, then every legal color trace has many restricting vertices. (The claim holds also for illegal traces, provided that they are not too long, but the following claim suffices for our needs.)

**Claim 4.8.5** *If there exists a legal color trace  $\bar{\tau} = ((v_1, c_1), \dots, (v_t, c_t))$  such that less than  $\epsilon|V|/3$  vertices are restricting w.r.t.  $\bar{\tau}$ , then  $G$  is  $\epsilon$ -close to being  $k$ -colorable.*

**Proof:** Consider the following  $k$ -partition of  $V$ , denoted  $\chi: V \rightarrow [k]$ . For each  $i \in [t]$ , we set  $\chi(v_i) = c_i$ . For any *non-restricting* vertex  $v$ , we set  $\chi(v) = c$  such that  $(v, c)$  is *non-restricting* w.r.t.  $\bar{\tau}$ . (Such a color  $c$  exists, since otherwise  $v$  must be restricting. For all other vertices (i.e., the restricting ones), we set the color  $\chi$  arbitrarily.) We upper bound the number of violating edges (i.e., edges with both endpoints assigned the same color) by charging edges that are incident at a non-restricting vertex  $v$  to the pair  $(v, \chi(v))$ , and charging edges that connect two restricting vertices to the corresponding vertex pair. Denoting the set of non-restricting (w.r.t.  $\bar{\tau}$ ) vertices by  $V'$ , the total charge is upper-bounded by

$$\sum_{v \in V'} |R_{\bar{\tau}}(v, \chi(v))| + \binom{|V \setminus V'|}{2} < |V| \cdot \frac{\epsilon|V|}{3} + \binom{\epsilon|V|/3}{2}$$

and the claim follows. ■

**Definition 4.8.6** (open traces): *A legal color trace  $\bar{\tau}$  is called open if there exist at least  $\epsilon|V|/3$  vertices that are restricting w.r.t.  $\bar{\tau}$ .*

Claim 4.8.5 implies that if  $G$  is  $\epsilon$ -far from being  $k$ -colorable, then each color trace must be either illegal or open. The following “restriction procedure” is aimed at producing a set of traces with a frontier (i.e., maximal traces) that are all illegal traces, where a trace is in the frontier of the set (i.e., is maximal) if it is not a proper prefix of any other trace in the set. The restriction procedure may only augment an existing trace  $\bar{\tau} = ((v_1, c_1), \dots, (v_t, c_t))$  by a vertex-color pair  $(v, c)$  if  $F_{\bar{\tau}}(v)$  ( $= \{c_i : i \in [t] \wedge \{v, v_i\} \in E\}$ ) equals  $\{c_i : i \in [t] \wedge \{v, v_i\} \in E \cap S\}$ , where  $S$  is the sequence of samples provided to the tester.<sup>26</sup> This means that the sample  $S$  contains “evidence” for each color contained in the forbidden set  $F_{\bar{\tau}}(v)$  of the vertex  $v$ , where the said evidence for a color  $c' \in F_{\bar{\tau}}(v)$  is a pair  $\{v, v_i\} \in E \cap S$  such that  $c_i = c'$  (i.e., the pair  $\{v, v_i\}$  is in the sample and is an edge of the graph and so  $v$  cannot be legally colored by  $c_i$ ). We stress that the following restriction procedure is merely a mental experiment conducted in the analysis.

**Restriction Procedure (as a mental experiment).** The procedure is given a sequence of labeled samples, where each sample point is a pair of vertices and the label indicates whether or not the edge is in the graph; that is, the sample point  $(u, v)$  is labeled by  $g(u, v)$  such that  $g(u, v) = 1$  if  $\{u, v\} \in E$  and  $g(u, v) = 0$  otherwise. (If  $(u, v)$  is not in the sample set  $S$ , then  $g$  is undefined on  $(u, v)$ .) The procedure is initiated with an empty trace, which is typically an open trace (since otherwise

---

<sup>26</sup>This is the point where we depart from [AK02]: The algorithm in [AK02] uses a single sample of  $\tilde{O}(1/\epsilon^2)$  vertices and queries all vertex pairs in this sample. Only these vertices will be placed in color traces, and the algorithm has evidence for each color contained in the forbidden set of a vertex  $v$  used to augment  $\bar{\tau}$ . In contrast, our algorithm obtains a random sample  $S$  of vertex pairs, and the restriction procedure presented below is confined to evidence that is revealed by  $S$ . In particular, the restriction procedure of [AK02] (re)uses the same vertices on all traces, whereas our restriction procedure uses different vertices when augmenting different traces.

$|E| < \epsilon \cdot \binom{|V|}{2}$ ). As long as the set of traces contains a *maximal* trace  $\bar{\tau} = ((v_1, c_1), \dots, (v_t, c_t))$  that is both *legal* and *open*<sup>27</sup>, the procedure takes the following steps.

1. If there exists no vertex  $v$  that satisfies the following two conditions, then the procedure halts with an error message. The conditions (regarding  $v$  and  $\bar{\tau}$ ) are:

- (a)  $v$  is restricting w.r.t.  $\bar{\tau}$ ; and
- (b) for every  $c \in F_{\bar{\tau}}(v)$  there exists  $i \in [t]$  such that  $c_i = c$  and  $g(v, v_i) = 1$ .  
(Recall that  $g(v, v_i) = 1$  means that  $(v, v_i)$  is in the sample set  $S$  and  $\{v, v_i\} \in E$ .)

Otherwise (i.e., some  $v$  satisfies the above condition), the procedure picks such a vertex  $v$ .

2. For every  $c \in [k]$ , the procedure augments the set of traces with  $((v_1, c_1), \dots, (v_t, c_t), (v, c))$ .  
(Indeed, this trace is illegal if  $c \in F_{\bar{\tau}}(v)$ , and this illegality is visible in the sample.)

If the current set of traces contains no maximal trace that is both legal and open, then the procedure halts and outputs this set of traces.

**Claim 4.8.7** *The set of color traces that the procedure constructs contains traces of length at most  $3k/\epsilon$ .*

**Proof:** The procedure only augments traces by using restricting vertices, but any restricting vertex w.r.t. a given trace increases the sum of the sizes of the forbidden sets by at least  $\epsilon|V|/3$ . Specifically, suppose that the trace  $\bar{\tau} = ((v_1, c_1), \dots, (v_t, c_t))$  is contained in the output, and let  $\bar{\tau}^{(i)}$  denote its  $i$ -long prefix (i.e., the trace  $((v_1, c_1), \dots, (v_i, c_i))$ ). Then, for every  $i \in [t]$ , it holds that

$$\begin{aligned} \sum_{v \in V} |F_{\bar{\tau}^{(i)}}(v)| &= |R_{\bar{\tau}^{(i-1)}}(v_i, c_i)| + \sum_{v \in V} |F_{\bar{\tau}^{(i-1)}}(v)| \\ &\geq \frac{\epsilon|V|}{3} + \sum_{v \in V} |F_{\bar{\tau}^{(i-1)}}(v)| \end{aligned}$$

and it follows that  $t \leq k|V|/(\epsilon|V|/3)$ . ■

**Claim 4.8.8** *Suppose that  $G$  is  $\epsilon$ -far from being  $k$ -colorable. If the restriction procedure halts without error, then the sample  $S$  contains a subgraph that is not  $k$ -colorable. Furthermore, this subgraph contains only vertices that appear in the set of color traces output by the procedure.*

**Proof:** By Claim 4.8.5 each color trace must be either illegal or open (since  $G$  is  $\epsilon$ -far from being  $k$ -colorable). Hence, the (non-error) output of the procedure is a set of color traces such that all maximal traces in this set are illegal. Furthermore, for each vertex that occurs on any trace, the sample contains evidence for the forbidden colors w.r.t. the relevant prefix of the trace. Thus, an illegal trace asserts that the corresponding coloring of its vertices is not a legal  $k$ -coloring in the subgraph  $R$  revealed by the sample. Now, since all maximal traces are illegal and the set of traces corresponds to a  $k$ -ary tree (i.e., for each non-maximal trace  $\bar{\tau}$ , there exists a  $v$  such that, for each  $c \in [k]$ , the set of traces contains the trace  $(\bar{\tau}, (v, c))$ ), it follows that there is no legal  $k$ -coloring of  $R$  (or rather of the subgraph of  $R$  induced by the vertices that appear on this set of traces). ■

---

<sup>27</sup>Recall that this condition (as well as the other conditions below) is not tested by the actual algorithm (i.e., the one outlined in the beginning of this proof). This condition is only considered by the mental experiment that we conduct in our analysis. Furthermore, we may assume that all legal traces are open, since Claim 4.8.5 asserts that this is the case if  $G$  is  $\epsilon$ -far from being  $k$ -colorable.

**Claim 4.8.9** *The probability that the restriction procedure halts with error is at most  $1/3$ , where the probability is taken uniformly over the choice of the sample  $S$ .*

**Proof:** It will be more convenient to consider a sample  $S$  that is chosen by selecting each vertex pair with probability  $p = s_k(n, \epsilon)/2n$ , independent of all other choices. (With very high probability, such a sample of independently chosen pairs contains at most  $2pn = s_k(n, \epsilon)$  pairs, whereas the probability that the restriction procedure halts with error decreases when the sample size increases.)<sup>28</sup> The restriction procedure makes at most  $k^{3k/\epsilon}$  iterations, where in each iteration it tries to extend a legal open trace  $\bar{\tau} = ((v_1, c_1), \dots, (v_t, c_t))$ . At this point, the set of restricting vertices, denoted  $V'$ , has size at least  $\epsilon|V|/3$ . For each vertex  $v \in V'$ , let  $I_v \subseteq \{i \in [t] : \{v, v_i\} \in E\}$  be a set of  $|F_{\bar{\tau}}(v)|$  indices such that for every  $c \in F_{\bar{\tau}}(v)$  there exists  $i \in I_v$  such that  $c_i = c$  (and  $\{v, v_i\} \in E$ ). Now, the probability that the sample contains  $\{(v, v_i) : i \in I_v\}$  equals  $p^{|I_v|}$ , and these events are independent for the various  $v$ 's in  $V'$  (since each event refers to the existence of samples that are adjacent at a different vertex  $v \in V'$ ).<sup>29</sup> Thus, the current iteration fails with probability at most

$$\prod_{v \in V'} (1 - p^{|F_{\bar{\tau}}(v)|}) < (1 - p^k)^{\epsilon|V|/3}$$

which is  $\exp(-p^k \cdot \epsilon|V|)$ . Using  $p = \Omega(|V|^{-1/k}/\epsilon)$ , we have  $\exp(-p^k \cdot \epsilon|V|) \ll k^{-3k/\epsilon}$  and the claim follows. ■

Combining Claims 4.8.9 and 4.8.8, it follows that *if  $G$  is  $\epsilon$ -far from  $k$ -colorable, then, with probability at least  $2/3$  over the choice of the sample  $S$ , the vertices that appear in the output of the restriction procedure induce a subgraph of  $R$  that is not  $k$ -colorable, where  $R$  is the subgraph of  $G$  that is visible in the sample.* Using Claim 4.8.7, the number of vertices in this output (and hence in the induced subgraph) is at most  $f'_k(\epsilon) = k^{3k/\epsilon}$ . This establishes the claim that  $k$ -colorability has a sample-based tester of sample complexity  $s_k(n, \epsilon)$  and time complexity  $n^{f'_k(\epsilon)} \cdot \exp(f'_k(\epsilon))$ .

In order to improve the time complexity of the tester, we turn the restriction procedure from a mental experiment to an actual procedure to be implemented by the tester. An immediate difficulty that arises is that the tester may not be able to determine the forbidden sets of various vertices and the (approximate) number of restricting vertices (w.r.t. various color traces). Our solution is just to guess these unknowns at random (and to compensate for the lower probability of success by a suitable number of repetitions (to be discussed below)).

**Restriction Procedure (as an actual implementation).** The procedure is initiated with an empty trace, and assumes that each trace is legal unless some evidence is found against this assumption (see below). (For every legal trace, the procedure also assumes that the trace is open, and so we shall not refer at all to the question of whether a trace is open.) (Note that the procedure does not know the forbidding set of a vertex, nor its neighbors on the trace, and so it cannot tell whether or not a trace is actually legal, let alone whether or not a trace is open.) In particular, as long as the set of traces contains a *maximal* trace  $\bar{\tau} = ((v_1, c_1), \dots, (v_t, c_t))$  that is *assumed* to be legal, the procedure takes the following steps.

1. If  $t \geq 3k/\epsilon$ , then the procedure halts with an error message.

Otherwise (i.e.,  $t < 3k/\epsilon$ ), if  $t \leq k$ , then the procedure sets  $I = [t]$ , otherwise the procedure randomly selects  $I$  uniformly among all  $k$ -subsets of  $[t]$ . The procedure selects at random a vertex  $v$  among all vertices  $v$  such that for every  $i \in I$  the pair  $(v, v_i)$  is in the sample.

---

<sup>28</sup>See further justification in Appendix A.2.

<sup>29</sup>Also note that these events are unconditioned by any other events that were considered when constructing  $\bar{\tau}$ , since these events refer to samples that connect pairs of  $v_i$ 's.

(If there exists no vertex  $v$  that satisfies the above condition, then the procedure halts with an error message. This is quite unlikely (since  $(s_k(n, \epsilon)/n)^k \cdot |V| \gg 1$ ), and otherwise we continue to the next step.)

(Motivation: Our hope is that the vertex  $v$  selected above is restricting w.r.t.  $\bar{\tau}$ , and that  $\{c_i : i \in I \wedge \{v, v_i\} \in E\} = F_{\bar{\tau}}(v)$ .)

2. Let  $v$  be the vertex selected in the previous step, and let  $C_v \subseteq F_{\bar{\tau}}(v)$  be the set of colors that cannot be assigned to  $v$  due to an edge  $\{v, v_i\}$  that is visible in the sample; that is,  $C_v = \{c_i : i \in I \wedge \{v, v_i\} \in E\}$ , where  $I$  is the set selected in the previous step. For every  $c \in [k]$ , the procedure augments the set of traces with the trace  $((v_1, c_1), \dots, (v_t, c_t), (v, c))$ . For each such new trace, the procedure assumes that it is legal unless  $c \in C_v$ ; that is,  $((v_1, c_1), \dots, (v_t, c_t), (v, c))$  is assumed to be legal if and only if  $c \notin C_v$ .

If the current set of traces contains no maximal trace that is assumed to be legal, then the procedure halts outputting this set of traces.

Note that we augmented the original Step 1 such that the procedure never outputs a set of traces that contains a trace that is longer than  $3k/\epsilon$ . Thus, the restriction procedure can be implemented in time  $O(k^{3k/\epsilon} \cdot s_k(n, \epsilon))$ .

Turning to the analysis of the output of the procedure, note that whenever a set of traces is output, the maximal traces in the output are all illegal (and their illegality is visible in the sample). Thus, we obtain a  $k^{3k/\epsilon}$ -vertex subgraph of  $G$  that is not  $k$ -colorable such all edges of this subgraph appear in the sample. This will happen if (and only if) all random guesses made by the restriction procedure are successful (i.e., satisfy our hopes as stated in the foregoing motivation). Indeed, the probability of this random event is quite low, but our final tester will consist of invoking the restriction procedure for an adequate number of times. We refer to such a single invocation as a *basic tester*. We stress that the basic tester accepts whenever the procedure halts with an error message, and rejects only if the procedure outputs a set of traces (which by the foregoing discussion yields a subgraph of  $R$  that is induced by the vertices that reside on these traces such that this subgraph is not  $k$ -colorable).

Assuming that  $G$  is  $\epsilon$ -far from being  $k$ -colorable, we lower bound the probability that the basic tester rejects. This boils down to lower-bounding the probability that the actual restriction procedure correctly emulates the mental experiment. The key to the analysis is provided by the following claim.

**Claim 4.8.10** *Suppose that  $\bar{\tau} = ((v_1, c_1), \dots, (v_t, c_t))$  is a legal and open trace produced by the restriction procedure, and that the procedure executes Step 1 on  $\bar{\tau}$ . Then, with probability at least  $(\epsilon/3k)^{2k+1}$ , the procedure selects a vertex that is restricting w.r.t.  $\bar{\tau}$ , where the probability is taken uniformly over the choice of the sample  $S$  and the choices made in Step 1.*

*Proof:* Again, it will be more convenient to consider a sample  $S$  that is chosen by selecting each vertex pair with probability  $p = s_k(n, \epsilon)/2n$ , independent of all other choices.<sup>30</sup> For every  $I \subseteq [t]$ , let  $S_I$  be the set of vertices  $v$  such that for every  $i \in I$  the pair  $(v, v_i)$  is in the sample  $S$ . Note that each  $S_I$  is a random set of vertices that is obtained by selecting each vertex with probability exactly  $p^{|I|}$ . Turning to the vertices that are restricting w.r.t.  $\bar{\tau}$ , we denote by  $V_J$  the set of restricting vertices  $v$  such that  $\{c_j : j \in J\} = F_{\bar{\tau}}(v)$ . Then, there exists a set  $J \subseteq [t]$  of size at most  $k$  such that  $\frac{|V_J|}{|V|} > \frac{\epsilon/3}{(3k/\epsilon)^k} > 2 \cdot (\epsilon/3k)^{k+1}$ , since the union of all these  $V_J$ 's contains all restricting

---

<sup>30</sup>Again, the correspondence is not accurate, but it suffices for our purposes: See Appendix A.2.

vertices. Letting  $I \subseteq [t]$  be an arbitrary  $\min(k, t)$ -superset of  $J$ , and suppose that  $I$  is selected at Step 1 (which happens with probability at least  $(\epsilon/3k)^k$ ). Then, with very high probability over the choice of  $S$ , the set  $S_I \cap V_J$  constitutes at least a  $(\epsilon/3k)^{k+1}$  fraction of  $S_I$ , where we use  $(\epsilon/3k)^{k+1} \cdot p^k |V| \gg 1$  (cf. Claim 4.8.9). Hence, for  $J$  fixed as above, a set  $I \supseteq J$  is selected (in Step 1) with probability at least  $(\epsilon/3k)^{k+1}$ , the random  $v \in R_I$  (also selected in Step 1) is in  $V_J$  with probability at least  $(\epsilon/3k)^k$ , and the claim follows. ■

Indeed, as suggested in the foregoing motivation, it follows that the actual restriction procedure emulates the imaginary one with probability at least  $((\epsilon/3k)^{2k+1})^{k^{3k/\epsilon}} = 1/f_k(\epsilon)$ , which establishes the time bound claimed in the theorem (since it suffices to repeat the basic test for  $f_k(\epsilon)$  times). But this argument means that the sample complexity is also multiplied by a factor of  $f_k(\epsilon)$ .

To obtain a better upper bound on the sample complexity, we modify the restriction procedure such that the choice of  $I$  is made by branching in parallel (to  $(3k/\epsilon)^k$  branches) rather than by a random selection (among  $(3k/\epsilon)^k$  options). Note that the proof of Claim 4.8.10 actually shows that the probability that a restricting vertex is selected (on the adequate branch) is at least  $(\epsilon/3k)^{k+1}$ .

Next, rather than selecting at random one vertex (in Step 1 of the restriction procedure), we select at random  $(3k/\epsilon)^{2k}$  vertices (among all vertices  $v$  such that for every  $i \in I$  the pair  $(v, v_i)$  is in the sample), and branch on these  $(3k/\epsilon)^{2k}$  chosen vertices too, where all is done using the same sample. We can use the same sample on all branches, because we actually care only about what happens on one good branch. Hence, we get back to the sample complexity of the simpler version (i.e., the one that uses an imaginary procedure). For sake of clarity, we spell out the revised restriction procedure.

**Restriction Procedure (revised parallel implementation).** For  $K \stackrel{\text{def}}{=} (3k/\epsilon)^k$  and  $M \stackrel{\text{def}}{=} k^{3k/\epsilon}$ , we run in parallel  $(K \cdot K^2)^M$  parallel copies of the restriction procedure described above, all using the same sample. The copies correspond to choices  $((I_1, i_1), \dots, (I_M, i_M))$ , where each  $I_j$  is a  $k$ -subset of  $[3k/\epsilon]$  and  $i_j \in [K^2]$ . A generic copy, which corresponds to such a sequence of pairs, uses the sets  $I_1, \dots, I_M$  instead of the random choices of  $k$ -subsets made in Step 1 such that when augmenting the  $j^{\text{th}}$  trace it uses the  $k$ -subset  $I_j$ . These copies also share their randomness, denoted  $(r_1, \dots, r_M)$ , which is used to select the vertices  $v$  in Step 1 such that  $r_j$  is used when augmenting the  $j^{\text{th}}$  trace. The generic copy that corresponds to  $((I_1, i_1), \dots, (I_M, i_M))$  uses the  $i_j^{\text{th}}$  block of  $r_j$  when augmenting the  $j^{\text{th}}$  trace (i.e.,  $r_j = (r_{j,1}, \dots, r_{j,K^2})$  and the  $i_j^{\text{th}}$  block of  $r_j$  means  $r_{j,i_j}$ ). Thus, the modified Step 1 of this generic copy reads:

Suppose that this is the  $j^{\text{th}}$  time that Step 1 is executed. Then (assuming  $t < 3k/\epsilon$  and  $t > k$ )<sup>31</sup>, sets  $I = I_j$  and uses the randomness  $r_{j,i_j}$  in order to select  $v$  (among all vertices  $v$  such that for every  $i \in I$  the pair  $(v, v_i)$  is in the sample).

Each of the copies produces an output as the original restriction procedure, and the parallel restriction procedure outputs a set of traces (rather than an error message) if any of the copies has output such a set. ■

**Beyond  $k$ -Colorability.** The case of  $k$ -Colorability represents almost all that can be tested via one-sided error testers of query complexity that only depends on  $\epsilon$ . Recall that by [GT03, Theorem 3], the only such (non-trivial)<sup>32</sup> properties are (1) the property of being a clique, and

<sup>31</sup>The other cases are handled exactly as before; that is, the copy halts with error if  $t \geq 3k/\epsilon$  and sets  $I = [t]$  if  $t \leq k$ .

<sup>32</sup>A property is called non-trivial (for testing) if, for every  $\epsilon > 0$  and all sufficiently large  $n$ , there exists both an  $n$ -sized object in the property and an  $n$ -sized object that is  $\epsilon$ -far from the property.

(2) properties that imply  $k$ -Colorability. The first property (which consists of a single graph for every size) can be tested by  $O(1/\epsilon)$  random samples, and the result of Theorem 4.8 extends to the second class (as detailed next). Properties of the second class correspond to  $k$ -vertex graphs such that the property associated with the graph  $H = ([k], A)$  consists of all graphs  $G = (V, E)$  that can be  $k$ -partitioned such that edges appear only between vertices that reside in the  $i^{\text{th}}$  and  $j^{\text{th}}$  part when  $(i, j) \in A$ . Such graphs  $G$  are called  $H$ -embeddable.

**Definition 4.9** ( $H$ -embeddable graphs): *Let  $H = ([k], A)$  be a simple graph (with no self-loops). The graph  $G = (V, E)$  is  $H$ -embeddable if there exists a  $k$ -partition of  $V$  into  $(V_1, \dots, V_k)$  such that  $E \cap (V_i \times V_j) \neq \emptyset$  only if  $(i, j) \in A$ .*

Note that  $k$ -Colorability corresponds to the set of graphs that are embedded in the  $k$ -vertex clique.

**Theorem 4.10** (on the complexity of sample-based testers for  $H$ -Embeddability): *Let  $H = ([k], A)$  be a simple graph (with no self-loops). Then there exists a (one-sided error) sample-based tester of time complexity  $f_k(\epsilon) \cdot n^{1-(1/2k)}$  for  $H$ -Embeddability, where  $f_k(\epsilon) \stackrel{\text{def}}{=} \exp(\exp(\tilde{O}(k/\epsilon)))$ . Furthermore, this tester uses  $s_k(n, \epsilon) = (k/\epsilon^2) \cdot n^{1-(1/2k)}$  samples.*

**Proof Sketch:** Building upon the proof of Theorem 4.8, we let  $F'_\tau(v) \stackrel{\text{def}}{=} \bigcup_{c \in F_\tau(v)} \{c' \in [k] : h(c, c') = 0\}$ , where  $h(c, c') = 0$  if and only if  $\{c, c'\} \notin A$ . That is, while  $F_\tau(v)$  denotes the set of colors of the vertices on the color trace  $\tau$ , here the set  $F'_\tau(v) \supseteq F_\tau(v)$  is the set of colors that are forbidden for  $v$  (in a legal extension of the coloring implicit in  $\tau$ ). All definitions and claims (with one exception discuss below) will refer to the modified sets of forbidden colors  $F'_\tau(v)$ . The only exception is the “evidence condition” in Item 1b of the Restriction Procedure, where we still require that *for every  $c \in F_\tau(v)$  (rather than  $c \in F'_\tau(v)$ ) there exists  $i \in [t]$  such that  $c_i = c$  and  $g(v, v_i) = 1$*  (where  $g(v, v_i) = 1$  means that  $(v, v_i)$  is in the sample and  $\{v, v_i\} \in E$ ). ■

## 5 Sample-based testers for linearity

Recall that “testing linearity” actually refers to testing homomorphism between two groups. For two groups  $(G, +)$  and  $(H, \odot)$ , the function  $h: G \rightarrow H$  is a **group homomorphism** if for every  $x, y \in G$  it holds that  $h(x + y) = h(x) \odot h(y)$ . Recall that group homomorphism can be tested by  $O(1/\epsilon)$  queries regardless of the size of the groups [BLR93], and that by Theorem 3.5 there exists a (one-sided error) sample-based tester of sample complexity  $s(n, \epsilon) = \text{poly}(1/\epsilon) \cdot n^{1/3}$ . As noted in Section 3.3, a more efficient sample-based tester can be obtained directly.

**Theorem 5.1** (the sample-complexity of testing group homomorphism):

1. *For any two groups  $(G, +)$  and  $(H, \odot)$ , group homomorphism from  $G$  to  $H$  has a one-sided error sample-based tester of sample complexity  $O(1/\epsilon + \log |G|)$ .*
2. *There exist groups  $(G, +)$  and  $(H, \odot)$ , such that sample-based testing group homomorphism from  $G$  to  $H$  requires  $\Omega(1/\epsilon + \log |G|)$  samples, provided  $\epsilon \geq 1/|G|$ .*

**Proof:** The proof of the upper bound (i.e., Item 1) is based on the fact that, with very high probability, the partial sums of  $O(\log |G|)$  random elements of  $G$  generates all elements of  $G$ . In fact, a stronger statement holds:

**Claim 5.1.1** *Suppose that  $r_1, \dots, r_t$  are uniformly and independently distributed in  $G$ . Then, with probability at least  $1 - (\delta^{-2}|G|/(2^t - 1))$ , for every  $v \in G$ , it holds that  $|\{I \subseteq [t] : I \neq \emptyset \wedge \sum_{i \in I} r_i = v\}| = (1 \pm \delta) \cdot (2^t - 1)/|G|$ , where  $\sum_{i \in I} r_i$  is a shorthand for  $r_{i_1} + \dots + r_{i_{|I|}}$  such that  $i_1 < i_2 < \dots < i_{|I|}$  is an ordering of the elements of  $I$ .*

(We are careful about the meaning of  $\sum_{i \in I}$ , since  $G$  may not be Abelian.)

**Proof:** For every non-empty  $I \subseteq [t]$ , we consider the random variable  $R_I = R_I(r_1, \dots, r_t)$  defined as  $\sum_{i \in I} r_i$ . Each of these random variables is uniformly distributed in  $G$ . Furthermore, these random variables are pairwise independent, since for every  $I \setminus J \neq \emptyset$  and  $u, v \in G$  we have  $\Pr[\sum_{i \in I} r_i = u \mid \sum_{j \in J} r_j = v] = 1/|G|$  (by fixing the values of all  $r_i$  except for some  $i \in I \setminus J$ ). Now, fix an arbitrary element  $v \in G$ , and let  $\zeta_I = 1$  if  $R_I = v$  and  $\zeta_I = 0$  otherwise. Then,  $\mathbb{E}[\zeta_I] = 1/|G|$  and the  $\zeta_I$ 's are pairwise independent. Letting  $\bar{\zeta}_I = \zeta_I - \mathbb{E}[\zeta_I]$ , and applying Chebyshev's Inequality we get

$$\begin{aligned} \Pr \left[ \left| \sum_{\emptyset \neq I \subseteq [t]} \bar{\zeta}_I \right| > \delta \cdot (2^t - 1)/|G| \right] &< \frac{\mathbb{E} \left[ \left( \sum_I \bar{\zeta}_I \right)^2 \right]}{(\delta \cdot (2^t - 1)/|G|)^2} \\ &< \frac{\mathbb{E} \left[ \sum_I \zeta_I^2 \right]}{(\delta \cdot (2^t - 1)/|G|)^2} \\ &= \frac{(2^t - 1)/|G|}{\delta^2 \cdot (2^t - 1)^2 / |G|^2} \end{aligned}$$

and the claim follows. ■

By Claim 5.1.1, for  $t = O(\log |G|)$  and for any homomorphism  $h: G \rightarrow H$ , the values of  $h$  at  $t$  random samples, determine the value of  $h$  on all of  $G$ . This suggests the following sample-based tester. On input an  $f$ -labeled sample  $((r_1, f(r_1)), \dots, (r_t, f(r_t)), (r_{t+1}, f(r_{t+1})), \dots, (r_{t+t'}, f(r_{t+t'})))$ , where  $t' = O(1/\epsilon)$ , the tester proceeds as follows:

1. Reconstruct a function  $\tilde{f}$  based on the labeled sample and check that  $\tilde{f}$  is a homomorphism. Specifically, for every  $x \in G$  find a (non-empty) set  $I \subseteq [t]$  such that  $\sum_{i \in I} r_i = x$ , and set  $\tilde{f}(x) = \odot_{i \in I} f(r_i)$ . If  $\tilde{f}$  is not defined on all points in  $G$ , then accept. If  $\tilde{f}$  is not a group homomorphism, then reject.

We continue to the next step only if  $\tilde{f}$  is defined over all of  $G$  and is a group homomorphism. (Note that we do not check whether the value defined based on  $I$  matches the value that could have been defined based on a different  $I'$  such that  $\sum_{i \in I} r_i = x$ ; such a check is natural but not needed.)

2. Test that  $f = \tilde{f}$ . That is, accept if and only if for every  $i \in [t']$  it holds that  $f(r_{t+i}) = \tilde{f}(r_{t+i})$ .

If  $f$  is a group homomorphism and  $\tilde{f}$  was defined on  $x \in G$ , then necessarily  $\tilde{f}(x) = \odot_{i \in I} f(r_i) = f(x)$  for all  $I$ 's such that  $\sum_{i \in I} r_i = x$ . Hence, the test always accepts group isomorphisms. Assume, on the other hand, that  $f$  is  $\epsilon$ -far from being a group isomorphism. Then, with probability at least  $5/6$ , for every  $x \in G$ , there exists an  $I \subseteq [t]$  such that  $\sum_{i \in I} r_i = x$ . In this case, unless the tester rejects already in Step 1, it holds that  $\tilde{f}$  is defined over all of  $G$  and is a group homomorphism. Hence,  $f$  and  $\tilde{f}$  differ on at least an  $\epsilon$  fraction of  $G$ , and it follows that Step 2 rejects with probability at least  $5/6$ . Having established the upper bound, we now turn to the lower bound.

**Claim 5.1.2** *Any sample-based tester for group homomorphism from  $\mathbb{Z}_2^m$  to  $\mathbb{Z}_2$  requires  $\Omega(1/\epsilon + m)$  samples.*

In this case, the group homomorphisms correspond to linear functions.

**Proof:** The key observation is that, with very high probability, a sample of  $m/2$  random points in  $\mathbb{Z}_2^m$  yields  $m/2$  linearly independent vectors. In this case, the labels provided to these points by a random linear function are identically distributed to the labels provided by a totally random function (which, with very high probability, is  $1/5$ -far from any linear function). Hence, a lower bound of  $m/2$  follows. To obtain an  $\Omega(1/\epsilon)$  lower bound consider a linear function that is modified on an  $2\epsilon$  fraction of its domain, and note that  $\Omega(1/\epsilon)$  samples are required for hitting this modified region. ■

This completes the proof of the theorem. ■

**Digest.** The sample-based tester presented in the proof of Theorem 5.1 is based on an exact learning algorithm: Indeed, when given  $O(\log |G|)$  random samples labeled by a group homomorphism  $f$ , with very high probability, Step 1 reconstructs  $f$ . Furthermore, when Step 1 fails, it indicates this event (by halting with a partially defined  $f$ ). Thus, an alternative presentation of the (one-sided error) tester may proceed along the reduction of testing to (proper) learning presented in [GGR98, Proposition 3.1], while relying on the features of the aforementioned learning algorithm (i.e., its zero-error features). Indeed, it follows that *testing linearity is not easier than (exact) learning of linear functions*.

## 6 Invariance under the symmetric group

In this section we consider the case of properties that are invariant under the symmetric group acting on the functions' domain. That is, we consider properties  $\Pi = \bigcup_{n \in \mathbb{N}} \Pi_n$  such that for every  $n \in \mathbb{N}$  and every permutation  $\pi: D_n \rightarrow D_n$  it holds that  $f \in \Pi_n$  if and only if  $f \circ \pi \in \Pi_n$ , where  $(f \circ \pi)(x) = f(\pi(x))$ . We call such properties *symmetric*.

### 6.1 Queries versus samples

We first show that in the context of testing symmetric properties, queries add little power beyond samples. This is quite intuitive when considering sampling without repetitions, but our sampling model is with repetitions and the gap between these two models is significant for  $\Omega(\sqrt{n})$  samples. The proof of the following result starts by validating the former intuition and closing the latter gap.

**Theorem 6.1** (samples are as good as queries for testing symmetric properties): *Suppose that  $\Pi$  is a symmetric property that has a standard tester of query complexity  $q$ . Then,  $\Pi$  has a sample-based tester of query complexity  $s'$  such that  $s'(n, \epsilon) = O(q(n, \epsilon))$  if  $q(n, \epsilon) < n/2$  and  $s'(n, \epsilon) = O(n \log n)$  otherwise. Furthermore, if the standard tester has one-sided error, then so does the sample-based tester.*

**Proof:** The basic idea is to emulate the execution of the standard tester on a random isomorphic copy of the function that we are required to test. That is, given the standard tester  $T$ , and wishing to test the function  $f: D_n \rightarrow R_n$ , we select uniformly a permutation  $\pi: D_n \rightarrow D_n$  and emulate an execution of  $T$  while providing it with oracle access to the function  $f \circ \pi$ . That is, if  $T$  makes the query  $i$ , then we make the query  $\pi(i)$ . (We assume, without loss of generality, that  $T$  never makes the same query twice.) Thus, we perfectly emulate an execution of  $T^{f \circ \pi}$ , while making queries that

are almost uniformly and independently distributed,<sup>33</sup> and our decision regarding  $f$  is as good as  $T$ 's decision regarding  $f \circ \pi$ . Noting that the distance of  $f$  to  $\Pi_n$  equals the distance of  $f \circ \pi$  to  $\Pi_n$ , we are almost done. The only problem is that the sequence of  $q$  queries is uniformly distributed among all  $q$ -long sequences of *distinct* elements in  $D_n$ , whereas we wish it to be uniformly distributed in  $D_n^q$ .

We start by addressing this problem in the case of  $q \geq n/2$ . In this case, we merely use  $O(n \log n)$  independently and uniformly distributed queries in order to obtain the value of  $f$  on the entire domain. We may fail with some small constant probability, and in such a case we accept without taking any other action. Otherwise, we just decide based on our full knowledge of  $f$ .

Turning to the case of  $q < n/2$ , we first try to obtain the value of  $f$  on a uniformly distributed sequence of  $q$  distinct elements in  $D_n$  by making  $3q$  independently and uniformly distributed queries in  $D_n$ . With very high probability, we shall succeed, since for  $q < n/2$  it holds that a uniformly distributed sequence  $\bar{s} = (s_1, \dots, s_{3q}) \in D_n^{3q}$  contains at least  $q$  distinct elements (i.e.,  $|\{s_i : i \in [3q]\}| \geq q$ ). In case we fail, we accept without taking any other action. Otherwise, we emulate the execution of  $T$  using the first  $q$  distinct elements that we obtained. ■

## 6.2 The role of the range size

We observe that symmetric properties can be presented as properties of the frequencies of occurrences of the various values in the range. Specifically, let  $\Pi = \bigcup_{n \in \mathbb{N}} \Pi_n$  be a symmetric property such that  $\Pi_n$  contains functions from  $D_n$  to  $R_n$ , and let  $m = |R_n|$ . Actually, assume, without loss of generality, that  $D_n = [n]$  and  $R_n = [m]$ . Then, there exists a set  $S_n \subset [n]^m$  such that  $f \in \Pi_n$  if and only if  $(\text{cnt}_1(f), \dots, \text{cnt}_m(f)) \in S_n$ , where  $\text{cnt}_i(f) \stackrel{\text{def}}{=} |\{j \in [n] : f(j) = i\}|$ . It follows that testing  $\Pi_n$  amounts to testing whether the corresponding frequencies sequence is in the corresponding  $S_n$ , which means that we can expect the complexity to depend only on  $m$  and  $\epsilon$  (and not on  $n$ , at least for large enough  $n$ 's).

The foregoing observation suggests to consider the complexity of testing symmetric properties also as a function of the size of the range of the function (in addition to its dependence on the size of the domain). This requires presenting properties as sequences indexed by two (independent) size parameters,  $n$  and  $m$ , representing the sizes of the domain and range, respectively. That is, we consider classes of the form  $\Pi = \bigcup_{n, m \in \mathbb{N}} \Pi_{n, m}$  such that  $\Pi_{n, m}$  contains functions with domain  $D_n$  and range  $R_m$ , where  $|D_n| = n$  and  $|R_m| = m$ . The tester will be given both  $n$  and  $m$  (as well as the proximity parameter  $\epsilon$ ), and its sample complexity will be considered as a function of  $n, m$  and  $\epsilon$ . Note that complexity measures that depend on parameters other than the size of the domain are not uncommon in studies of properties of functions (see, e.g., the bounded-degree graph model of [GR02] and testing monotonicity of functions over arbitrary (product) domains and arbitrary ranges [GGL<sup>+</sup>00, DGL<sup>+</sup>99]).

**Theorem 6.2** (the complexity of testing symmetric functions is at most linear in the size of the range): *Let  $\Pi = \bigcup_{n, m \in \mathbb{N}} \Pi_{n, m}$  be a property of symmetric functions as above. Then the sample complexity of testing  $\Pi$ , denoted  $s: \mathbb{N}^2 \times (0, 1] \rightarrow \mathbb{N}$ , satisfies  $s(n, m, \epsilon) = O(m/\epsilon^2)$ .*

**Proof:** Let  $R_m = \{r_1, \dots, r_m\}$ . Generalizing the foregoing argument, we note that there exists a set  $S_{n, m} \subset [n]^m$  such that  $f \in \Pi_{n, m}$  if and only if  $(\text{cnt}_1(f), \dots, \text{cnt}_m(f)) \in S_{n, m}$ , where  $\text{cnt}_i(f) \stackrel{\text{def}}{=} |\{j \in D_n : f(j) = r_i\}|$ . Hence, a function  $f$  is  $\epsilon$ -far from  $\Pi_{n, m}$  if and only if for every  $(s_1, \dots, s_m) \in S_{n, m}$  it holds that  $\frac{1}{n} \sum_{i=1}^m |s_i - \text{cnt}_i(f)| > 2\epsilon$ . Consequently, distinguishing between the case that

<sup>33</sup>Indeed, here we use the equivalent formulation of sample-based testers as provided in Proposition 2.4

$f \in \Pi_{n,m}$  and the case that  $f$  is  $\epsilon$ -far from  $\Pi_{n,m}$  reduces to obtaining a  $\epsilon$ -approximation of the sequence  $(\text{cnt}_1(f), \dots, \text{cnt}_m(f))$ ; that is, a sequence  $(n_1, \dots, n_m)$  such that  $\frac{1}{n} \sum_{i=1}^m |n_i - \text{cnt}_i(f)| \leq \epsilon$ . Such an approximation can be obtained using a sample of size  $O(m/\epsilon^2)$ ; cf. [Dev83, Lemma 3]. Below, for the sake of illustration, we prove a somewhat weaker bound.

As observed in the foregoing discussion, the problem at hand is actually one of approximately learning a probability distribution over a set of  $m$  elements when given samples drawn independently from this distribution (i.e., testing whether  $f$  belongs to  $\Pi_{n,m}$  amounts to approximating the distribution in which the  $i^{\text{th}}$  element appears with probability  $\text{cnt}_i(f)/n$  with respect to the variation distance). We represent a probability distribution over  $[m]$  as a vector  $(p_1, \dots, p_m)$  such that  $p_i$  is the probability of the  $i^{\text{th}}$  element. We wish to find a vector  $(\tilde{p}_1, \dots, \tilde{p}_m)$  such that with probability at least  $2/3$ , it holds that  $\sum_{i \in [m]} |\tilde{p}_i - p_i| \leq \epsilon$ .

The obvious solution is to approximate each  $p_i$  up to an additive deviation of  $\gamma/m$ , but this will require a sample size of  $\Omega((m/\epsilon)^2)$ . Instead, letting  $I = \{i \in [m] : p_i \geq 1/m\}$  and  $\gamma = \epsilon/2$ , it suffices to have  $\tilde{p}_i = (1 \pm \gamma) \cdot p_i$  for every  $i \in I$  and  $\tilde{p}_i = p_i \pm \gamma/m$  for the other  $i$ 's (i.e.,  $i \in [m] \setminus I$ ), since

$$\sum_{i \in [m]} |\tilde{p}_i - p_i| = \sum_{i \in I} |\tilde{p}_i - p_i| + \sum_{i \in [m] - I} |\tilde{p}_i - p_i| \leq \sum_{i \in I} \gamma p_i + (m - |I|) \cdot \gamma/m < 2\gamma.$$

Using a sample of size  $s = O((m \log m)/\epsilon^2)$ , we obtain  $\Pr[\tilde{p}_i \neq (1 \pm \gamma)p_i] < 1/3m$  for each  $i \in I$ , and  $\Pr[\tilde{p}_i \neq p_i \pm \gamma/m] < 1/3m$  for each  $i \in [m] \setminus I$ . Both bounds are obtained by using a multiplicative Chernoff bound for Bernoulli trials with success probability  $p_i$ : For  $i \in I$  we consider a multiplicative deviation of  $\gamma$  and use  $p_i \geq 1/m$  (which implies  $\exp(-p_i \gamma^2 s/3) < 1/3m$ ), whereas for  $i \in [m] \setminus I$  we consider a multiplicative deviation of  $\delta = \gamma/m p_i$  and use  $p_i \cdot \delta^2 s = \epsilon^2/m^2 p_i s > \gamma^2 s/m = \Omega(\log m)$  (while assuming  $p_i \geq \gamma/m$ , whereas in case  $p_i < \gamma/m$ , we can just use the bound obtained for  $p_i = \gamma/m$ ). ■

### 6.3 Relation to testing distributions

The proof of Theorem 6.2 (implicitly) reduces the task of testing a symmetric property of functions to the task of testing distributions (to be reviewed next), which is in turn reduced to the task of (approximately) learning distributions. Our objective here is to go in the opposite direction; that is, to cast the model of testing distributions as a special case of testing symmetric properties of functions.

We stress that the model of testing (properties of) distributions has little in common with the model of testing (properties of) functions (especially, in its standard incarnation that refers to a fixed (i.e., uniform) distribution on the function's domain). Nevertheless, we show that the study of testing distributions can be reduced to the study of testing functions. Needless to say, this does not mean that one may not gain clarity and insight when working in the model of testing distributions. It only means that the latter model does fit within the more general framework of testing properties of functions.

Recall that in the context of testing distributions (cf., [BFR<sup>+</sup>00]), one is given samples of an unknown distribution and is asked to determine whether the distribution has some property (of distributions) or is statistically far from any distribution that has this property (i.e., the variation distance is large).<sup>34</sup> In the context of testing distributions, the key parameter is an upper bound

<sup>34</sup>The variation distance between  $D$  and  $D'$  is  $\frac{1}{2} \cdot \sum_v |\Pr_{r \sim D}[r=v] - \Pr_{r' \sim D'}[r'=v]|$ , where  $\Pr_{r \sim D}[r=v]$  denotes the probability that an element distributed according to  $D$  equals  $v$ .

on the size of the support of the distribution. Thus, we shall consider properties of the form  $\Pi = \bigcup_{m \in \mathbb{N}} \Pi_m$ , where  $\Pi_m$  is a class of distributions over a predetermined set  $S_m$  of size  $m$ . In the testing problem, the tester is given the parameter  $m$ , a distance parameter  $\epsilon$ , and samples from an unknown distribution  $D$  with a support that is a subset of  $S_m$  (i.e.,  $D$  assigns positive probability mass only to elements of  $S_m$ ).

**Definition 6.3** (testing distributions, following [BFR<sup>+</sup>00]): *We say that a distribution  $D$  is  $\epsilon$ -far from  $\Pi_m$  if for every  $D'$  in  $\Pi_m$  the variation distance between  $D$  and  $D'$  is at least  $\epsilon$  (i.e.,  $\frac{1}{2} \cdot \sum_v |\Pr_{r \sim D}[r = v] - \Pr_{r' \sim D'}[r' = v]| \geq \epsilon$ ). Let  $\Pi = \bigcup_{m \in \mathbb{N}} \Pi_m$  be such that each distribution in  $\Pi_m$  has support that is a subset of  $S_m$  and  $|S_m| = m$ . A distribution tester of sample complexity  $s: \mathbb{N} \times (0, 1] \rightarrow \mathbb{N}$  for  $\Pi$  is a probabilistic algorithm  $T$  that satisfies the following two conditions for every  $m \in \mathbb{N}$  and  $\epsilon > 0$ :*

1. *For every distribution  $D$  in  $\Pi_m$ , it holds that*

$$\Pr_{r_1, \dots, r_s \sim D}[T(m, \epsilon, (r_1, \dots, r_s)) = 1] \geq \frac{2}{3},$$

*where  $s = s(m, \epsilon)$  and  $r_1, \dots, r_s$  are drawn independently such that each  $r_i$  is distributed according to  $D$ .*

2. *For every distribution  $D$  with support that is a subset of  $S_m$ , if  $D$  is  $\epsilon$ -far from  $\Pi_m$  then*

$$\Pr_{r_1, \dots, r_s \sim D}[T(m, \epsilon, (r_1, \dots, r_s)) = 0] \geq \frac{2}{3},$$

*where again  $s = s(m, \epsilon)$  and  $r_1, \dots, r_s$  are drawn independently such that each  $r_i$  is distributed according to  $D$ .*

Recall that any property of distributions as in Definition 6.3 can be tested using a sample of size  $O(m/\epsilon^2)$ , see [Dev83, Lemma 3] as well as the proof of Theorem 6.2 (which establishes a somewhat weaker bound). This upper bound is obtained via a learning algorithm, and the study of distribution testing is focused on improving upon this upper bound.

It is convenient to describe distributions as above by  $m$ -long sequences of non-negative number that sum-up to 1; that is, the sequence  $(p_1, \dots, p_m)$  represents a distribution in which the  $i^{\text{th}}$  element of  $S_m$  occurs with probability  $p_i$ . Thus, variation distance (or statistical distance) between distribution corresponds to half the  $L_1$ -norm of the difference between the two corresponding sequences. This is indeed closely related to the frequency sequences used in Section 6.2.

The basic intuition is that testing a property of distributions  $\Pi = \bigcup_{m \in \mathbb{N}} \Pi_m$  is the same as (sample-based) testing a property of functions  $\Pi' = \bigcup_{n, m \in \mathbb{N}} \Pi'_{n, m}$  that contains all functions having a frequency sequence that corresponds to a distribution in  $\Pi$ . Note that the property  $\Pi'$  is symmetric by its very definition. The foregoing intuition is sound provided that (1) it does not matter that the frequencies that correspond to functions with domain  $D_n$  are restricted to be multiples of  $1/n$ , and (2) the  $\Pi'$ -tester gains little by being provided also with uniformly distributed arguments of the function (rather than only with the value of the function on such arguments). Both conditions hold provided that  $n$  is sufficiently large (i.e., larger than some fixed polynomial in  $m/\epsilon$ ).

**Theorem 6.4** (distribution testing reduces to sample-based testing of symmetric properties): *Let  $\Pi = \bigcup_{m \in \mathbb{N}} \Pi_m$  be a property of distributions such that all distributions in  $\Pi_m$  have a support that is a subset of  $S_m$ , where  $|S_m| = m$ . Denote the sample complexity of testing  $\Pi$  by  $s: \mathbb{N} \times (0, 1] \rightarrow \mathbb{N}$ . Then, there exists a symmetric property of functions  $\Pi' = \bigcup_{n, m \in \mathbb{N}} \Pi'_{n, m}$  such that the following two conditions hold:*

1. All functions in  $\Pi'_{n,m}$  have domain  $[n]$  and range  $S_m$ .
2. There exists a constant  $c$  such that for every  $n \geq c \cdot m^2 / \epsilon^4$ , the sample complexity of testing  $\Pi'$ , denoted  $s': \mathbb{N}^2 \times (0, 1) \rightarrow \mathbb{N}$ , satisfies  $s(m, \epsilon) = O(s'(n, m, \epsilon/2))$  and  $s(m, \epsilon) = \Omega(s'(n, m, 2\epsilon))$ .

Indeed, the sample complexity of testing the function class will not depend on the standard (domain) size parameter  $n$ , but rather on the auxiliary parameter  $m$  (representing the size of the range). The proof of Theorem 6.4 provides reductions between settings of parameters for one problem (e.g.,  $m$  and  $\epsilon$  for  $\Pi$ ) and settings of parameters for the other problem (e.g.,  $n, m$  and  $\epsilon$  for  $\Pi'$ ). In particular, for any value of  $\epsilon$ , the equivalence is between  $\epsilon$ -testing  $\Pi_m$  and  $\Theta(\epsilon)$ -testing  $\Pi'_{\text{poly}(m/\epsilon), m}$ .

**Proof:** The proof is based on the association of properties of distributions over  $S_m$  with symmetric properties of function that range over  $S_m$ . The symmetric condition allows to disregard the identity of the argument that evaluates to a specific value, provided that the sample complexity is smaller than the square root of the domain size (since otherwise the issue of sampling without repetitions gets in our way). For the association to work, we need to assume that the tested distributions are of the form  $(q_1, \dots, q_m)$  such that each  $q_v$  is a multiple of  $1/n$ . In this case the said distribution can be associated with a function  $f: [n] \rightarrow S_m$  such that for every  $v \in S_m$  it holds that  $|\{j \in [n] : f(j) = v\}| = q_v \cdot n$ . For  $n$  that is sufficiently large (w.r.t.  $m$  and  $1/\epsilon$ ), the statistical distance between any distribution on  $S_m$  and such a distribution  $(q_1, \dots, q_m)$  is sufficiently small.

Fixing the property (of distributions)  $\Pi$  as in the hypothesis, we shall consider the following (symmetric) property (of functions)  $\Pi' = \bigcup_{n,m \in \mathbb{N}} \Pi'_{n,m}$ : *A function  $f: [n] \rightarrow S_m$  is in  $\Pi'_{n,m}$  if and only if there exists a distribution  $(p_1, \dots, p_m)$  in  $\Pi_m$  such that the max-norm distance between  $(p_1 \cdot n, \dots, p_m \cdot n)$  and  $(\text{cnt}_1(f), \dots, \text{cnt}_m(f))$  is at most 1, where, as before,  $\text{cnt}_i(f)$  is the number of elements  $j \in [n]$  such that  $f(j)$  equals the  $i^{\text{th}}$  elements in  $S_m$ .*

**Claim 6.4.1** *For every  $n \geq \max(20s(m, \epsilon/2), 2m/\epsilon)$ , it holds that  $s'(n, m, \epsilon) = O(s(m, \epsilon/2))$ .*

**Proof:** Fixing  $\epsilon$ , consider a distribution tester, denoted  $T_{\text{ds}}$ , for  $\Pi_m$  that uses  $t = s(m, \epsilon/2)$  samples. We construct an  $\epsilon$ -tester, denoted  $T_{\text{fn}}$ , for functions in  $\Pi'_{n,m}$ . The tester  $T_{\text{fn}}$  is given a sequence of uniformly distributed samples that are labeled by some function  $f: [n] \rightarrow S_m$ ; that is,  $T_{\text{fn}}$  is given a sequence  $((r_1, f(r_1)), \dots, (r_t, f(r_t)))$ , where  $r_1, \dots, r_t$  are uniformly and independently distributed in  $[n]$ . Suppose that  $T_{\text{fn}}$  just invokes  $T_{\text{ds}}$  on  $(f(r_1), \dots, f(r_m))$ , and outputs whatever  $T_{\text{ds}}$  does. Then, effectively,  $T_{\text{fn}}$  invokes  $T_{\text{ds}}$  on samples taken from the distribution that is described by  $\bar{q} = (q_1, \dots, q_m)$  such that  $q_v = \Pr_{r \in [n]}[f(r) = v]$ .

Assume first that  $f \in \Pi'_{n,m}$ . Then, there exists  $\bar{p} = (p_1, \dots, p_m)$  in  $\Pi_m$  such that  $\max_{i \in [m]} \{|p_i - q_i|\} \leq 1/n$ . Hence, the variation distance between a sample of  $t$  elements from  $\bar{q}$  and  $t$  elements from  $\bar{p}$  is at most  $t/n \leq 1/20$ . It follows that  $T_{\text{ds}}$  accepts (when invoked with samples of  $\bar{q}$ ) with probability at least  $(2/3) - (t/n) > 0.6$ , and hence  $T_{\text{fn}}$  accepts  $f$  with probability at least 0.6. (We shall use straightforward error reduction in order to regain the error bound of  $1/3$ .)

On the other hand, if  $f: [n] \rightarrow S_m$  is  $\epsilon$ -far from  $\Pi'_{n,m}$ , then  $\bar{q}$  is  $\epsilon/2$ -far from any distribution in  $\Pi_m$ , because otherwise  $\bar{q}$  is  $((\epsilon/2) + (m/n))$ -close to a distribution  $\bar{p} \in \{i/n : i \in \{0, \dots, n\}\}^m$  that is  $(m/n)$ -close to  $\Pi_m$  (and contradiction follows since  $m/n \leq \epsilon/2$ , whereas  $\bar{p}$  corresponds to the frequency sequence of a function  $g: [n] \rightarrow S_m$  in  $\Pi'_{n,m}$ ). In this case,  $T_{\text{ds}}$  (and so also  $T_{\text{fn}}$ ) rejects with probability at least  $2/3$ . ■

**Claim 6.4.2** *For every  $n = \Omega(m^2/\epsilon^4)$ , it holds that  $s(m, \epsilon) = O(s'(m, n, \epsilon/2))$ .*

**Proof:** Fixing  $\epsilon$ , consider a function tester, denoted  $T_{\text{fn}}$ , for  $\Pi'_{n,m}$  that uses  $t = s'(n, m, \epsilon/2)$  samples. By Theorem 6.2, it holds that  $s'(n, m, \epsilon) = O(m/\epsilon^2)$ . We construct an  $\epsilon$ -tester, denoted  $T_{\text{ds}}$ , for distributions in  $\Pi_m$ . The tester  $T_{\text{ds}}$  is given a sequence, denoted  $(v_1, \dots, v_t)$ , of samples from distribution  $\bar{p} = (p_1, \dots, p_m)$  that ranges over  $S_m$ . Suppose that  $T_{\text{ds}}$  invokes  $T_{\text{fn}}$  on  $((r_1, v_1), \dots, (r_t, v_t))$ , where  $r_1, \dots, r_t$  are uniformly and independently distributed in  $[n]$ . Since  $t < \sqrt{n}/3$ , these  $r_i$ 's are distinct with probability at least 0.9.

Let  $\bar{q} = (q_1, \dots, q_m) \in \{i/n : i \in \{0, \dots, n\}\}^m$  be a vector that approximates  $\bar{p}$  in the sense that  $\max_{i \in [m]} \{|p_i - q_i|\} \leq 1/n$ . Then,  $\bar{q}$  is  $m/n$ -close to  $\bar{p}$ . It will be instructive to consider an invocation of  $T_{\text{ds}}$  on a sample drawn according to  $\bar{q}$ . In such a case, essentially (i.e., ignoring the event of colliding  $r_i$ 's), the tester  $T_{\text{ds}}$  invokes  $T_{\text{fn}}$  on random samples labeled by some function  $f: [n] \rightarrow S_m$  such that  $\Pr_{r \in [n]}[f(r) = v] = q_v$ .

Now, if  $\bar{p}$  is in  $\Pi_m$ , then  $f \in \Pi'_{n,m}$ . Furthermore, the variation distance between a sample of  $t$  elements from  $\bar{q}$  and  $t$  elements from  $\bar{p}$  is at most  $t/n < 0.01$ , since  $n = \Omega(t^2)$ . Thus, when  $T_{\text{fn}}$  is given a sequence generated according to  $\bar{p}$  (augmented with uniformly selected  $r_i$ 's), it accepts with probability at least  $(2/3) - (t/n) - 0.1 > 0.6$ , where the 0.1 term is due to the event of colliding  $r_i$ 's. It follows that  $T_{\text{ds}}$  accepts  $\bar{p}$  with probability at least 0.6. (Again, we shall use straightforward error reduction in order to regain the error bound of 1/3.)

On the other hand, if  $\bar{p}$  is  $\epsilon$ -far from  $\Pi_m$ , then  $f$  is  $\epsilon/2$ -far from  $\Pi'_{n,m}$ . This holds because otherwise  $g \in \Pi'_{n,m}$  that is  $\epsilon/2$ -close to  $f$  yields a distribution  $\bar{p}_g$  that is  $m/n$ -close to  $\Pi_m$  as well as  $((\epsilon/2) + (m/n))$ -close to  $\bar{p}$  (since it is  $\epsilon/2$ -close to the frequency vector of  $f$ ), which implies that  $\bar{p}$  is  $((\epsilon/2) + (2m/n))$ -close to  $\Pi_m$  and contradiction follows since  $m/n < \epsilon/4$ . Now, switching (again) between  $\bar{q}$  and  $\bar{p}$ , we infer that  $T_{\text{fn}}$  rejects with probability at least  $(2/3) - (t/n) - 0.1 > 0.6$  when given a sequence generated according to  $\bar{p}$  (augmented with uniformly selected  $r_i$ 's), and it follows that  $T_{\text{ds}}$  rejects  $\bar{p}$  with probability at least 0.6. ■

This completes the proof of the theorem. ■

## 7 Distribution-Free Testing

Recall that distribution-free testing is a generalization of the standard notion of testing functions where the distance measure is defined according to an arbitrary distribution of the function's domain. In this case, the tester is given samples drawn according to this (unknown) distribution, where these samples are labeled by the tested function. We stress that the task is to test properties of the function, not of the distribution.<sup>35</sup>

For a distribution  $\mu$ , by writing  $r_1, \dots, r_t \sim \mu$  we mean that  $r_1, \dots, r_t$  are selected independently such that each  $r_i$  is distributed according to  $\mu$ . Denoting by  $\mu^t$  the  $t$ -wise Cartesian product of  $\mu$ , we can also write  $(r_1, \dots, r_t) \sim \mu^t$ .

**Definition 7.1** (distribution-free sample-based tester, following [GGR98]): *Let  $\Pi$  and  $s: \mathbb{N} \times (0, 1] \rightarrow \mathbb{N}$  be as in Definition 2.3. A (sample-based) distribution-free tester of sample complexity  $s$  for a property  $\Pi$  is a probabilistic algorithm  $T$  that satisfies the following two conditions for every  $n \in \mathbb{N}$  and every distribution  $\mu_n$  over  $D_n$ :*

1. For every  $f \in \Pi_n$  (and every  $\epsilon > 0$ ), it holds that

$$\Pr_{r_1, \dots, r_s \sim \mu_n}[T(n, \epsilon, (r_1, f(r_1)), \dots, (r_s, f(r_s))) = 1] \geq \frac{2}{3},$$

---

<sup>35</sup>Indeed, the latter subject (i.e., testing distributions) is discussed in Section 6.3.

where  $s = s(n, \epsilon)$  and  $(r_1, \dots, r_s)$  is distributed according to  $\mu_n^s$ .

2. For every  $\epsilon > 0$ , and every  $f: D_n \rightarrow R_n$  that is  $\epsilon$ -far from  $\Pi_n$  with respect to the distribution  $\mu_n$  (i.e.,  $\delta_{\Pi_n}^{\mu_n}(f) > \epsilon$ ), it holds that

$$\Pr_{r_1, \dots, r_s \sim \mu_n} [T(n, \epsilon, (r_1, f(r_1)), \dots, (r_s, f(r_s))) = 0] \geq \frac{2}{3},$$

where again  $s = s(n, \epsilon)$  and  $(r_1, \dots, r_s)$  is distributed according to  $\mu_n^s$ . We say that  $f: D_n \rightarrow R_n$  is  $\epsilon$ -far from  $\Pi_n$  with respect to the distribution  $\mu_n$  if  $\delta_{\Pi_n}^{\mu_n}(f) > \epsilon$ , where

$$\delta_{\Pi_n}^{\mu_n}(f) \stackrel{\text{def}}{=} \inf_{g \in \Pi_n} \{\Pr_{x \sim \mu_n} [f(x) \neq g(x)]\}. \quad (7)$$

Again, if the tester accepts every function in  $\Pi$  with probability 1, then we say that it has one-sided error.

Definition 2.3 refers to the case that  $\mu_n$  is uniform on  $D_n$ . In what follows, the (unqualified) term “sample-based tester” refers to testers that work only for this (uniform) distribution, as defined in Definition 2.3.

## 7.1 A general result

We start by proving the first item in Theorem 1.3. Namely, we show that the sample complexity of distribution-free testing cannot be significantly smaller than the sample complexity of one-sided error testing (with respect to the uniform distribution). Specifically, we show that the latter is at most nearly-quadratic in the former.

**Theorem 7.2** (distribution-free testers imply one-sided error testers under the uniform distribution): *For every property  $\Pi$ , if  $\Pi$  has a sample-based distribution-free tester of sample complexity  $s: \mathbb{N} \times (0, 1] \rightarrow \mathbb{N}$ , then  $\Pi$  has a one-sided error sample-based tester of sample complexity  $\tilde{O}(s^2)$  (under the uniform distribution).*

The sample complexity of the one-sided error sample-based tester is actually  $O(s^2)$  if  $s < \sqrt{n}/25$  and  $O(n \log n)$  otherwise. While an upper bound of  $O(s^2)$  is quite likely, we shall show (in Theorem 7.4) that an upper bound of  $O(s)$  cannot hold in general. We stress that Theorem 7.2 (as well as the last sentence) refer to an upper bound that holds for *any* property (rather than to specific cases, as studied in Section 7.2).

**Proof:** Let  $T_{\text{df}}$  be a sample-based distribution-free tester of sample complexity  $s: \mathbb{N} \times (0, 1] \rightarrow \mathbb{N}$  for  $\Pi$ . If  $s = \Omega(\sqrt{n})$ , then with high probability, a labeled (uniformly distributed) sample of size  $O(s^2 \log s)$  allows to fully reconstruct the tested function. This directly implies the existence of a one-sided error sample-based tester of sample complexity  $\tilde{O}(s^2)$ . Therefore, assume that  $s$  is strictly smaller than  $\sqrt{n}$ , and in particular,  $s < \sqrt{n}/25$ . We may assume that the error probability of  $T_{\text{df}}$  is at most  $1/6$ , rather than at most  $1/3$ , since this can be obtained by straightforward error reduction, which merely increases  $s$  by a constant factor.

Consider the following sample-based algorithm  $T$ , of sample complexity  $t = 500s^2$  (observe that  $t < n$ ). On input parameters  $n, \epsilon$  and a sequence  $((r_1, v_1), \dots, (r_t, v_t))$ , where the  $v_i$ 's are determined by the tested function, algorithm  $T$  accepts if and only if there exists  $g \in \Pi_n$  such that

$g(r_i) = v_i$  for every  $i \in [t]$ . We claim that  $T$  is a one-sided error tester for  $\Pi$  under the uniform distribution.

By its construction,  $T$  accepts every function in  $\Pi$  with probability 1. Hence, if  $T$  is not a one-sided error tester for  $\Pi$  (under the uniform distribution), then it must be that it accepts some no-instance (i.e., a function far from  $\Pi$ ) with too high probability (i.e., probability exceeding  $1/3$ ). Thus, we assume towards a contradiction that there exist  $n, \epsilon$  and  $f: D_n \rightarrow R_n$  such that  $f$  is  $\epsilon$ -far from  $\Pi_n$  (under the uniform distribution) and  $T$  accepts  $f$  with probability at least  $1/3$  (when given  $t = t(n, \epsilon)$  uniformly distributed samples that are labeled according to  $f$ ). In other words, for at least one third of the possible choices of  $\bar{r} = (r_1, \dots, r_t) \in D_n^t$ , there exists  $g_{\bar{r}} \in \Pi_n$  such that  $g_{\bar{r}}(r_i) = f(r_i)$  for every  $i \in [t]$ .

For each such  $\bar{r} = (r_1, \dots, r_t)$ , consider the distribution  $\mu_{\bar{r}}$  that is uniform over the corresponding (multi)set  $\{r_1, \dots, r_t\}$ . The key observation is that  $T_{\text{df}}$  must accept a sequence of  $s$  samples drawn from  $\mu_{\bar{r}}$  that are labeled according to  $f$  with probability at least  $5/6$ , since these labels are consistent with  $g_{\bar{r}} \in \Pi$ . (The fact that  $f$  is  $\epsilon$ -far from  $\Pi$  according to the uniform distribution on  $D_n$  is irrelevant here, and in particular  $\delta_{\Pi_n}^{\mu_{\bar{r}}}(f) = 0$ .)<sup>36</sup> Thus, the following holds.

**Claim 7.2.1** *Let  $T$  and  $f$  be as in the foregoing paragraph, and let  $X = (X_1, \dots, X_s)$  denote a random sequence produced by first selecting uniformly  $\bar{r} \in D_n^t$ , and then picking  $s$  independent samples from the distribution  $\mu_{\bar{r}}$ . Then, on input  $n, \epsilon$  and  $((X_1, f(X_1)), \dots, (X_s, f(X_s)))$ , algorithm  $T_{\text{df}}$  accepts with probability at least  $\frac{1}{3} \cdot \frac{5}{6}$ .*

On the other hand, on input  $n, \epsilon$  and  $((Y_1, f(Y_1)), \dots, (Y_s, f(Y_s)))$ , where the  $Y_i$ 's are independently and uniformly distributed in  $D_n$ , algorithm  $T_{\text{df}}$  must reject with probability at least  $5/6$ , because  $f$  is  $\epsilon$ -far from  $\Pi_n$  (under the uniform distribution). However, (as shown next in Claim 7.2.2) the two distributions (i.e.,  $X$  and  $Y$ ) are 0.1-close (recall that  $t = 500s^2$  by the definition of  $T$  and that we are analyzing the case that  $s < \sqrt{n}/25$ ). This yields a contradiction, because it implies that on input  $n, \epsilon$  and  $((X_1, f(X_1)), \dots, (X_s, f(X_s)))$ , algorithm  $T_{\text{df}}$  rejects with probability at least  $5/6 - 0.1 = 11/15$ , while by Claim 7.2.1 is rejects with probability at most  $1 - 5/18$ .

**Claim 7.2.2** *Suppose that  $500s^2 \leq t < n$ , and let  $X = (X_1, \dots, X_s)$  and  $Y = (Y_1, \dots, Y_s)$  be as above; that is,  $X$  is as defined in Claim 7.2.1 and  $Y$  is a sequence of independently and uniformly distributed elements of  $D_n$ . Then, the variation distance between  $X$  and  $Y$  is at most 0.1.*

**Proof:** Note that conditioned on  $X$  containing no collision (i.e.,  $X_i = X_j$  for  $i \neq j$ ), the sequence  $X$  is uniformly distributed over all  $s$ -long sequences of distinct elements of  $D_n$ . The same holds for  $Y$ , and thus it suffices to prove that, in both  $X$  and  $Y$ , the probability that the same element appears twice is at most  $1/20$ . The claim is obvious for  $Y$ , since the collision probability is at most  $\binom{s}{2} \cdot \frac{1}{n} < 1/1000$ .

Turning to  $X$ , we note that the argument is equally easy in the case that  $40t^2 < n$  (since  $\binom{t}{2} \cdot \frac{1}{n} + \binom{s}{2} \cdot \frac{1}{t} < 1/80 + 1/1000$ ), alas we only have  $t < n$ . Nevertheless, we first upper bound by  $1/40$  the probability that a uniformly chosen sequence  $\bar{r} = (r_1, \dots, r_t) \in D_n^t$  contains more than  $20t$  collisions (i.e., pairs  $(i, j)$  such that  $r_i = r_j$ ). This follows by noting that the expected number of collisions is  $\binom{t}{2} \cdot \frac{1}{n} < \frac{t}{2}$ . Fixing an arbitrary sequence  $\bar{r}$  with at most  $20t$  collisions, we consider the probability that a sample of  $s$  elements taken from  $\bar{r}$  has a collision. Such collisions may be due either to a collision in  $\bar{r}$  or to picking the same position in  $\bar{r}$  twice. Thus, this probability is at most  $\binom{s}{2} \cdot \left(\frac{20t}{t^2} + \frac{1}{t}\right) < 1/40$ , and the claim follows. ■

<sup>36</sup>We stress that our argument is not based on the fact that  $\delta_{\Pi_n}^{\mu_{\bar{r}}}(f) = 0$  but rather on the fact that the samples labeled by  $f$  can be thought to be labeled by  $g_{\bar{r}} \in \Pi$ .

Since we have addressed both the case that  $s \geq \sqrt{n}/25$  and the case that  $s < \sqrt{n}/25$ , the theorem follows. ■

**Remark 7.3** (a general phenomenon): *The proof of Theorem 7.2 is actually oblivious to the specifics of the promise problem at stake (i.e., the fact that we refer to property testing). Indeed, it can be stated for any promise problem  $(P, Q)$ , where the promise is that the input is in  $P$  and the question is whether it is in  $Q$ :*

*If a promise problem  $(P, Q)$  has a sample-based distribution-free two-sided error decider of sample complexity  $s$ , then  $(P, Q)$  has a one-sided error sample-based decider of sample complexity  $\tilde{O}(s^2)$  (under the uniform distribution).*

*The proof shows that if  $f \in P \setminus Q$  is accepted with probability greater than  $1/3$  by the canonical one-sided error decider, which accepts a function if and only if the sample is consistent with some  $g \in P \cap Q$ , then a distribution-free decider for  $(P, Q)$  cannot have sample complexity  $s$ , because when testing  $f$  it must err with probability greater than  $1/6$  either on the uniform distribution on some set of size  $O(s^2)$  or on the uniform distribution on  $f$ 's entire domain.*

## 7.2 Specific results

In this section we consider several cases that indicate that there is no fixed relation between the sample complexity of one-sided error testers under the uniform distribution and the sample complexity of distribution-free testers. The first three cases are based on natural properties. We recall the notation that was introduced in the introduction:  $\text{OSE}(\Pi)$  denotes the sample complexity of one-sided error sample-based testing  $\Pi$  (under the uniform distribution); that is,  $\text{OSE}(\Pi)$  is a function  $s(n, \epsilon)$  that represents the size of the sample taken by the best such tester on input parameters  $n$  and  $\epsilon$ . Likewise,  $\text{DF}(\Pi)$  denotes the sample complexity of distribution-free (sample-based) testing  $\Pi$ . Indeed, Theorem 7.2 asserts that for every  $\Pi$  it holds that  $\text{OSE}(\Pi) = \tilde{O}(\text{DF}(\Pi)^2)$ .

We first show that a general upper bound as in Theorem 7.2 (i.e., an upper bound on  $\text{OSE}$  in terms of  $\text{DF}$ ) cannot be linear (let alone sub-linear). That is, there exist (natural) properties  $\Pi$  such that  $\text{OSE}(\Pi) = \Omega(\text{poly}(\epsilon) \cdot \text{DF}(\Pi) \log \text{DF}(\Pi))$ , as stated in Item 2.

**Theorem 7.4** (distribution-free sample-complexity may be lower than the one-sided error sample complexity under the uniform distribution): *There exists a property  $\Pi$  such that*

1. *Any one-sided error tester for  $\Pi$  (under the uniform distribution) must have query complexity  $q(n, \epsilon) = \Omega(n)$  for every  $\epsilon < 1/3$ .*
2. *There exists a (sample-based) distribution-free tester for  $\Pi$  of sample complexity  $s(n, \epsilon) = \text{poly}(1/\epsilon) \cdot \frac{n}{\log n}$ .*

*For example,  $\Pi_n$  may be the set of Boolean functions over  $[n]$  that evaluate to 1 more often than to 0.*

It follows that there are properties for which the one-sided error sample complexity (under the uniform distribution) is higher than the distribution-free sample complexity. Note that  $\Pi$  is a symmetric property, and so queries are not more powerful for it than samples (see Theorem 6.1).

**Proof:** We consider the class, denoted  $\Pi_n$ , of Boolean functions over  $[n]$  such that  $f: [n] \rightarrow \{0, 1\}$  is in  $\Pi_n$  if and only if  $|\{i \in [n] : f(i) = 1\}| > n/2$ .

**Claim 7.4.1** *Any one-sided error tester for  $\Pi_n$  (under the uniform distribution) has query complexity  $q(n, \epsilon) \geq n/2$  for any  $\epsilon < 1/2$ .*

**Proof:** A one-sided error tester may reject only if the answers it has obtained are not consistent with any function in  $\Pi_n$ , which means that it must see at least  $n/2$  zeros. On the other hand, when having oracle access to a function with at least  $(0.5 + \epsilon) \cdot n$  zeros, which is  $\epsilon$ -far from  $\Pi_n$ , the tester must reject. The claim follows. ■

**Claim 7.4.2** *There exists a (sample-based) distribution-free tester of sample complexity  $s(n, \epsilon) = \tilde{O}(1/\epsilon^2) \cdot \frac{n}{\log n}$  for  $\Pi$ .*

**Proof:** We shall rely on an algorithm of Valiant and Valiant [VV11], which approximates the support size (as well as the entropy) of an unknown distribution over  $[n]$  based on  $O(n/\log n)$  samples of the distribution. Specifically, we use the following result.

**Theorem 1 in [VV11]** (rephrased)<sup>37</sup>: *There exists an algorithm  $A_{\text{vv}}$  that on input a parameter  $\delta$  and  $\tilde{O}(1/\delta^2) \cdot n/\log n$  samples of a random variable, denoted  $X_n$ , that ranges over  $[n]$ , outputs a representation of a distribution  $Y_n$  that with overwhelmingly high probability is  $\delta$ -close to a distribution  $X'_n$  that equals  $X_n$  up to relabeling.*

By saying that  $X$  equals  $X'$  up to relabeling we mean that there exists a permutation  $\pi: [n] \rightarrow [n]$  such that for every  $v \in [n]$  it holds that  $\Pr_{x \sim X'}[x = v] = \Pr_{x \sim X}[x = \pi(v)]$ .

(Actually, the representation of  $Y_n$  is succinct so to allow the algorithm to run in time that is linear in the number of samples that it obtains, but this is irrelevant to us here.) Now, suppose that, on input  $n$  and  $\epsilon$ , we get  $s = s(n, \epsilon)$  samples drawn according to an unknown distribution  $\mu_n$  and labeled by an unknown Boolean function  $f$ . Our task is to decide whether  $f$  is in  $\Pi_n$  or is  $\epsilon$ -far from  $\Pi_n$  with respect to the distribution  $\mu_n$ . We propose the following algorithm.

1. The algorithm uses a small portion of the label sample (i.e.,  $O(1/\epsilon^2)$  samples) in order to approximate the probability  $p_0 \stackrel{\text{def}}{=} \Pr_{r \sim \mu_n}[f(r) = 0]$  up to an additive error of  $\epsilon/4$ . If the estimate, denoted  $\tilde{p}$ , is below  $\epsilon/2$ , then the algorithm halts with output 1 (i.e., it accepts).

Assuming the algorithm did not halt in Step 1, we may assume that  $p_0 > \epsilon/4$ .

2. The algorithm defines  $X_n$  as the distribution  $\mu_n$  conditioned on  $f$  evaluating to 0; that is,  $\Pr_{x \sim X_n}[x = v] = \Pr_{x \sim \mu_n}[x = v | f(x) = 0]$ . It invokes algorithm  $A_{\text{vv}}$  with parameter  $\delta = \epsilon/(4\tilde{p})$  providing it with  $\tilde{O}(1/\delta^2) \cdot n/\log n$  samples of  $X_n$ , and obtaining a description of a distribution  $Y_n$ .

Samples of  $X_n$  are obtained from the  $s$  labeled samples provided to the main algorithm; specifically, the main algorithm forwards the first  $s_0 \stackrel{\text{def}}{=} \tilde{O}(1/\delta^2) \cdot n/\log n$  samples that are labeled 0 to  $A_{\text{vv}}$ . With very high probability, the number of samples labeled 0 is at least  $p_0 \cdot s/2 \geq \tilde{p} \cdot s/4 = \epsilon s/(16\delta)$ , which is at least  $s_0$ , because we can set  $s_0 < (\delta/\epsilon)^2 \cdot s$  (which is smaller than  $(\epsilon/16\delta) \cdot s$ , since  $\delta \leq \epsilon/4$ ).

---

<sup>37</sup>The statement in [VV11, Theorem 1] is for the *relative earth-mover distance* between the histograms corresponding to the distributions (see Definitions 4 and 5 in [VV11]). However (as stated following Definition 5 in [VV11]), this gives an upper bound on the variation distance as stated here. A proof of this fact appears in [VV15, Appendix A] (as the special case of  $\tau = 0$ ).

3. The algorithm accepts if and only if  $Y_n$  is  $\delta$ -close to a distribution that has support size smaller than  $n/2$ .

We now turn to the analysis of the foregoing algorithm. Consider first the case that  $f \in \Pi_n$ . If  $p_0 \leq \epsilon/4$ , then, with high constant probability, it holds that  $\tilde{p} < \epsilon/2$  (by an additive Chernoff bound), and the algorithm accepts. Otherwise (i.e.,  $p_0 > \epsilon/4$ ), with high constant probability, it holds that  $\tilde{p} \leq 2p_0$  (by a multiplicative Chernoff bound). Assuming the algorithm does not accept due to  $\tilde{p}$  being smaller than  $\epsilon/2$  (so that it continues to its second step), with high probability the distribution  $Y_n$  (output by  $A_{vv}$ ) is  $\delta$ -close to a relabeling of  $X_n$ , which in turn has support size smaller than  $n/2$ . Hence, in either cases, the algorithm accepts with high probability.

Now suppose that  $f$  is  $\epsilon$ -far from  $\Pi_n$  with respect to the distribution  $\mu_n$ . In this case  $p_0 > \epsilon$ , since otherwise  $f$  is  $\epsilon$ -close to the all-ones function, which is in  $\Pi_n$ . Now, with high probability, it holds that  $\tilde{p} > p_0/2 \geq \epsilon/2$ , which means that the algorithm proceeds to Step 2. Furthermore,  $X_n$  is  $\epsilon/p_0$ -far from having support size smaller than  $n/2$ , since otherwise we can modify  $f$  in a corresponding manner in contradiction to the hypothesis (i.e., if  $X_n$  is  $\epsilon/p_0$ -close to  $X'_n$  that has support smaller than  $n/2$ , then setting  $f'(x) = 1$  if and only if  $x$  is not in the support of  $X'_n$  yields a function in  $\Pi_n$  that is  $\epsilon$ -close to  $f$  w.r.t the distribution  $\mu_n$ ). Recalling that  $\delta = \epsilon/4\tilde{p}$  (and  $\tilde{p} > p_0/2$ , which implies that  $\epsilon/p_0 > \epsilon/2\tilde{p} = 2\delta$ ), it follows that  $X_n$  is  $2\delta$ -far from having support size smaller than  $n/2$ . Now, with very high probability, the distribution  $Y_n$  (output by  $A_{vv}$ ) is  $\delta$ -close to a relabelling of  $X_n$ , which implies that  $Y_n$  is  $(2\delta - \delta)$ -far from any distribution having support smaller than  $n/2$ , which in turn implies that our algorithm rejects. ■

Theorem 7.4 follows. ■

**On the non-tightness of Theorem 7.2.** Having just shown that OSE may be larger than DF, we now turn to cases in which OSE is at most linear in DF, which in particular means that in these cases Theorem 7.2 is not tight. The first case is of a property  $\Pi$  such that  $\text{OSE}(\Pi) = \Theta(\text{DF}(\Pi))$  holds (Item 3 in Theorem 1.3), and the other cases are ones in which OSE is significantly smaller than DF.

**Theorem 7.5** (distribution-free sample-complexity may equal the one-sided error sample complexity under the uniform distribution): *There exists a property  $\Pi$  such that*

1. *Any sample-based tester for  $\Pi$  (under the uniform distribution) has sample complexity  $s(n, \epsilon) = \Omega(\sqrt{n/\epsilon})$ , provided that  $\epsilon \geq c/n$  for a sufficiently large constant  $c$ . Indeed, this lower bound holds also for two-sided error testers.*
2. *There exists a (sample-based) distribution-free tester for  $\Pi$  of sample complexity  $s(n, \epsilon) = O(\sqrt{n/\epsilon})$ . Furthermore, this tester has one-sided error.*

*For example,  $\Pi_n$  may be the set of monotone (non-decreasing) functions over  $[n]$ .*

It follows that there are properties for which the one-sided error sample complexity (under the uniform distribution) is at least linear in the distribution-free sample complexity. Indeed, for this assertion it would suffice to have a lower bound for one-sided error testers (under the uniform distribution) and an upper bound for two-sided error (distribution-free) testers. However, the stronger assertion of Theorem 7.5 implies that all four complexity measures are linearly related: The one-sided and two-sided error sample complexities for both the uniform distribution and the distribution-free case are all linearly related.

**Proof:** For an arbitrary ordered set  $R_n$ , we consider the class, denoted  $\Pi_n$ , of monotone non-decreasing functions from  $[n] \rightarrow R_n$ ; that is,  $f: [n] \rightarrow R_n$  is in  $\Pi_n$  if and only if  $f(i) \leq f(j)$  for every  $i < j$ .

**Claim 7.5.1** For  $R_n = [n]$ , any sample-based tester for  $\Pi_n$  (under the uniform distribution) has sample complexity  $s(n, \epsilon) = \Omega(\sqrt{n/\epsilon})$ , provided that  $\epsilon \geq c/n$  for a sufficiently large constant  $c$ .

Our guess is that Claim 7.5.1 was discovered before by several researchers, but we are unaware of a prior publication of it.

**Proof:** Assume for simplicity that  $n$  is even, and let  $n' = n/2$ . Furthermore, assume that  $\epsilon > c/n$  for a sufficiently large constant  $c$ . We consider only functions  $f$  that satisfy  $f(2i-1), f(2i) \in \{2i-1, 2i\}$  for every  $i \in [n']$ . For any  $m \in \{0, \dots, n'\}$ , consider the class of functions  $F_m$  such that  $f \in F_m$  if  $\{f(2i-1), f(2i)\} = \{2i-1, 2i\}$  for every  $i \leq m$  and  $f(2i-1) = f(2i) \in \{2i-1, 2i\}$  for every  $i > m$ . Note that  $F_0 \subset \Pi_n$ , while for every  $m$  that is at least a sufficiently large constant, with high constant probability a function selected uniformly in  $F_m$  is  $m/3n$ -far from  $\Pi_n$ . We shall show that  $\Omega(n/\sqrt{m})$  samples are required in order to distinguish a function uniformly selected in  $F_0$  from a function uniformly selected in  $F_m$ .

Let  $X = ((X_1, g(X_1)), \dots, (X_t, g(X_t)))$  denote a sequence of pairs such that the  $X_i$ 's are uniformly and independently distributed in  $[n]$  and  $g$  is uniformly distributed in  $F_0$ . Likewise,  $Y = ((Y_1, f(Y_1)), \dots, (Y_t, f(Y_t)))$  denotes a sequence where the  $Y_i$ 's are uniformly and independently distributed in  $[n]$  and  $f$  is uniformly distributed in  $F_m$ . Let  $\chi((z_1, v_1), \dots, (z_t, v_t)) = 1$  if there exists  $i \in [m]$  such that  $|\{j : z_j \in \{2i-1, 2i\}\}| > 1$  and  $\chi((z_1, v_1), \dots, (z_t, v_t)) = 0$  otherwise. Then, condition on  $\chi = 0$ , the random variables  $X$  and  $Y$  are identically distributed; that is, for every  $\gamma$ , it holds that  $\Pr[X = \gamma | \chi(X) = 0] = \Pr[Y = \gamma | \chi(Y) = 0]$ . On the other hand  $\Pr[\chi(X) = 1] = \Pr[\chi(Y) = 1]$  is upper bounded by  $\binom{t}{2} \cdot \frac{2m}{n} \cdot \frac{1}{n} < \frac{t^2 m}{n^2}$ . Thus, for  $t < n/\sqrt{10m}$ , the variation distance between  $X$  and  $Y$  is smaller than  $1/10$ . Using  $m = 3\epsilon n$ , the claim follows. ■

**Claim 7.5.2** For any ordered set  $R_n$ , there exists a (sample-based) distribution-free tester of sample complexity  $s(n, \epsilon) = O(\sqrt{n/\epsilon})$  and one-sided error for  $\Pi$ .

**Proof:** The algorithm simply accepts the function  $f$  if and only if the sample obtained is consistent with a monotone (non-decreasing) function (i.e., there exists a monotone function  $g$  such that  $f(r) = g(r)$  for every  $r$  that appears in the sample). Hence, this algorithm accepts every  $f \in \Pi_n$  with probability 1. We now consider an arbitrary function  $f$  that is  $\epsilon$ -far from  $\Pi_n$  with respect to some distribution  $\mu$  (i.e.,  $\delta_{\Pi_n}^\mu(f) > \epsilon$ ).

Following [DGL<sup>+</sup>99], we consider the (“violation”) graph  $G_f = ([n], E_f)$  such that  $(i, j) \in E_f$  if and only if  $i < j$  and  $f(i) > f(j)$ . (The edges of  $G_f$  are called violating edges.) Note that every vertex cover of this graph has weight greater than  $\epsilon$ , where the weight of any set of vertices is the sum of the probabilities assigned to its elements under  $\mu$ . (Otherwise, modifying the function on this vertex cover, in an adequate manner, yields a monotone function that is  $\epsilon$ -close to  $f$ .)<sup>38</sup> Fixing a minimum weight vertex cover  $C$  of  $G_f$ , we note that every  $S \subseteq C$  must have weight that is not larger than the weights of the set of its neighbors (i.e., of the set  $\bigcup_{v \in S} \Gamma_f(v)$ , where  $\Gamma_f(v)$  denotes the set of neighbors of  $v$  in  $G_f$ ).<sup>39</sup> We shall show that, with high probability, a sample of  $s = O(\sqrt{n/\epsilon})$  vertices (i.e., points in  $[n]$ ), drawn independently from  $\mu$ , contains a violating edge.

<sup>38</sup>If  $i$  is in the vertex cover and  $j < i$  is the largest integer that is not in the cover, then we let  $f(i) = f(j)$ . This simple modification rule benefit from the fact that the domain is totally ordered. Nevertheless, a similar modification is possible for any partially ordered domain; see [FLN<sup>+</sup>02, Lemma 1].

<sup>39</sup>Otherwise, we can replace  $S$  by  $\bigcup_{v \in S} \Gamma_f(v)$ , obtaining a vertex cover of smaller weight. (Note that edges with an endpoint in  $S$  are covered by the other endpoint, which must be in  $\bigcup_{v \in S} \Gamma_f(v)$ .)

First, we consider a sample, denoted  $S_1$ , of  $s/2$  vertices drawn independently from  $\mu$ . The expected weight of  $S_1 \cap C$  is at least  $(s/2) \cdot \epsilon/n$ , and, with very high probability, its weight is at least  $\epsilon s/4n$ . Fixing such a set  $S_1$ , we consider a second sample, denoted  $S_2$ , of  $s/2$  vertices drawn independently from  $\mu$ . Since the weight of the set of neighbors of  $S_1 \cap C \subseteq S_1$  is at least  $\epsilon s/4n$ , each of these  $s/2$  samples is a neighbor of some vertex of  $S_1$  with probability at least  $\epsilon s/4n$ . Hence, the probability that  $S_1 \times S_2$  contains no violating edge is at most  $(1 - (\epsilon s/4n))^{s/2}$ , which is very small since  $\epsilon s^2/8n \gg 1$ . The claim follows. ■

Theorem 7.5 follows. ■

We next establish Item 4 in Theorem 1.3.

**Theorem 7.6** (distribution-free sample-complexity may be exponential in the one-sided error sample complexity under the uniform distribution): *There exists a property  $\Pi = \bigcup_{\ell \in \mathbb{N}} \Pi_{2^\ell}$  such that*

1. *Any sample-based distribution-free tester for  $\Pi$  has sample complexity  $s(2^\ell, \epsilon) = \Omega(\sqrt{\ell}/\epsilon)$ . Indeed, this lower bound holds also for two-sided error testers.*
2. *There exists a sample-based one-sided error tester for  $\Pi$  (under the uniform distribution) having sample complexity  $s(2^\ell, \epsilon) = \tilde{O}(\log \ell/\epsilon)$ .*

For example,  $\Pi_{2^\ell}$  may be the set of monomials over  $\{0, 1\}^\ell$ .

It follows that there are properties for which the one-sided error sample complexity (under the uniform distribution) is logarithmic in the distribution-free sample complexity (up to a factor of  $1/\epsilon$ ). We comment that the upper bound in Item 2 is tight: Any sample-based tester for  $\Pi$  (under the uniform distribution) has sample complexity  $s(2^\ell, \epsilon) = \Omega(\log \ell/\epsilon)$ , provided that  $\epsilon > 2^{-(\ell - \ell^{0.01})}$ . (This lower bound holds also for two-sided error testers; see Appendix A.3 for details.)

**Proof:** We consider the property  $\Pi = \bigcup_{\ell \in \mathbb{N}} \Pi_{2^\ell}$  such that  $\Pi_{2^\ell}$  consists of Boolean functions  $f: \{0, 1\}^\ell \rightarrow \{0, 1\}$  that can be represented by monomials (i.e.,  $f(x_1, \dots, x_\ell)$  can be written as a conjunction of a subset of the literals  $\{x_1, \bar{x}_1, \dots, x_\ell, \bar{x}_\ell\}$ ). Glasner and Servedio [GS09] proved that the distribution-free *query* complexity of testing  $\Pi_{2^\ell}$  is  $\ell^{0.2 - o(1)}$ ; here we shall present a much simpler proof (of a higher lower bound) for the case that only samples are allowed. Turning to the upper bound, we note that Parnas *et al.* [PRS02] proved that the query complexity of this property under the uniform distribution is  $O(1/\epsilon)$ ; here we present a higher upper bound, but for sample-based tester. We stress that our sample-based tester has one-sided error.

**Claim 7.6.1** *Any sample-based distribution-free tester for  $\Pi$  has sample complexity  $s(2^\ell, \epsilon) = \Omega(\sqrt{\ell}/\epsilon)$ .*

We note that a higher lower bound of  $\Omega(\ell^{2/3})$  follows from a more complex proof of Feige (priv. comm., 2011).

**Proof:** We first establish the lower bound for  $\epsilon = \Omega(1)$ , and later extend it to a general  $\epsilon$ . To this end we define two sets,  $\mathcal{Y}$  and  $\mathcal{N}$ , of pairs  $(f, D)$ , where  $f$  is a Boolean function over  $\{0, 1\}^\ell$  and  $D$  is a distribution over  $\{0, 1\}^\ell$ . For each pair  $(f, D)$  in  $\mathcal{Y}$ , the function  $f$  is a monomial, and for each pair  $(f, D)$  in  $\mathcal{N}$  the function  $f$  is  $\Omega(1)$ -far from being a monomial with respect to the underlying distribution  $D$ . In fact, the set  $\mathcal{N}$  consists of a single pair  $(f^{\text{no}}, D^{\text{no}})$ . By Definition 7.1, any (sample based) distribution-free testing algorithm must reject  $f^{\text{no}}$  with high constant probability when given a sample generated by  $D^{\text{no}}$  and labeled by  $f^{\text{no}}$ , and must accept with high constant probability every  $f$  such that  $(f, D) \in \mathcal{Y}$  when given a sample generated by  $D$  and labeled by  $f$ . We

show that for  $s \leq \sqrt{\ell}/8$ , the variation distance between a labeled sample of size  $s$  that is generated by  $D^{\text{no}}$  and labeled by  $f^{\text{no}}$  and a labeled sample of size  $s$  that is generated and labeled according to a pair  $(f, D)$  that is selected uniformly in  $\mathcal{Y}$ , is a small constant. This implies a lower bound of  $\Omega(\sqrt{\ell})$  (for constant  $\epsilon$ ) on the sample complexity of any (sample based) distribution-free tester for monomials. Details on the construction of  $\mathcal{Y}$  and  $\mathcal{N}$  follow.

Assume for simplicity that  $\ell$  is even and consider a partition of the variables  $x_1, \dots, x_\ell$  into pairs  $(x_{2i-1}, x_{2i})$  for  $i \in [\ell/2]$ . We start by defining  $D^{\text{no}}$  and  $f^{\text{no}}$ . For each  $i \in [\ell/2]$  and  $(b, b') \in \{00, 01, 10\}$ , let  $\alpha^i(b, b')$  be the string that is all 1s except for bits in positions  $2i - 1$  and  $2i$ , which are  $b$  and  $b'$ , respectively; that is,  $\alpha^i(b, b') = (11)^{i-1}bb'(11)^{\ell/2-i}$ . The distribution  $D^{\text{no}}$  is uniform over the set of strings  $S^{\text{no}} \stackrel{\text{def}}{=} \{\alpha^i(b, b') : i \in [\ell/2], (b, b') \in \{00, 01, 10\}\}$ , and for each  $i \in [\ell/2]$ , we have  $f^{\text{no}}(\alpha^i(00)) = 0$ , and  $f^{\text{no}}(\alpha^i(01)) = f^{\text{no}}(\alpha^i(10)) = 1$ . It is irrelevant to the argument how  $f^{\text{no}}$  is defined outside of  $S^{\text{no}}$ , so we set it arbitrarily to 0.

We observe that  $f^{\text{no}}$  is  $1/3$ -far from being a monomial with respect to  $D^{\text{no}}$ . This is true since for every (disjoint) triple of strings  $\alpha^i(00), \alpha^i(01), \alpha^i(10)$  the value of the function must be modified on at least one of the three strings so that the function become consistent with some monomial.

We turn to defining the set  $\mathcal{Y}$ . For each  $\beta \in \{00, 01, 10\}^{\ell/2}$ , the set  $\mathcal{Y}$  contains the pair  $(f^\beta, D^\beta)$ , where the distribution  $D^\beta$  is uniform over the set of  $\ell/2$  strings  $S^\beta \stackrel{\text{def}}{=} \{\alpha^i(\beta_{2i-1}, \beta_{2i}) : i \in [\ell/2]\}$  (so that  $S^\beta \subset S^{\text{no}}$ ), and  $f^\beta$  is the (monotone) monomial that is the conjunction of the following variables. For each  $i \in [\ell/2]$ , if  $\beta_{2i-1} = 1$ , then  $x_{2i-1}$  is a variable of  $f^\beta$  and  $x_{2i}$  is not; if  $\beta_{2i} = 1$ , then  $x_{2i}$  is a variable of  $f^\beta$  and  $x_{2i-1}$  is not; and, if  $\beta_{2i-1} = \beta_{2i} = 0$ , then both  $x_{2i-1}$  and  $x_{2i}$  are variables of  $f^\beta$ . By this definition,  $f^\beta(\alpha^i(\beta_{2i-1}, \beta_{2i})) = 1$  when  $\beta_{2i-1}\beta_{2i} \in \{01, 10\}$  and  $f^\beta(\alpha^i(\beta_{2i-1}, \beta_{2i})) = 0$  when  $\beta_{2i-1} = \beta_{2i} = 0$ .

Note that for each  $\beta \in \{00, 01, 10\}^{\ell/2}$ , the functions  $f^{\text{no}}$  and  $f^\beta$  agree on all strings in  $S^\beta$ , and so the distance between them with respect to  $D^\beta$  is 0. This is in contrast to the fact that  $f^{\text{no}}$  and  $f^\beta$  disagree on a third of the strings in  $S^{\text{no}}$ , and so the distance between them with respect to  $D^{\text{no}}$  is  $1/3$ . Also note that for a fixed choice of  $\beta$ , the distribution over samples generated by  $D^\beta$  and labeled by  $f^\beta$  differs significantly from the distribution over samples generated by  $D^{\text{no}}$  and labeled by  $f^{\text{no}}$ , even for constant size samples. This is true simply because with high constant probability we get a sample point from  $S^{\text{no}} \setminus S^\beta$ . However, if we also select  $\beta \in \{00, 01, 10\}^{\ell/2}$  uniformly at random, then the corresponding distribution over labeled samples will be close to that generated by  $D^{\text{no}}$  and labeled by  $f^{\text{no}}$  as long the sample size is sufficiently smaller than  $\sqrt{\ell}$ . Details follow.

Consider the following two distributions over labeled samples  $((r^1, v^1), \dots, (r^s, v^s))$  of size  $s$  where  $r^j \in \{0, 1\}^\ell$  and  $v^j \in \{0, 1\}$  for each  $1 \leq j \leq s$ . The first distribution,  $L^{\text{no}}$ , generates a labeled sample simply by selecting each  $r^j$  independently according to  $D^{\text{no}}$  and setting  $v^j = f^{\text{no}}(r^j)$ . The second distribution,  $L^{\text{yes}}$ , generates a labeled sample by first selecting  $\beta \in \{00, 01, 10\}^{\ell/2}$  uniformly at random and then generating each  $r^j$  independently according to  $D^\beta$  and setting  $v^j = f^\beta(r^j)$ .

An alternative (equivalent) formulation of the generation process of samples distributed according to  $L^{\text{no}}$  is the following: For each  $1 \leq j \leq s$ , select  $i \in [\ell/2]$  uniformly at random, select  $(b, b') \in \{00, 01, 10\}$  uniformly at random, and let  $r^j = \alpha^i(b, b')$  (and  $v^j = f^{\text{no}}(r^j)$ ). Similarly, we can generate a sample distributed according to  $L^{\text{yes}}$  in the following manner. Initially, set  $\beta = ?^\ell$  (indicated that no bit in  $\beta$  is determined). For each  $1 \leq j \leq s$ , select  $i \in [\ell/2]$  uniformly at random. If  $\beta_{2i-1} = ?$ , then select  $(b, b') \in \{00, 01, 10\}$  uniformly at random and set  $\beta_{2i-1} = b$ ,  $\beta_{2i} = b'$ . Now let  $r^j = \alpha^i(\beta_{2i-1}, \beta_{2i})$  and let  $v^j = 1$  if  $(\beta_{2i-1}, \beta_{2i}) \in \{01, 10\}$  and  $v^j = 0$  otherwise. Given this alternative description, we observe that conditioned on the event that no  $i \in [\ell/2]$  is selected more than once, the two distributions on samples are exactly the same. By the (lower bound of) the ‘‘Birthday Paradox’’, if  $s \leq \sqrt{\ell}/c$  and  $c$  is a sufficiently large constant, then, with high constant probability over both generation processes, no  $i$  is selected more than once. It follows that the

statistical distance between  $L^{\text{no}}$  and  $L^{\text{yes}}$  is a small constant.

It remains to extend the argument to any  $\epsilon \leq 1/3$ . This is done by adding the string  $1^\ell$  to  $S^{\text{no}}$  as well as to each  $S^\beta$  and setting  $f^{\text{no}}(1^\ell) = 1$  (note that  $f^\beta(1^\ell) = 1$  for every  $\beta$  because  $f^\beta$  is a monotone monomial). The distributions are modified as follows:  $D^{\text{no}}(1^\ell) = D^\beta(1^\ell) = 1 - 3\epsilon$ , and  $D^{\text{no}}$  as well as each  $D^\beta$  is uniform over the remaining strings in its support.

By this modification, the distance with respect to  $D^{\text{no}}$  between  $f^{\text{no}}$  and any monomial is at least  $\epsilon$ . By combining the argument given for the case  $\epsilon = 1/3$  with the fact that the probability that each sample point differs from  $1^\ell$  (both when the sample is generated according to  $D^{\text{no}}$  and when it is generated according to  $D^\beta$  for some  $\beta$ ) is  $O(\epsilon)$ , the lower bound of  $\Omega(\sqrt{\ell}/\epsilon)$  follows. ■

**Claim 7.6.2** *There exists a sample-based one-sided error tester for  $\Pi$  (under the uniform distribution) having sample complexity  $s(2^\ell, \epsilon) = \tilde{O}(\log \ell/\epsilon)$ .*

**Proof:** As noted at the end of Section 2.2, the operation of a sample-based one-sided tester is totally determined by its sample: Being sample-based, this tester has no control over its access to the function, and having one-sided error it has no real control on its decision; that is, without loss of generality, the tester accepts if and only if the labeled sample that it has obtained is consistent with some function in  $\Pi$ . Thus, the claim amounts to proving that *if  $f: \{0, 1\}^\ell \rightarrow \{0, 1\}$  is  $\epsilon$ -far from  $\Pi$ , then, with probability at least  $2/3$ , no function in  $\Pi$  is consistent with the labeling given by  $f$  to a random sample of  $\tilde{O}(\log \ell/\epsilon)$  points.* Fixing  $\ell$  and  $\epsilon$  for the rest of the discussion, and letting  $t \stackrel{\text{def}}{=} \lceil \log_2(2/\epsilon) \rceil$ , the proof is based on two observations:

1. If  $f$  is  $\epsilon$ -far from  $\Pi_{2^\ell}$  and  $g \in \Pi_{2^\ell}$  is a monomial that contains exactly  $t$  literals, then  $\Pr_{x \in \{0,1\}^\ell} [f(x) \neq 0 = g(x)] > \epsilon/2$ , where here and throughout the entire proof all probabilities are taken uniformly over  $\{0, 1\}^\ell$ .

This holds since  $f$  is  $\epsilon$ -far from  $g$ , whereas  $\Pr[f(x) \neq 1 = g(x)] \leq \Pr[g(x) = 1] \leq 2^{-t} \leq \epsilon/2$ .

2. If  $g \in \Pi_{2^\ell}$  contains more than  $t$  literals, then there exists  $h \in \Pi_{2^\ell}$  that contains exactly  $t$  literals such that  $h^{-1}(0)$  is a subset of  $g^{-1}(0)$ ; that is, if  $h(x) = 0$  (for some  $x$ ), then  $g(x) = 0$ .

This holds by considering an arbitrary monomial  $g' \in \Pi_{2^\ell}$  that contains exactly  $t$  of the literals that appear in  $g$ .

Let  $\Pi^{1s}$  (resp.,  $\Pi^{\text{eq}}$ ) denote the set of monomials over  $\{0, 1\}^\ell$  containing less than (resp., exactly)  $t$  literals. Suppose that  $f$  is  $\epsilon$ -far from  $\Pi_{2^\ell}$ , and let  $s = O(t \log \ell/\epsilon) = \tilde{O}(\log \ell/\epsilon)$ . Then, using  $\Pr_r[f(r) = g(r)] \leq 1 - \epsilon$  for every  $g \in \Pi^{1s}$ ,

$$\begin{aligned} & \Pr_{r_1, \dots, r_s \in \{0,1\}^\ell} [\forall g \in \Pi^{1s} \exists i \in [s] \text{ s.t. } f(r_i) \neq g(r_i)] \\ &= 1 - \Pr_{r_1, \dots, r_s \in \{0,1\}^\ell} [\exists g \in \Pi^{1s} \text{ s.t. } \forall i \in [s] f(r_i) = g(r_i)] \\ &\geq 1 - \sum_{j \in [t-1]} \binom{2^\ell}{j} \cdot (1 - \epsilon)^s > \frac{5}{6}. \end{aligned} \tag{8}$$

Similarly, using Observation 1, we have

$$\begin{aligned} & \Pr_{r_1, \dots, r_s \in \{0,1\}^\ell} [\forall g \in \Pi^{\text{eq}} \exists i \in [s] \text{ s.t. } f(r_i) \neq 0 = g(r_i)] \\ &\geq 1 - \binom{2^\ell}{t} \cdot (1 - \epsilon/2)^s > \frac{5}{6}. \end{aligned} \tag{9}$$

Fixing any sequence  $(r_1, \dots, r_s)$  that satisfies the conditions in both Equation (8) and Equation (9), and using Observation 2, we conclude that for every  $g \in \Pi_{2^\ell}$  there exists  $i \in [s]$  such that  $f(r_i) \neq g(r_i)$ , where for  $g \in \Pi_{2^\ell} \setminus \Pi^{1s}$  we actually have  $f(r_i) \neq 0 = g(r_i)$ . The claim follows. ■

Theorem 7.6 follows from Claims 7.6.1, and 7.6.2. ■

**A detour: A generalization of the argument underlying the proof of Claim 7.6.2.** Let us consider an arbitrary property  $\Pi$  of functions over  $[n]$ , where in Claim 7.6.2  $[n] \equiv \{0, 1\}^\ell$ . For every two functions  $f$  and  $g$  over  $[n]$ , let  $\Delta(f, g) \stackrel{\text{def}}{=} \{x \in [n] : f(x) \neq g(x)\}$ . As in the proof of Claim 7.6.2, for a random sample of  $s$  points that are labeled by a function  $f$  that is  $\epsilon$ -far from  $\Pi$ , we wish to lower bound the probability that  $f$  disagrees with each  $g \in \Pi$  on some point in that sample. The issue is improving over the obvious lower bound of  $1 - |\Pi| \cdot (1 - \epsilon)^s$ . The argument establishing Claim 7.6.2 made implicit use of a suitable notion of a cover of such  $f$ 's (with respect to  $\Pi$ ). For  $B: \mathbb{N} \times [0, 1] \rightarrow \mathbb{N}$  and  $\rho: [0, 1] \rightarrow [0, 1]$ , we say that a collection  $\mathcal{S}_\Pi(f)$  of subsets (of  $[n]$ ) is a  $(B, \rho)$ -cover of  $f$  (w.r.t.  $\Pi$ ) if the following two conditions hold:

1. For every  $g \in \Pi$  there exists  $S \in \mathcal{S}_\Pi(f)$  such that  $S \subseteq \Delta(f, g)$ .
2. If  $f$  is  $\epsilon$ -far from  $\Pi$ , then  $|\mathcal{S}_\Pi(f)| \leq B(n, \epsilon)$  and for every  $S \in \mathcal{S}_\Pi(f)$  it holds that  $|S| \geq \rho(\epsilon) \cdot n$ .

(In the proof of Claim 7.6.2, we implicitly referred to the collection  $\mathcal{S}_\Pi(f)$  that contained all  $\Delta(f, g)$  for  $g \in \Pi^{1s}$  and all  $\{x : f(x) \neq 0 = g(x)\} \subseteq \Delta(f, g)$  for  $g \in \Pi^{\text{eq}}$ .) We stress that the collection  $\mathcal{S}_\Pi(f)$  is *not* constructed by the tester, it is merely used in its analysis, which asserts that a sample of  $s(n, \epsilon) = O(\rho(\epsilon)^{-1} \log B(n, \epsilon))$  random points suffices provided that  $\mathcal{S}_\Pi(f)$  is a  $(B, \rho)$ -cover of  $\Pi$  w.r.t.  $f$ . To see that this is the case, consider any function  $f$  (over the  $n$ -element domain) that is  $\epsilon$ -far from  $\Pi$ . Then, by Condition 2, with probability at least  $1 - (1 - \rho(\epsilon))^{s(n, \epsilon)} \geq 2/3$ , a sample of  $s(n, \epsilon)$  random points hits each subset in the collection  $\mathcal{S}_\Pi(f)$ . In this case, by Condition 1, for every  $g \in \Pi$  this sample hits  $\Delta(f, g)$ .

**Back to the main thread.** Turning back to the gap between distribution-free sample-complexity and one-sided error sample complexity (under the uniform distribution), we note that an ever larger gap exists. Specifically, distribution-free sample-based testing may be extremely hard (i.e., requires a linear number of queries), while for the same property one-sided error testing (under the uniform distribution) may be extremely easy (or even trivial). The next theorem established Item 5 in Theorem 1.3.

**Theorem 7.7** (distribution-free sample-based testing can be extremely harder than one-sided error sample-based testing under the uniform distribution): *There exists a property  $\Pi$  such that*

1. *Any distribution-free tester for  $\Pi$  has query complexity  $q(n, \Omega(1)) = \Omega(n)$ . Indeed, this lower bound holds also for two-sided error testers that make queries in addition to obtaining samples from the distribution.*
2. *There exists a sample-based one-sided error tester for  $\Pi$  (under the uniform distribution) having sample complexity  $s(n, \epsilon) = \tilde{O}(1/\epsilon)$ .*

*In fact, we may select  $\Pi$  such that any function is  $1/n$ -close to  $\Pi_n$ .*

Note that in the latter case, testing  $\Pi_n$  has sample complexity  $\tilde{O}(1/\epsilon)$ , where the latter expression smoothens the actual complexity that is  $s(n, \epsilon) = 0$  if  $\epsilon > 1/n$  and  $s(n, \epsilon) = O(n \log n)$  otherwise.

**Proof:** Our construction of  $\Pi_n$  is a disjoint union of an extremely easy-to-test property, denoted  $\Pi_{n-1}^E$ , and an extremely hard-to-test property, denoted  $\Pi_{n-1}^H$ . Specifically,  $\Pi_{n-1}^E$  may be trivial (i.e., contains all functions from  $[n-1]$  to  $\{0,1\}$ ), whereas  $\Pi_{n-1}^H$  may be a property that is extremely hard to test (i.e., requires a linear number of queries even for two-sided error testing under the uniform distribution, cf., e.g., [GGR98, Section 4.1]).<sup>40</sup> The disjoint union of these two properties is captured by  $\Pi_n$  such that  $f: \{0, 1, \dots, n-1\}$  is in  $\Pi_n$  if either  $f(0) = 0$  and  $f' \in \Pi_{n-1}^E$  or  $f(0) = 1$  and  $f' \in \Pi_{n-1}^H$ , where  $f'$  is the restriction of  $f$  to  $[n-1]$  (i.e.,  $f'(i) = f(i)$  for  $i \in [n-1]$ ).

Indeed, testing  $\Pi_n$  is almost trivial *under the uniform distribution*, since in this case any  $f: [n] \rightarrow \{0,1\}$  is  $1/n$ -close to  $\Pi_n$  (by merely resetting  $f(0) = 0$ ). Thus, for  $\epsilon > 1/n$ , the tester may just accept (without looking at anything else), whereas for  $\epsilon \leq 1/n$  it may just reconstruct the function based on  $O(n \log n)$  labeled samples (and rule accordingly). Hence:

**Claim 7.7.1** *The one-sided error sample complexity of  $\Pi$  (under the uniform distribution) is  $s(n, \epsilon) = \tilde{O}(1/\epsilon)$ .*

The forgoing phenomenon (of each function being  $1/n$ -close to  $\Pi_n$ ) *does not necessarily hold with respect to an arbitrary distribution*: Consider, for example, a distribution that assign probability mass of  $1/2$  to  $i = 0$ . In this case, if  $f(0) = 1$  and  $\epsilon < 1/2$ , we cannot afford to reset  $f(0) = 0$  and must test the residual  $f'$  with respect to  $\Pi_{n-1}^H$ . In this case (i.e.,  $\epsilon < 1/2$ ), testing  $\Pi_n$  with respect to the distribution  $D_n$  that assigns probability  $1/2$  to  $i = 0$  and is uniform on the rest of  $[n]$ , amounts to testing  $\Pi_{n-1}^H$ . Formally:

**Claim 7.7.2** *Let  $s(n, \epsilon)$  denote the sample complexity of one-sided error testing  $\Pi_n$  under the distribution  $D_n$ , and  $s'(n-1, \epsilon)$  denote the sample complexity of one-sided error testing  $\Pi_{n-1}^H$  under the uniform distribution. Then,  $s'(n-1, \epsilon) \leq s(n, \epsilon/2)$ . Ditto with respect to the query complexity of testers that are allowed queries.*<sup>41</sup>

Indeed, the sample-based tester for  $\Pi_{n-1}^H$  invokes the sample-based tester for  $\Pi_n$ , while providing it with an adequately distributed sample. Specifically, for each  $j \in [t]$ , with probability  $1/2$  it places  $(0,1)$  as the  $j^{\text{th}}$  labeled sample of the invoked tester, and otherwise it just uses the  $j^{\text{th}}$  labeled sample that it got. Hence, given a uniformly distributed sample labeled by  $f'$ , the new tester creates a sample labeled by the corresponding  $f$  (i.e.,  $f(0) = 1$  and  $f(i) = f'(i)$  for every  $i \in [n-1]$ ), where this sample is distributed according to  $D_n$ . A similar transformation holds for testers that make queries. Using the hardness of  $\Pi_{n-1}^H$ , the theorem follows. ■

## Acknowledgements

We wish to thank several readers for their useful comments, including, in particular, the reviewers of this (journal) version of the paper.

## References

- [AK02] N. Alon and M. Krivelevich. Testing  $k$ -colorability. *SIAM Journal on Discrete Math*, 15(2):211–227, 2002.

<sup>40</sup>More generally, we might have use any easy to test  $\Pi_{n-1}^E$ , provided  $\Pi_{n-1}^H \subset \Pi_{n-1}^E$ . In this case, the main claims would have hold, although it would not necessarily be the case that any function is  $1/n$ -close to  $\Pi_n$ .

<sup>41</sup>Note that since we deal with a fixed distribution in each of the two problems, queries subsume samples.

- [AKNS01] N. Alon, M. Krivelevich, I. Newman, and M. Szegedy. Regular languages are testable with a constant number of queries. *SIAM Journal on Computing*, pages 1842–1862, 2001.
- [BBBY12] M. Balcan, E. Blais, A. Blum, and L. Yang. Active property testing. In *Proceedings of the Fifty-Third Annual Symposium on Foundations of Computer Science (FOCS)*, pages 21–30, 2012.
- [BFR<sup>+</sup>00] T. Batu, L. Fortnow, R. Rubinfeld, W. Smith, and P. White. Testing that distributions are close. In *Proceedings of the Forty-First Annual Symposium on Foundations of Computer Science (FOCS)*, pages 259–269, 2000.
- [BLR93] M. Blum, M. Luby, and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of the ACM*, 47:549–595, 1993.
- [Dev83] L. Devroye. The equivalence of weak, strong and complete convergence in L1 for kernel density estimates. *Annals of Statistics*, 11(3):896–904, 1983.
- [DGL<sup>+</sup>99] Y. Dodis, O. Goldreich, E. Lehman, S. Raskhodnikova, D. Ron, and A. Samorodnitsky. Improved testing algorithms for monotonicity. In *Proceedings of the Third International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*, pages 97–108, 1999.
- [FGL13] E. Fischer, Y. Goldhirsh, , and O. Lachish. Some properties are not even partially testable. Technical Report TR13-082, Electronic Colloquium on Computational Complexity (ECCC), 2013.
- [FGL14] E. Fischer, Y. Goldhirsh, , and O. Lachish. Partial tests, universal tests and decomposability. In *Proceedings of the Fifth Innovations in Computer Science conference (ITCS)*, pages 474–483, 2014.
- [FLN<sup>+</sup>02] E. Fischer, E. Lehman, I. Newman, S. Raskhodnikova, R. Rubinfeld, and A. Samorodnitsky. Monotonicity testing over general poset domains. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on the Theory of Computing (STOC)*, pages 474–483, 2002.
- [GGL<sup>+</sup>00] O. Goldreich, S. Goldwasser, E. Lehman, D. Ron, and A. Samorodnitsky. Testing monotonicity. *Combinatorica*, 20(3):301–337, 2000.
- [GGN10] O. Goldreich, S. Goldwasser, and A. Nussboim. On the implementation of huge random objects. *SIAM Journal on Computing*, 39(7):2761–2822, 2010.
- [GGR98] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45(4):653–750, 1998.
- [GK11] O. Goldreich and T. Kaufman. Proximity oblivious testing and the role of invariances. In *Proceedings of the Fifteenth International Workshop on Randomization and Computation (RANDOM)*, pages 579–592, 2011.
- [Gol] O. Goldreich. Introduction to testing graph properties. In [Gol10].
- [Gol10] O. Goldreich, editor. *Property Testing: Current Research and Surveys*. Springer, 2010. LNCS 6390.

- [GR02] O. Goldreich and D. Ron. Property testing in bounded degree graphs. *Algorithmica*, 32(2):302–343, 2002.
- [GR11] O. Goldreich and D. Ron. On proximity oblivious testing. *SIAM Journal on Computing*, 40(2):534–566, 2011.
- [GR13] T. Gur and R. Rothblum. Non-interactive proofs of proximity. Technical Report TR13-078, Electronic Colloquium on Computational Complexity (ECCC), 2013.
- [GS09] D. Glasner and R. A. Servedio. Distribution-free testing lower bound for basic boolean functions. *Theory of Computing*, 5(1):191–216, 2009.
- [GS12] O. Goldreich and I. Shinkar. Two-sided error proximity oblivious testing. In *Proceedings of the Sixteenth International Workshop on Randomization and Computation (RANDOM)*, pages 565–578, 2012.
- [GT03] O. Goldreich and L. Trevisan. Three theorems regarding testing graph properties. *Random Structures and Algorithms*, 23(1):23–57, 2003.
- [Knu] D.E. Knuth. *The Art of Computer Programming Vol. 2 (Seminumerical Algorithms)*. Addison-Wesley Publishing Company, Inc. 1969 (first edition) and 1981 (second edition).
- [KR00] M. Kearns and D. Ron. Testing problems with sub-learning sample complexity. *Journal of Computer and System Sciences*, 61(3):428–456, 2000.
- [KS08] T. Kaufman and M. Sudan. Algebraic property testing: The role of invariances. In *Proceedings of the Fourtieth Annual ACM Symposium on the Theory of Computing (STOC)*, pages 403–412, 2008.
- [KT00] J. Katz and L. Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *Proceedings of the Thirty-Second Annual ACM Symposium on the Theory of Computing (STOC)*, pages 80–86, 2000.
- [KV94] M. Kearns and U. Vazirani. *An introduction to Computational Learning Theory*. MIT Press, 1994.
- [PRS02] M. Parnas, D. Ron, and A. Samorodnitsky. Testing basic boolean formulae. *SIAM Journal on Discrete Math*, 16(1):20–46, 2002.
- [Ron08] D. Ron. Property testing: A learning theory perspective. *Foundations and Trends in Machine Learning*, 1(3):307–402, 2008.
- [Ron10] D. Ron. Algorithmic and analysis techniques in property testing. *Foundations and Trends in Theoretical Computer Science*, 5:73–205, 2010.
- [RS96] R. Rubinfeld and M. Sudan. Robust characterization of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.
- [RS06] S. Raskhodnikova and A. Smith. A note on adaptivity in testing properties of bounded degree graphs. Technical Report TR06-089, Electronic Colloquium on Computational Complexity (ECCC), 2006.
- [Sud] M. Sudan. Invariance in property testing. In [Gol10], pages 211–227.

- [Val84] L. G. Valiant. A theory of the learnable. *CACM*, 27(11):1134–1142, November 1984.
- [VV11] G. Valiant and P. 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 the Theory of Computing (STOC)*, pages 685–694, 2011. See *ECCC* TR10-180 for the algorithm, and TR10-179 for the lower bound.
- [VV15] G. Valiant and P. Valiant. Instance optimal learning. *CoRR*, abs/1504.05321, 2015.

## Appendix: In Passing

### A.1 A relaxed notion of POT

The following definition is a hybrid of the standard definition of a tester and a POT: Specifically, the tester gets the proximity parameter  $\epsilon$  and its behavior and performance may depend on  $\epsilon$ , but its query complexity is fixed.

**Definition A.1** (relaxed POT): *Let  $\Pi$  and  $\varrho$  be as in Definition 2.2. A relaxed POT with detection probability  $\varrho$  for  $\Pi$  is a probabilistic oracle machine  $T$  that makes a constant number of queries and satisfies the following two conditions:*

1. *For every  $n \in \mathbb{N}$ ,  $\epsilon > 0$ , and  $f \in \Pi_n$ , it holds that  $\Pr[T^f(n, \epsilon) = 1] \geq c$ .*
2. *For every  $n \in \mathbb{N}$ ,  $\epsilon > 0$ , and  $f: D_n \rightarrow R_n$  that is  $\epsilon$ -far from  $\Pi_n$ , it holds that  $\Pr[T^f(n, \epsilon) = 1] \leq c - \varrho(\epsilon)$ .*

*Again, the constant  $c$  is called the threshold probability, and the tester is said to have one-sided error if  $c = 1$ .*

Note that detection probability of a relaxed POT is not required to increase with the distance of the tested object to the property. (Instead, the detection probability is determined by the proximity parameter  $\epsilon$ , which is given to the tester, and the detection guarantee applies only to functions that are  $\epsilon$ -far from  $\Pi$ .) Nevertheless, if  $c = 1$ , then such a tester can be used to derive a POT (in which the detection probability does increase with the distance of the tested object to the property).

**Proposition A.2** *If  $\Pi$  has a one-sided error relaxed POT with detection probability  $\varrho$ , then  $\Pi$  has a one-sided error POT with detection probability  $\varrho'$  such that  $\varrho'(\delta) = O(\varrho(\delta/2)/\log^2(2/\delta))$ . Furthermore, the resulting POT preserves the fairness of the relaxed POT, where fairness is as defined in the beginning of Section 3.*

**Proof:** Let  $T$  be a (one-sided error) relaxed POT with detection probability  $\varrho$  for  $\Pi$ , and consider the following oracle machine: On input  $n$  and access to the oracle  $f$ , the machine selects  $i \in [\lceil \log_2 n \rceil]$  with probability  $1/(i+1)^2$ , invokes  $T^f(n, 2^{-i})$ , and outputs whatever  $T$  outputs.<sup>42</sup> Note that if  $f \in \Pi_n$ , then the new tester always outputs 1 (regardless of its choice of  $i$ ). On the other hand, if  $f: D_n \rightarrow R_n$  is not in  $\Pi_n$ , then  $\delta_\Pi(f) \geq 1/n$  and  $i = \lceil \log_2(1/\delta_\Pi(f)) \rceil$  is selected with probability at least  $(1 + \lceil \log_2(1/\delta_\Pi(f)) \rceil)^{-2}$ . Hence,  $T$  outputs 1 with probability at most  $1 - \frac{\varrho(2^{-i})}{(1 + \lceil \log_2(1/\delta_\Pi(f)) \rceil)^2} \leq 1 - \frac{\varrho(\delta_\Pi(f)/2)}{(1 + \lceil \log_2(1/\delta_\Pi(f)) \rceil)^2}$ , where the inequality is due to the monotonicity of  $\varrho$ . ■

**Remark A.3** *The proof of Proposition A.2 does not extend to the case that the threshold probability is smaller than 1 (i.e.,  $c < 1$ ). The problem is that in such a case, it is not guaranteed that functions outside of  $\Pi$  are rejected with probability at least  $1 - c$  when the relaxed POT is given a proximity parameter that is smaller than their distance from  $\Pi$ . The proof does extend to two-sided error POTs that satisfy the latter condition.*

<sup>42</sup>With probability  $1 - \sum_{i \in [\log_2 n]} (i+1)^{-2} > 0$ , the machine just outputs 1 (or alternatively it selects  $i \in [\log_2 n]$  arbitrarily and behaves accordingly).

## A.2 Relating various notions of sampling

In this section, we relate various notions of sampling. In all cases we are given a number of sample points from a domain  $D_n$  as well as the size of this domain  $n = |D_n|$ . Definition 2.3 refers to the natural case in which we are given a fixed number of sample points, where each sample point is uniformly and independently distributed in  $D_n$ . Thus, the given sample may contain repetitions, and in fact it will contain repetitions with a fair (or very high) probability if its size exceeds  $\sqrt{|D_n|}$ .

It is often more convenient to consider a sample of a fixed size that is drawn without repetitions; that is, a set of a fixed size, denoted  $m$ , that is selected uniformly among all  $m$ -subsets of  $D_n$ . Likewise, we sometime consider a sample of  $D_n$  generated by picking each element with some probability  $p$  (e.g.,  $p = m/n$ ) independently of all other choices. In this section we relate these three notions, showing that each one can emulate the others up-to small deviations, which are typically insignificant. We first relate sampling with and without repetitions.

**Proposition A.4** (sampling with repetitions implies sampling without repetitions): *Suppose that  $m \leq n$  and that we are given  $m$  uniformly and independently distributed elements of  $D_n$ , denoted  $r_1, \dots, r_m$ . Then, with probability at least  $1 - \exp(-m)$ , this sample contains at least  $m/3$  distinct elements. In this case, if we let  $t \leq m$  be the smallest integer such that  $|\{i \in [t] : r_i\}| = m/3$ , and output the  $m/3$ -set  $\{i \in [t] : r_i\}$ , then we get a uniformly selected sample of size  $m/3$  without repetitions.*

**Proof:** The probability of the bad event is upper bounded by

$$\binom{n}{m/3} \cdot \left(\frac{m/3}{n}\right)^m < (9n/m)^{m/3} \cdot (m/3n)^m \leq (1/3)^{m/3}$$

where the first inequality is due to  $t! > (t/3)^t$ . Given that the bad event did not occur, and denoting the  $i^{\text{th}}$  element placed in the output set by  $s_i$ , observe that  $s_i$  is distributed uniformly in  $D_n \setminus \{s_1, \dots, s_{i-1}\}$ . ■

**Proposition A.5** (sampling without repetitions implies sampling with repetitions): *Suppose that  $m \leq n$  and that we are given a set of  $m$  elements selected uniformly among all  $m$ -subsets of  $D_n$ . Let  $(e_1, \dots, e_m)$  be a random ordering of the elements of this set. Consider a selection of  $(r_1, \dots, r_m)$  such that, for every  $i \in [m]$ , with probability  $(n - i + 1)/n$  we select  $r_i = e_i$  and otherwise we select  $r_i$  uniformly in  $\{e_1, \dots, e_{i-1}\}$ . Then,  $(r_1, \dots, r_m)$  is a sequence of  $m$  uniformly and independently distributed elements of  $D_n$ .*

**Proof:** We shall show that for every  $i \in [m]$  and for every  $v_1, \dots, v_i \in D_n$ , it holds that

$$\Pr[r_i = v_i \mid (r_1, \dots, r_{i-1}) = (v_1, \dots, v_{i-1})] = 1/n.$$

In fact, for each sequence of  $i - 1$  distinct  $u_1, \dots, u_{i-1} \in D_n$ , we consider the probability

$$\Pr[r_i = v_i \mid (r_1, \dots, r_{i-1}) = (v_1, \dots, v_{i-1}) \ \& \ (e_1, \dots, e_{i-1}) = (u_1, \dots, u_{i-1})].$$

Now, if  $v_i \in \{u_1, \dots, u_{i-1}\}$ , then  $r_i = v_i$  with probability  $(1 - \frac{n-i+1}{n}) \cdot \frac{1}{i-1} = \frac{1}{n}$ . On the other hand, if  $v_i \notin \{u_1, \dots, u_{i-1}\}$ , then  $r_i = v_i$  with probability  $\frac{n-i+1}{n} \cdot \frac{1}{n-(i-1)} = \frac{1}{n}$  (where  $\frac{1}{n-(i-1)}$  represents the probability that  $e_i = v_i$  when  $e_i$  is selected uniformly in  $D_n \setminus \{u_1, \dots, u_{i-1}\}$ ). ■

**Two versions of sampling without repetitions.** In the foregoing, “sampling without repetitions” referred to selecting a random subset of a fixed size within  $D_n$ . An alternative notion of sampling without repetitions arises when each element is selected with some fixed probability, independently of all other choices.

**Construction A.6** (individual selection implies sampling a set of fixed size): *Suppose that each point is selected to be included in the sample with probability  $p \in (0, 1]$  independently of all other choices. Then, with probability at least  $1 - \exp(-pn)$ , at least  $m \stackrel{\text{def}}{=} pn/2$  points are selected. In this case, we output a random  $m$ -subset of the set of selected points.*

**Construction A.7** (sampling a set of fixed size implies individual selection): *Suppose that  $m \leq n$  and that we are given a set  $S$  of  $m$  elements selected uniformly among all  $m$ -subsets of  $D_n$ . Setting  $p = m/2n$ , let  $B_k(n, p)$  denote the probability of obtaining  $k$  successes when making  $n$  independent experiments with success probability  $p$  each. Consider selecting  $i$  with probability  $B_i(n, p)$ , and outputting a random  $i$ -subset of  $S$ , if  $i \leq |S|$ , and outputting nothing otherwise. Then, the output distribution is  $\exp(-m)$ -close to the distribution of a sample that is selected by including each point in the sample with probability  $p$  independently of all other choices.*

Sampling according to the binomial distribution  $B(n, p)$  (i.e., selecting  $i$  with probability  $B_i(n, p)$ ) can be approximated up to  $\epsilon$  in  $\text{poly}(\log(n/\epsilon))$ -time; see, e.g., [Knu, Section 3.4.1] (and [GGN10, Appendix A.1]).

### A.3 On the sample complexity of testing monomials

The following result is implicit in [BBBY12].

**Proposition A.8** *Let  $\Pi = \bigcup_{\ell \in \mathbb{N}} \Pi_{2^\ell}$ , where  $\Pi_{2^\ell}$  is the set of all  $\ell$ -variable Boolean functions that are monomials. Then, any sample-based tester for  $\Pi$  (under the uniform distribution) has sample complexity  $s(2^\ell, \epsilon) = \Omega(\log \ell/\epsilon)$ , provided that  $\epsilon > 2^{-(\ell - \ell^{0.01})}$ . This lower bound holds also for two-sided error testers.*

**Proof:** Let us consider first the case  $\epsilon = 1/3$ . In this case we consider the task of distinguishing a random Boolean function  $f: \{0, 1\}^\ell \rightarrow \{0, 1\}$  from a uniformly distributed dictatorship function (i.e.,  $d_i(x) = x_i$  for  $i \in [\ell]$ ) based on  $s = 0.5 \log_2 \ell$  labeled samples, denoted  $r_1, \dots, r_s$ . (This is exactly the task analyzed in [BBBY12, Theorem VI.8].)<sup>43</sup> The labels assigned by a random function are uniformly and independently distributed in  $\{0, 1\}$ , since the probability of a repetition (i.e.,  $|\{r_1, \dots, r_s\}| < s$ ) in such a small sample is negligible. We now analyze the distribution of  $s$  samples that are labeled by  $d_i$  for a uniformly distributed  $i \in [\ell]$ .

We consider an iterative process of assigning these labels, while initializing  $I = [\ell]$ . Specifically, for  $j = 1, 2, \dots, s$ , the label of the  $j^{\text{th}}$  sample is determined by selecting  $i_j$  uniformly in  $I$ , labeling the  $j^{\text{th}}$  sample according to  $d_{i_j}$ , and resetting  $I$  to equal  $\{i \in I : d_i(r_j) = d_{i_j}(r_j)\}$ . Denoting the set  $I$  after the  $j^{\text{th}}$  iteration by  $I_j$ , note that  $I_j = \{i \in I_{j-1} : r_{j,i} = r_{j,i_j}\}$ , where  $r_{j,i}$  is the  $i^{\text{th}}$  bit of  $r_j$ . Hence, for every  $j \in [s]$ , with probability at least  $1 - \exp(-\Omega(|I_{j-1}|^{1/3}))$ , it holds that  $|I_j| = (0.5 \pm |I_{j-1}|^{-1/3}) \cdot |I_{j-1}|$ . It follows that the deviation of these labels from the uniform distribution is upper-bounded by  $\exp(-\omega(\ell/2^s)^{1/3}) = \exp(-\omega(\ell^{1/6}))$ , and  $s(2^\ell, 1/3) \geq 0.5 \log_2 \ell$  follows, since (w.v.h.p.) a random function is  $1/3$ -far from  $\Pi$ .

<sup>43</sup>In fact, the furthermore clause of [BBBY12, Theorem VI.8] refers explicitly to distinguishing dictatorship functions from random functions. Still, for the sake of self-containment, we reproduce the argument here.

For the general case (of  $\epsilon < 1/6$ ), we designate  $t \stackrel{\text{def}}{=} \log_2(3/\epsilon)$  of the variables and apply the foregoing construction to the remaining  $\ell - t$  variables. For each function  $f': \{0, 1\}^{\ell-t} \rightarrow \{0, 1\}$  (as considered above), we consider the function  $f(x_1, \dots, x_\ell) = \bigwedge_{i \in [t]} x_i \wedge f'(x_{t+1}, \dots, x_\ell)$ . Note that if  $f'$  is a dictator function then  $f$  is a monomial, whereas if  $f'$  is  $1/3$ -far from  $\Pi$  then  $f$  is  $\epsilon$ -far from  $\Pi$ . We thus obtain a lower bound of  $s(2^\ell, \epsilon) = \Omega(s(2^{\ell-t}, 1/3)/\epsilon)$ , and the claim follows. ■