

Three Theorems Regarding Testing Graph Properties*

Oded Goldreich,^{1,†} Luca Trevisan²

¹Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel; e-mail: oded@wisdom.weizmann.ac.il

²EECS—Computer Science Division, University of California, Berkeley, California 94720; e-mail: luca@eecs.berkeley.edu

Received 9 December 2001; accepted 18 September 2002

DOI 10.1002/rsa.10078

ABSTRACT: Property testing is a relaxation of decision problems in which it is required to distinguish YES-instances (i.e., objects having a predetermined property) from instances that are far from any YES-instance. We presents three theorems regarding testing graph properties in the adjacency matrix representation. More specifically, these theorems relate to the project of characterizing graph properties according to the complexity of testing them (in the adjacency matrix representation). The first theorem is that there exist *monotone* graph properties in \mathcal{NP} for which testing is very hard (i.e., requires to examine a constant fraction of the entries in the matrix). The second theorem is that every graph property that can be tested making a number of queries that is independent of the size of the graph can be so tested by uniformly selecting a set of vertices and accepting iff the induced subgraph has some fixed graph property (which is not necessarily the same as the one being tested). The third theorem refers to the framework of *graph partition problems*, and is a characterization of the subclass of properties that can be tested using a one-sided error tester making a number of queries that is independent of the size of the graph. © 2003 Wiley Periodicals, Inc. *Random Struct. Alg.*, 23: 23–57, 2003

Keywords: property testing; graph properties; monotone properties; NP

Correspondence to: O. Goldreich

*An extended abstract has appeared in the proceedings of the 42nd FOCS, Las Vegas, NV, October 2001.

†Supported by the MINERVA Foundation, Germany.

© 2003 Wiley Periodicals, Inc.

1. INTRODUCTION

Property testing (cf. [12, 8]) is a natural notion of approximation for decision problems: For a predetermined property (or decision), the task is to distinguish whether a given instance has this property (i.e., is a YES-instance) or is “far” from any instance having the property.

This work is concerned with testing *graph properties* in the *adjacency matrix representation* (cf. [8]). A tester for a predetermined graph property Π is a (randomized) algorithm that is given a size parameter N and a distance parameter ε as well as oracle access to the adjacency matrix of an N -vertex graph G ; that is, query $(u, v) \in [N] \times [N]$ is answered by a bit indicating whether or not the edge (u, v) is present in the graph. The algorithm is required to accept (with probability at least $\frac{2}{3}$) any graph having property Π , and reject (with probability at least $\frac{2}{3}$) any graph that is ε -far from having property Π , where distance between (N -vertex) graphs is defined as the fraction of edges (over $\binom{N}{2}$) on which the graphs differ.

Our focus is on the *query complexity* of testing some graph properties. We consider two extreme cases. In one case, the query complexity of testing depends only on the distance parameter ε (and is independent of the size of the graph N). In this case, we say that **testing is very easy**. In the other extreme case, for some constant $\varepsilon > 0$, any tester must make $\Omega(N^2)$ queries (and is thus not significantly better than a trivial tester that inspect the entire graph). In this case, we say that **testing is very hard**.

1.1. Our Main Results

Our first main result (cf. Theorem 1) is that there exist *monotone* graph properties in \mathcal{NP} for which testing is very hard (i.e., requires $\Omega(N^2)$ queries). This improves over an analogous result of Goldreich, Goldwasser, and Ron [8, Property 10.2.3.2] that established the same lower bound for non-monotone graph properties (in \mathcal{NP}). In fact, our result is obtained by a simple extension of their technique. This resolves a natural open problem (raised by several researchers, and most recently by Y. Dodis).

Our second main result (cf. Theorem 2) refers to graph properties that can be tested very easily (i.e., using a number of queries that only depends on the distance parameter ε). We show that such graph properties can be so tested by uniformly selecting a set of vertices of size depending only on ε and accepting if and only if the induced subgraph has some graph property (which is not necessarily the same as the one being tested). This improves over a previous result of Alon et al. [2], who only assert that a tester may just inspect a random induced subgraph but do not assert that the decision may depend only on a property of that subgraph (rather than also on the tester’s coins). Our result extend to any query complexity so that if the original tester had query complexity q , then the new tester has query complexity that is polynomial in q . Furthermore, the transformation preserves one-sided error. It follows that the query complexity of testing graph properties is polynomially related to the query complexity of performing such testing via nonadaptive testers (and while preserving one-sided error). This improves over the naive transformation of adaptive testers to non-adaptive ones (which incurs an exponential blow-up in the query complexity).

Our third main result (cf. Theorem 3) refers to the framework of *graph partition problems* introduced by Goldreich, Goldwasser, and Ron [8]. This natural framework refers to graph properties that require the existence of a partition of the graph’s vertices

such that certain bounds on the sizes of parts and the number of edges between them must hold. In particular, the framework include several natural graph properties (e.g., k -Colorability as well as threshold versions of Max-Clique, Max-Cut, and Bisection). Goldreich, Goldwasser, and Ron showed that every problem Π in this framework can be tested very easily (i.e., by making $\text{poly}(1/\epsilon)$ queries). Within this framework, we characterize the subclass of properties that can be tested very easily using a *one-sided error* tester (i.e., the tester must accept any graph having the property with probability 1). Details follow.

A **graph partition testing problem** is parameterized by a sequence of corresponding pairs of lower and upper bounds. Specifically, for some (implicit) parameter k , the sequence contains k pairs of vertex-sets densities and $k + \binom{k}{2}$ pairs of “edge-densities.” A graph has the specified (by the sequence) property if there exists a k -partition of its vertices such that the number of vertices in each component of the partition as well as the number of edges within each component and between each pair of components falls between the corresponding lower and upper bounds (in the sequence of parameters). For example, k -colorability can be expressed by setting the edge-density upper-bounds corresponding to edges within each part to equal zero. Goldreich, Goldwasser, and Ron [8] showed that every graph partition property (i.e., problem in the above framework) can be tested by making $\text{poly}(1/\epsilon)$ queries, but in general their tester has two-sided error probability. They also gave one-sided error testers for k -colorability (which operate by checking whether a random induced $\text{poly}(1/\epsilon)$ -vertex subgraph is k -colorable). We show that the class of graph partition properties that admit a one-sided error tester of query complexity that is independent of N consists of two subclasses:

1. *The main subclass:* Each property in the subclass corresponds to a k -vertex graph H . An N -vertex graph has the property if its vertices can be k -partitioned such that there are no edges among vertices residing in the same part and so that there are edges between vertices of the i th part and j th part only if (i, j) is an edge of H . (For example, k -colorability is expressed by letting H be the k -vertex clique.)
2. *The noninteresting subclass* consists of two graph properties: the *clique property* and the *trivial property*. The only N -vertex graph satisfying the *clique property* is the N -vertex clique, whereas all (but finitely many) graphs satisfy the *trivial property*.

We note that each property in the above class can be tested with one-sided error by uniformly selecting a set of $\text{poly}(1/\epsilon)$ vertices and accepting if and only if the induced subgraph satisfies the very same graph property being tested.

1.2. Perspective

For a wider perspective on property testing, see [7, 11]. Our results are related to the project of characterizing graph properties according to the complexity of testing them. This is a natural research project, alas a seemingly very difficult one (cf. [8, 2]). Our results carry good and bad news for this project.

Theorem 1 refutes the conjecture that all *monotone* graph properties can be tested very easily, a conjecture which could have been justified by the fact that (monotone) NP-hard problems such as k -Colorability and ρ -clique can be tested very easily (cf. [8, Sections 6.2 and 7]). Theorem 1 can be viewed as bad news for the “characterization project” (because

it means that yet another natural class of properties has both easily-testable and hard-to-test properties).

Theorem 2 provides a tool for the study of graph properties that can be tested very easily. It asserts that when conducting such a study, one may focus on canonical testers that operate in a relatively simple way. The usefulness of Theorem 2 is demonstrated in our proof of Theorem 3, a proof that repeatedly refers to the fact that the canonical tester accepts if and only if a random induced subgraph has a certain graph property.

Theorem 3 is of the type of results sought after by the “characterization project”; alas it refers only to a special class of graph properties (see above). Combined with previous results of [8, Section 9] and [2], this suggests that progress can be made with respect to subclasses of graph properties that can be expressed in some uniform structural formalism (rather than merely placed in a uniform complexity class such as \mathcal{NP}).

1.3. Organization

The above-mentioned results are stated formally in Section 2, and their proofs appear in the subsequent sections (i.e., in Sections 3, 4, and 5, respectively).

2. FORMAL SETTING

For any natural number n , we let $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$. We consider finite, undirected, labeled graphs without parallel edges. Without loss of generality, all N -vertex graphs have $[N]$ as their vertex set, and their edges are unordered pairs over $[N]$.

2.1. Graph Properties and Distance to Them

A graph property Π is a predicate defined over graphs that is preserved under graph isomorphism (i.e., if G has property Π and G' is isomorphic to G then G' has property Π).

We say that a graph $G = ([N], E)$ is ε -close to having property Π if there exists a graph $G' = ([N], E')$ having property Π such that the symmetric difference between E and E' is at most $\varepsilon \cdot \binom{N}{2}$. We say that a graph G is ε -far from having property Π if it is NOT ε -close to having property Π . A useful observation follows:

Claim 2.1. *If G is ε -close to (resp., ε -far from) having property Π , then so is any graph that is isomorphic to G .*

Proof. Suppose that $G = ([N], E)$ is ε -close to Π and $G' = \pi(G) \stackrel{\text{def}}{=} ([N], \{(\pi(u), \pi(v)) : (u, v) \in E\})$ for some permutation $\pi: [N] \rightarrow [N]$. Let H be a graph having property Π such that the graphs G and H differ on at most $\varepsilon \cdot \binom{N}{2}$ edges. Then, $H' = \pi(H)$ also has property Π and the graphs G' and H' differ on at most $\varepsilon \cdot \binom{N}{2}$ edges. ■

2.2. Testers for Graph Properties

Testers are oracle machines that are given as input a pair (N, ε) , where N is a size parameter and $\varepsilon > 0$ is a distance parameter, as well as oracle access to (the adjacency

matrix) of an N -vertex graph. An oracle machine T is called a **tester for property Π** if for every $G = ([N], E)$ and every ε , the following two conditions hold:

1. If G has property Π , then $\Pr[T^G(N, \varepsilon) = 1] \geq \frac{2}{3}$.
2. If G is ε -far from having property Π , then $\Pr[T^G(N, \varepsilon) = 1] \leq \frac{1}{3}$.

In both items, the probability space is that of the internal coin tosses of machine T , and a typical query (u, v) to oracle G is answered by 1 iff the edge (u, v) is in the graph G . The tester T (for Π) is said to have **one-sided error** if it always accepts graphs having the property Π ; that is, for every $G = ([N], E)$ having the property Π and every ε , it holds that $\Pr[T^G(N, \varepsilon) = 1] = 1$.

The query complexity of a tester T is a function $q: \mathbb{N} \times [0, 1] \rightarrow \mathbb{N}$ such that $q(N, \varepsilon)$ is an upper bound on the number of queries made by T on input (N, ε) and oracle access to the adjacency predicate of any N -vertex graph. The query complexity of a property Π is the minimum query complexity of testers for Π .

2.3. Statement of the Main Results

Existence of hard-to-test monotone graph properties (in \mathcal{NP}). A graph property Π is called **monotone** if adding any edge to any graph that has property Π results in a graph that has property Π . By saying that a graph property Π is in \mathcal{NP} , we mean the natural thing; that is, that the problem of deciding whether a given graph has property Π is in \mathcal{NP} (i.e., the set Π is in \mathcal{NP}).

Theorem 1. *There exists a monotone graph property Π in \mathcal{NP} such that the query complexity of Π , denoted q_Π , satisfies $q_\Pi(N, 0.1) = \Omega(N^2)$.*

Recall that $q_\Pi(N, \varepsilon)$ is a lower bound on the number of queries made by any tester for Π on input (N, ε) and oracle access to (the adjacency predicate of) any N -vertex graph. Theorem 1 is proven in Section 3.

Canonical forms of testers for graph properties. Let Π be any graph property and T be a tester for Π . We say that T is **canonical** if, for some function $s: \mathbb{N} \times [0, 1] \rightarrow \mathbb{N}$ and graph property Π' , the tester operates as follows: on input (N, ε) and oracle access to any N -vertex graph G , the tester T selects uniformly a set of $s(N, \varepsilon)$ vertices (in G), and accepts if and only if the corresponding induced subgraph (of G) has property Π' . Clearly, the query complexity of such a tester is $q(N, \varepsilon) = \binom{s(N, \varepsilon)}{2}$.

Theorem 2. *Let Π be any graph property having query complexity q_Π . Then Π has a canonical tester of query complexity $q'_\Pi(N, \varepsilon) = O(q_\Pi(N, \varepsilon)^2)$. Furthermore, if Π has a one-sided error tester of query complexity q , then Π has a canonical tester that has one-sided error and query complexity $\binom{2q}{2}$.*

In particular, it follows that in the context of testing graph properties, the query complexity of nonadaptive¹ algorithms is polynomially related to the query complexity of adaptive algorithms. Theorem 2 is proven in Section 4.

A characterization of graph partition properties that are easily testable with one-sided error. We refer to *graph partition properties* as described in the introduction and further discussed in Appendix A. We confine ourselves to “nontrivial” graph properties. That is, a graph property Π is nontrivial if for infinitely many N ’s there exist an N -vertex graph satisfying property Π as well as an N -vertex graph not satisfying property Π . For a k -vertex graph H and a graph G , we say that G is H -embeddable if the vertices of G can be k -partitioned such that there are no edges (of G) among vertices residing in the same part and so that there are edges between vertices of the i th part and j th part only if (i, j) is an edge of H . (For example, saying that G is C_k -embeddable, where C_k denotes the k -vertex clique, is equivalent to saying that G is k -colorable.)

Theorem 3. *Let Π be a nontrivial graph partition property that is testable with one-sided error and query-complexity independent of N . Then exactly one of the following two cases holds:*

1. *There exists a k -vertex graph H such that, for every sufficiently large graph G , the graph G satisfies Π if and only if G is H -embeddable.*
2. *For all sufficiently large N , an N -vertex graph has property Π if and only if it is an N -vertex clique.*

Theorem 3 is proven in Section 5.

2.4. Tools Used

In Section 3, we use efficient constructions of small-bias probability spaces, which were defined and first constructed by Naor and Naor [10]. An ε -biased sample space over $\{0, 1\}^n$, is a multiset S such that, for every nonempty set $I \subseteq [n]$, if $s = s_1 \cdots s_n$ is selected uniformly in S then

$$\left| \Pr[\bigoplus_{i \in I} s_i = 1] - \frac{1}{2} \right| \leq \varepsilon$$

Such sample spaces can be constructed in time $\text{poly}(n/\varepsilon)$; specifically, $|S| = (n/\varepsilon)^2$ suffices (cf. [3], following [10]). We will also use the following well-known fact (cf. [3, Appendix]):

Fact 2.2. *Let $\varepsilon \geq 0$, $t \in [n]$, and S be an ε -biased sample space over $\{0, 1\}^n$. Then, for every t -subset $I \subseteq [n]$ and every $\alpha = \alpha_1 \cdots \alpha_n \in \{0, 1\}^t$, if $s = s_1 \cdots s_n$ is selected uniformly in S , then it holds that $|\Pr[(\forall i \in I) s_i = \alpha_i] - 2^{-t}| \leq \varepsilon$.*

¹An oracle machine is called *nonadaptive* if it determines its queries based merely on its input and random-coins, independently of the answers to prior queries.

In addition, we will use the fact that certain “Laws of Large Numbers” hold also for ε -biased sample spaces. Specifically:

Fact 2.3. *Let $\varepsilon \geq 0$, $\delta > 0$, and S be an ε -biased sample space over $\{0, 1\}^n$. Then, for $s = s_1 \cdots s_n$ uniformly distributed in S it holds that $\Pr[\sum_{i=1}^n s_i \leq (0.5 - \delta) \cdot n] < (2\varepsilon + n^{-1})/\delta^2$.*

Proof. Using Chebyshev’s Inequality, we get (see, e.g., [5, Section 4.1]) that for any sequence of random variables X_1, \dots, X_n in $[0, 1]$, where each X_i has expectation μ , it holds that

$$\Pr\left[\sum_i X_i \leq (\mu - \delta) \cdot n\right] \leq \frac{n\mu + \sum_{i \neq j} \text{cov}[X_i, X_j]}{(\delta n)^2} \leq \frac{\mu/n}{\delta^2} + \frac{\max_{i \neq j} \{\text{cov}[X_i, X_j]\}}{\delta^2},$$

where $\text{cov}[X_i, X_j]$ is the covariance of X_i and X_j , which in our case (i.e., $X_i = s_i$) is less than 2ε . The claimed upper-bound follows. ■

3. MONOTONE GRAPH PROPERTIES MAY BE VERY HARD TO TEST

Throughout this section we consider the query complexity of testing, when setting the distance parameter to equal a constant (e.g., $\varepsilon = 0.1$). In contrast, the size of the graph (denoted N) is treated as a parameter. Thus, when we describe a set of N -vertex graphs, it is to be understood that we consider the union of all these sets (i.e., over all possible N ’s).

3.1. Motivation

Goldreich, Goldwasser, and Ron showed that there exist graph properties in $\mathcal{N}^{\mathcal{G}}$ for which any tester must inspect at least a constant fraction of the vertex-pairs [8, Property 10.2.3.2]. Their construction proceeds in two stages:

1. First, it is shown that certain sample spaces yield a collection of Boolean functions (i.e., a property of Boolean functions) that is hard to test (i.e., any tester must inspect at least a constant fraction of the function’s values).

On one hand, the sample space is relatively sparse (and thus a random function is far from any function in the resulting collection), but, on the other hand, it enjoys a strong pseudorandom feature (and so its projection on any constant fraction of the coordinates looks random). Thus, the functions in the class (which must be accepted with high probability) look random to any tester that inspect only a small constant fraction of the function’s values, whereas random functions are far from the class (and should be rejected with high probability). This yields a contradiction to the existence of a tester that inspect only a small constant fraction of the function’s values.

2. Next, the domain of the functions is associated with the set of unordered pairs of elements in $[N]$, and the collection of functions is “closed” under graph isomorphism (i.e., if a certain function on $\binom{N}{2}$ is in the collection then so is any function obtained from it by a relabeling of the elements of $[N]$).

The closure operation makes this collection correspond to a graph property (since it is now preserved under isomorphism). The parameters are such that the resulting collection (although likely to be $N!$ times bigger than the original one) is still sparse enough (and so a random graph is far from it). On the other hand, the indistinguishability feature is maintained.

Unfortunately, the above construction does not yield a monotone property (since the second stage inherits the possible non-monotonicity of the collection constructed at the first stage). We redeem the situation by adding a third stage in which the collection is “closed” under edge-additions. Clearly, this guarantees that the resulting (graph) property is monotone (and it also inherits the feature of being in \mathcal{NP}). Furthermore, using a crucial technical condition that we enforce on the collection constructed at the first stage, the resulting collection is also sparse. (Specifically, we force each initial function to have at least one third of one-entries.)

3.2. The Actual Construction

The actual construction follows the three stages sketched in the motivation above. For every N , we start by considering an efficiently constructible $0.1 \cdot 2^{-t}$ -biased sample space over $\binom{N}{2}$ -bit long strings (see Section 2.4), where $t \stackrel{\text{def}}{=} \frac{1}{200} N^2$. Recall that such sample spaces of size $O(\binom{N}{2}/(0.1 \cdot 2^{-t}))^2 = 2^{2t+o(t)}$ do exist. We actually omit from the sample space any sample that has less than one third of one-entries. (The importance of this modification and its effect will be analyzed below.) For each sample s in the (residual) space, we define a graph $G_s = ([N], E_s)$ by letting $(i, j) \in E_s$ if and only if the (i, j) th bit of s equals 1, where we consider any fixed (efficiently computable) order of the elements in $\{(i, j) : 1 \leq i < j \leq N\}$. We call these $2^{2t+o(t)}$ graphs, the **basic graphs**. Note that the set of basic graphs is not likely to be closed under isomorphism, and thus this collection does not constitute a graph property. On the other hand, the set of basic graphs is in \mathcal{NP} , because elements in the sample space can be generated in time polynomial in N (i.e., there exists a $\text{poly}(N)$ -time algorithm that given an $(2t + o(t))$ -bit long string produces an $\binom{N}{2}$ -bit long string in the sample space).²

Next, we consider the set of **secondary graphs** obtained by “closing” the set of basic graphs under isomorphism. That is, for every s in the sample space (equivalently, a basic graph G_s) and every permutation π over $[N]$, we consider the secondary graph $G_{s,\pi} = ([N], E_{s,\pi})$ that is defined so that $(\pi(u), \pi(v)) \in E_{s,\pi}$ iff $(u, v) \in E_s$. By construction, the set of secondary graphs is closed under isomorphism, and so this collection does constitute a graph property. The latter set is also in \mathcal{NP} (by incorporating the isomorphism π in the NP-witness).

Finally, we “close” the set of secondary graphs under edge-addition, obtaining our final set of graphs (which is, by construction, monotone): That is, for every secondary graph $G' = ([N], E')$, and every $E'' \supseteq E'$, we introduce the **final graph** $G'' = ([N], E'')$.

²In fact, using known constructions, membership in the (original) sample space can be decided in polynomial-time. Observing that the omitting condition is also decidable in polynomial-time, we conclude that the set of basic graphs is actually in \mathcal{P} .

Comment. At the point, the reason for the modification in the initial sample space may be clear: If, for example, the sample space had contained the all-zero string, then the set of final graphs would have contained all graphs, and testing membership in it would have been trivial.

Analysis. Our aim is to show that, although a random graph is far from the set of final graphs, no algorithm that makes $o(N^2)$ queries can distinguish a random graph from a graph selected among the final graphs (with some distribution that is not necessarily uniform). Since a tester for the set of final graphs must accept any final graph (with high probability) and reject a random graph (with high probability), we conclude that such tester must make $\Omega(N^2)$ queries. Specifically, throughout the rest of the analysis, we refer to testers of N -vertex graphs that should accept with probability at least $\frac{2}{3}$ every graph that has the property, and reject with probability at least $\frac{2}{3}$ every graph that is 0.1-far from having the property. Thus, we omit the distance parameter ε , which always equals 0.1, from all notations.

Claim 3.1. *The probability that a random graph is 0.1-close to some final graph is at most 0.1.*

The original proof can be found in our Technical Report [9]. The following proof was suggested to us by Amir Shpilka.

Proof. The key observation is that the set of final graphs is very sparse.³ Specifically, each basic graph gives rise to at most $N!$ secondary graphs, and each of the latter gives rise to at most $2^{2/3 \binom{N}{2}} < 2^{N^{2/3}}$ final graphs (because each secondary graph, which by the modified construction, has at least $\frac{1}{3} \binom{N}{2}$ edges, misses at most $\frac{2}{3} \binom{N}{2}$ edges). Thus, the number of final graphs is at most

$$2^{2t+o(t)} \cdot (N!) \cdot 2^{N^{2/3}} = 2^{((2/200)+o(1)+(1/3))+2 \cdot \binom{N}{2}} < 2^{0.7 \cdot \binom{N}{2}}, \quad (1)$$

where the above holds for all sufficiently large N . Thus, for all sufficiently large N , less than 10% of the graphs are 0.1-close to some final graph. ■

Claim 3.2. *Let M be a probabilistic oracle machine that makes at most t queries. Let R_N denote a random graph, and B_N denote a graph uniformly selected among the basic graphs. Then,*

$$|\Pr[M^{R_N}(N) = 1] - \Pr[M^{B_N}(N) = 1]| < 0.2.$$

Proof. We identify $\binom{N}{2}$ -bit long strings with N -vertex graphs (obtained as in the first stage of the construction). Let G_N denote a graph uniformly selected among all graphs in the sample space; that is, without discarding from the space those samples having less than one third of one-entries (equivalently, less than $\frac{1}{3} \binom{N}{2}$ edges). Thus, B_N is obtained from G_N by conditioning that G_N has at least $\frac{1}{3} \binom{N}{2}$ edges. Using the fact that a small bias sample space (as above) is almost pairwise independent, the probability that an element in it has

³This observation has originally escaped our attention due to the way we have developed our construction.

less than one third of one-entries is very small (e.g., tends to zero when $N \rightarrow \infty$, see Fact 2.3). Thus, the probability that G_N has at least $\frac{1}{3}\binom{N}{2}$ edges is overwhelmingly high (e.g., is at least 0.9), and so the statistical difference between G_N and B_N is very small (e.g., smaller than 0.1). It follows that

$$|\Pr[M^{G_N}(N) = 1] - \Pr[M^{B_N}(N) = 1]| < 0.1. \quad (2)$$

On the other hand, since the sample space underlying the construction of G_N has bias at most $0.1 \cdot 2^{-t}$, it follows (by Fact 2.2) that, for every t distinct (unordered) pairs $(u_1, v_1), \dots, (u_t, v_t) \in [N] \times [N]$ and every $\alpha_1, \dots, \alpha_t \in \{0, 1\}^t$, the probability that for every i query (u_i, v_i) to G_N is answered by α_i equals $2^{-t} \pm 0.1 \cdot 2^{-t}$. Now, observe that the internal coins of M together with the oracle answers to M determine oracle queries of M , and assume (without loss of generality) that M does not make redundant queries (i.e., it never repeats the same query, nor makes query (v, u) after making the query (u, v)). It follows that, for any fixed sequence of coins for M , any fixed sequence of t answers occurs with probability $2^{-t} \pm 0.1 \cdot 2^{-t}$ under G_N (rather than with probability 2^{-t} under R_N). Thus, for any fixed sequence of coins for M , the observed deviation of the t answers of G_N from the t answers of R_N is at most 0.1. We conclude that

$$|\Pr[M^{G_N}(N) = 1] - \Pr[M^{R_N}(N) = 1]| < 0.1. \quad (3)$$

Combining Eqs. (2) and (3), the claim follows. \blacksquare

Combining Claims 3.1 and 3.2, we obtain:

Theorem 3.3 (Theorem 1, Restated). *There exists a monotone graph property in $\mathcal{N}\mathcal{P}$ for which every tester requires $\Omega(N^2)$ queries (even when invoked with constant distance parameter).*

Proof. Consider the graph property, denoted Π , corresponding to the set of final graphs defined above. Recall that this set indeed corresponds to a monotone graph property in $\mathcal{N}\mathcal{P}$. Now, suppose that M is a tester for this property and that M makes less than $N^2/200$ queries (when invoked with distance parameter 0.1). Then, by Claim 3.2,

$$|\Pr[M^{R_N}(N) = 1] - \Pr[M^{B_N}(N) = 1]| < 0.2. \quad (4)$$

Now, since each graph in the support of B_N (i.e., each basic graph) has property Π , the tester must accept (i.e., output 1 on input) such graph with probability at least $\frac{2}{3}$. It follows that

$$\Pr[M^{B_N}(N) = 1] \geq \frac{2}{3} > 0.6 \quad (5)$$

On the other hand, the tester may accept with probability at most $\frac{1}{3}$ each graph that is 0.1-far from having property Π . By Claim 3.1, the probability that R_N is 0.1-far from having property Π is at least 0.9. It follows that

$$\Pr[M^{R_N}(N) = 1] \leq 0.9 \cdot \frac{1}{3} + 0.1 \cdot 1 = 0.4 \quad (6)$$

Combining Eq. (4)–(6), we reach a contradiction. The theorem follows. ■

4. CANONICAL FORMS OF GRAPH-PROPERTY TESTERS

We present two “canonization” transformations that can be applied to algorithms that test graph properties. The two “canonization” transformations are:

1. A transformation to testers that inspect a random induced subgraph.
2. A transformation (from the latter) to testers that decide according to whether the induced subgraph has some (possibly other) graph property.

Both transformations incur only a polynomial increase in the query complexity, and preserve one-sided error.

4.1. Moving to Testers That Inspect a Random Induced Subgraph

A natural way of transforming testers to the desired form proceeds by first making them nonadaptive and then arguing that without loss of generality they may inspect a (random) induced subgraph. Focusing on the first step, we note that the straightforward way of getting rid of the (possible) adaptivity of the original tester is to consider, for each i , all possible 2^{i-1} answers to the first $i-1$ original queries and perform (*nonadaptively*) all the corresponding (i.e., 2^{i-1}) queries that the original tester could have made as its i th query. This transformation incurs an exponential increase in the query complexity, and our goal is to present a transformation that only incurs a polynomial increase. Our transformation proceeds in a reverse order to the one outlined above: First we move to testers that inspect an induced subgraph, and only next do we get rid of the potential adaptivity of the tester. In the latter argument, we capitalized on the fact that (by definition) graph properties are preserved under isomorphism (i.e., renaming of vertex names).

Lemma 4.1. *Let Π be any graph property, and T be an arbitrary tester for Π . Suppose that T has query complexity $q(N, \varepsilon)$. Then there exists a tester for Π that selects a random subset of $2q(N, \varepsilon)$ vertices, denoted R , makes the queries $\{(u, v) : u, v \in R\}$, and decides based on the oracle answers (and its internal coin tosses). Thus, the new tester is non-adaptive and its query complexity is less than $2q(N, \varepsilon)^2$. Furthermore, if T has one-sided error then so does the new tester.*

Proof. First we transform T into an algorithm T' that belongs to a class of *vertex-uncovering* algorithms, defined as follows: A *vertex-uncovering algorithm* operates in iterations such that, in each iteration, depending on its coins and the answers obtained in previous iterations, the algorithm selects a new vertex, denoted v , and makes queries to all pairs (v, u) , where u is a vertex selected in some prior iteration. Clearly, T can be emulated by a vertex-uncovering algorithm, denoted T' , that makes at most $2q(N, \varepsilon)$ iterations, and thus at most $\binom{2q(N, \varepsilon)}{2} < 2q(N, \varepsilon)^2$ queries: Each query of T is emulated by two iterations of T' , while these iterations are not necessarily new ones.

We next consider an algorithm T'' obtained from T' as follows. When given oracle access to a graph $G = ([N], E)$, algorithm T'' first selects uniformly a permutation π over

$[N]$, and next invokes T' providing it with oracle access to the graph $\pi(G) \stackrel{\text{def}}{=} ([N], \pi(E))$, where $\pi(E) \stackrel{\text{def}}{=} \{(\pi(u), \pi(v)) : (u, v) \in E\}$. That is, when T' makes query (u, v) , algorithm T'' makes the query $(\pi(u), \pi(v))$ and feeds the answer to T' . Clearly, algorithm T'' also operates in a vertex-uncovering manner. Furthermore, the permutation π may be selected on-the-fly in the sense that whenever T' determines a new vertex v to be uncovered, algorithm T'' chooses $\pi(v)$ uniformly among all vertices of G that it (i.e., T'') did not inspect yet.

By the alternative implementation of T'' (i.e., when π is selected on-the-fly), it follows that the sequence of vertices uncovered by T'' is uniformly selected among all sequences of the same length.⁴ It follows that T'' inspects a random induced subgraph of G . Thus, it remains to show that T'' maintains the testing features of T' (and hence of T). For that purpose, it is more convenient to consider the first implementation of T'' (i.e., in which π is uniformly selected and fixed at the onset of T'').

- Let $G = ([N], E)$ be a graph having property Π . Then, for any permutation π , it is the case that the graph $\pi(G) \stackrel{\text{def}}{=} ([N], \{(\pi(u), \pi(v)) : (u, v) \in E\})$ has property Π . Thus, T' must accept the graph $\pi(G)$ with probability at least $\frac{2}{3}$. This means that (for every permutation π) conditioned on π being chosen in the onset of T'' , algorithm T'' accepts the graph G with probability at least $\frac{2}{3}$ (because, in this case, T'' just emulates for T' an oracle access to the graph $\pi(G)$). Since this holds for every π , it follows that T'' accepts G with probability at least $\frac{2}{3}$.
- Suppose that a graph $G = ([N], E)$ is ε -far from having property Π . Then (by Claim 2.1), for any permutation π , the graph $\pi(G)$ is ε -far from having property Π . The rest of the argument follows analogously to the above (where here we refer to an upper bound on the accepting probability). It follows that T'' accepts G with probability at most $\frac{1}{3}$.

The lemma follows. ■

Perspective. The main part of the proof of Lemma 4.1 is similar to an analogous statement proven by Bar-Yossef, Kumar, and Sivakumar [6] in the context of “sampling algorithms.” Consider, for example, the problem of approximating the average value of a function f defined over a huge space, say $f : \{0, 1\}^n \rightarrow [0, 1]$, when given only oracle access to the function. We seek algorithms that make relatively few oracle calls, and call them **samplers**. The proof of Lemma 4.1 can be easily modified (and in fact simplified) to prove that *the query complexity of nonadaptive samplers equals the query complexity of adaptive ones*. The key observation is that also here the relevant property (i.e., the average value of a function) is invariant under renaming (of the function’s arguments). In fact, this is exactly the way this statement (regarding query complexity of “sampling algorithms”) is proven in Lemma 9 of [6].

⁴A detailed proof of this simple fact can be found in our Technical Report [9]. In particular, it is useful to first consider the special case of a deterministic T' , and then to reduce to this case by fixing the coins of T' in all possible ways.

An open problem. The main part of the proof of Lemma 4.1 made essential use of the fact that T' is a vertex-uncovering algorithm. In particular, in such an algorithm, the queries are determined by a sequence of vertices, whereas for a random isomorphic copy of the graph all *vertex*-sequences are essentially equivalent. Things are different in case of a general algorithm, because the *query*-sequences [e.g., the 2-query sequences $((u, v), (v, w))$ and $((u, v), (u', v'))$] may not be equivalent (even w.r.t. a random isomorphic copy of the graph). This raises the question of whether or not nonadaptive testers are as powerful as adaptive ones, even when insisting on strictly the same query complexity (rather than allowing a polynomial relationship).

4.2. Moving to a Decision Determined by the Induced Subgraph

The current transformation is less generic than the previous one, since it applies only to nonadaptive (graph property) testers. For simplicity, our starting point is actually testers as resulting from Lemma 4.1 (i.e., that query the oracle on the edges of a random induced subgraph). The current transformation consists of three steps:

1. A transformation to testers that decide based on the subgraph they see, possibly by tossing new coins, but independently of the coins used to select the sample set of vertices. In particular, if the tester selects t vertices, then it decides based only on the induced t -vertex subgraph in which the vertices are labeled by the elements of $[t]$ according to some canonical order.
2. A transformation to testers that decide as above, but do so in a way that is independent of the labeling of the induced subgraph. That is, the decision depends only on an unlabeled version of the induced subgraph.
3. A transformation to testers that decide according to whether or not the induced subgraph has some fixed graph property. That is, this transformation gets rid of the coins that were possibly used in the previous decision process.

That is, the first transformation gets rid of the possible dependence of the decision on the identity of the sampled vertices; that is, the resulting decision depends only on the subgraph seen by the tester (but possibly depends on the ordering of the vertices in this subgraph). The second transformation makes this decision identical for all isomorphic subgraphs (i.e., independent of this ordering), and the third transformation makes this decision deterministic.

In order to facilitate the third transformation, we reduce the error of the tester (from $\frac{1}{3}$ to $\frac{1}{6}$) before employing the second transformation. Note that error reduction is trivial in case the resulting decision does not have to be independent of the labels of the induced subgraph, because in that case we just apply the original test on several independently chosen induced random graphs. Since the first two transformation do not effect the error probability of the test, we will apply the third transformation to testers having error at most $\frac{1}{6}$. (See further comments in Section 4.3).

4.2.1. Moving to a Sample-Independent Decision. We first move to testers in which the sample of vertices only determines the queries, but plays no direct role in the final decision (which depends only on the answers to these queries).

Claim 4.2. *Let Π be any graph property, and T be a nonadaptive tester for Π that selects a random subset of $t = t(N, \varepsilon)$ vertices, makes queries to determine the induced subgraph, and decides based on the oracle answers (and its internal coin tosses). Then, without loss of generality, the tester T can be decomposed into two parts.*

- *The first part uniformly selects a set of vertices and queries the oracle for the induced subgraph, which it passes to the second part.*
- *The second part makes a decision based on the subgraph obtained from the first part and possibly depending on its own coins, but independent of the coins used by the first part.*

We stress that the resulting test preserves the error bounds of the original test on both graphs having property Π and being ε -far from it. Furthermore, if T has one-sided error then no coins are tossed in the second part.

Proof. Without loss of generality, we may decompose the random-tape of T into the form (S, r) , where S is the vertex set selected by T [with probability $p_S = 1/\binom{N}{t}$] and r is the residual randomness. (This step is quite generic: the original random-tape of T induces uniform distribution over the possible vertex-sets, and each possible vertex-set induces uniform distribution over all random-tapes that cause T to select this set.)

For each set of vertices S and each sequence of possible answers α (which is a symmetric relation over S), we denote by $q_{S,\alpha}$ the probability that T accepts, having selected the vertex set S and seeing the answer sequence α . Indeed, $q_{S,\alpha}$ is merely the fraction of r 's that make T accept when its sees answers α (after selecting S). For any fixed α , let $q_\alpha \stackrel{\text{def}}{=} \sum_S p_S q_{S,\alpha}$ be the expected value of $q_{S,\alpha}$ (for varying S). We now derive a tester that selects S as T does, but decides only according to the answers it gets (i.e., independently of S). Specifically, we consider an algorithm T' that selects S with probability p_S , and (for every α) upon obtaining the answer sequence α accepts with probability q_α .

Claim: For any graph G , when given oracle access G , the probability that T' accepts G equals the probability that T accepts a random isomorphic copy of G .

Proof: Let π be a uniformly distributed permutation over $[N]$, and consider the operation of T when given oracle access to $\pi(G)$. Note that T selects S uniformly, inspects the subgraph of $\pi(G)$ induced by S [which equals the subgraph of G induced by $\pi^{-1}(S)$], and accepts with probability $q_{S,\alpha}$, where α is the relation representing this subgraph. In terms of G , the situation is the same as if T would have selected uniformly and independently two sets S and S' (since S' corresponds to $\pi^{-1}(S)$ where π is a uniformly distributed permutation) and accepts with probability $q_{S,\alpha}$, where α is the relation representing the subgraph of G induced by S' . Taking another step, it is as if T would have selected uniformly a set S' and, upon obtaining the answer sequence α , accepts with probability $q_\alpha = E_S(q_{S,\alpha})$ (because S is selected independently of S' and is only used to determine $q_{S,\alpha}$). But this exactly what T' actually does when given oracle access to G . \square

As in the proof of Lemma 4.1, it follows that T' is a tester for Π . The main part of the claim follows.

Finally we observe that if T has one-sided error then, without loss of generality, all the q_α 's may be in $\{0, 1\}$ (and so no coins are needed for the final decision). The reason is that if any q_α is strictly smaller than 1 then the answer sequence α cannot occur when accessing a graph that has property Π (or else this graph is rejected with nonzero probability), and so setting this q_α to zero does not effect the performance on graphs having property Π and may only improve the performance on all other graphs. ■

4.2.2. Moving to a Isomorphism-Oblivious Decision. Claim 4.2 asserts that, without loss of generality, the final decision of a tester is obtained by a randomized computation that depends only on the oracle answers (which are obtained by querying all pairs in a random set). We now show that, without loss of generality, this decision is closed under isomorphism. In the special case of one-sided testers, this means that the final decision is by whether or not the induced subgraph satisfies some fixed graph property.

Claim 4.3. *Let Π and T be as in Claim 4.2. Then, without loss of generality, T is composed of two parts as in Claim 4.2, and in the second part the decision applied to the answer is closed under graph isomorphism. That is, when seeing an induced subgraph that equals H , the second part decides exactly as in case it sees a subgraph that is isomorphic to H . Again, the resulting test preserves the error bounds of the original test on both graphs having property Π and being ε -far from it. Furthermore, if T has one-sided error then, without loss of generality, there exists a graph property Π' such that the final decision of T amounts to checking whether or not the subgraph induced by its vertex sample has property Π' .*

Proof sketch. The proof is very similar to the proof of Claim 4.2. Denoting by q_α the probability that T accepts on answer sequence α , we consider an algorithm T' that accepts with probability that depends only on the class of graphs isomorphic to α . Specifically, let $g(\alpha)$ be the set of graphs that are isomorphic to the graph having an adjacency matrix that corresponds to α , and let q'_H be the expected value of q_α taken uniformly over all α such that $g(\alpha) \ni H$. Then, T' selects uniformly a set of vertices S , and accepts with probability q'_H , where H is the subgraph induced by S .

As in the proof of Claim 4.2, we show that, for any graph G , the probability that T' accepts G equals the probability that T accepts a random isomorphic copy of G . [Here we use the fact that when T accesses a random isomorphic copy of G and obtains the answer sequence α , it decides as if it has accessed G and obtained an answer that corresponds to a random element in $g(\alpha)$.] Finally, again, we observe that if T is one-sided then, without loss of generality, each q'_H is either 0 or 1. The claim follows. ■

4.2.3. Moving to a Deterministic Decision. Here, our aim is to make the second part of the tester guaranteed by Claim 4.3 be deterministic (i.e., be determined by whether or not the induced subgraph has some fixed graph property). By the furthermore parts of Claim 4.3, it suffices to consider two-sided error testers (because for one-sided error testers the final decision was already shown to be deterministic). Thus, our focus in this subsection is on two-sided error testers. In fact, we can directly handle testers as in Claim 4.2 (i.e., with a decision that is not closed under isomorphism of the induced subgraph), but for sake of simplicity we prefer to use Claim 4.3 as our starting point.

Lemma 4.4. *Let Π and T be as in Claim 4.3, and let t denote the number of vertices selected by T . Assume that the error probability of T is at most $\frac{1}{6}$. Then, there exists a tester T' and a graph property Π' so that T' selects a random set of t vertices and accepts if and only if the subgraph induced by it has property Π' .*

We stress that the error probability of the resulting tester (i.e., T') may be twice as big as the error probability of the original tester (i.e., T). However, since we assumed that T has error probability at most $\frac{1}{6}$, the resulting T' (having error probability at most $\frac{1}{3}$) is a valid tester. The following proof was suggested to us by an anonymous referee, and improves our original proof that appears in [9].

Proof. The idea is that the new tester T' will inspect the induced subgraph and accept the input graph if and only if the probability associated with the induced subgraph (by the original tester) is at least $\frac{1}{2}$. That is, let us denote by p_H the probability that T accepts a graph when seeing H as the induced (t -vertex) subgraph. Consider a new algorithm T' that when given oracle access to an input graph G , selects at random a set of t vertices, denoted S , and accepts if and only if $p_{G_S} \geq \frac{1}{2}$, where G_S is the subgraph of G induced by S .

We first note that the acceptance probability of T' , on any input, is at most twice as much as the acceptance probability of T . This is because the subgraph induced by S contributes (a unit) to the acceptance probability of T' only if it contributes at least $\frac{1}{2}$ to the acceptance probability of T . Since each graph that is ε -far from Π is accepted by T with probability at most $\frac{1}{6}$, it follows that such a graph is accepted by T' with probability at most $2 \cdot \frac{1}{6} = \frac{1}{3}$. On the other hand, for each graph that has property Π , it holds that $\mathbb{E}_S(p_{G_S}) \geq \frac{5}{6}$. Thus, $\Pr_S[p_{G_S} < \frac{1}{2}] < \frac{1}{3}$, which means that the probability that T' selects S such that $p_{G_S} < \frac{1}{2}$ is smaller than $\frac{1}{3}$. The lemma follows.⁵ ■

4.2.4. Question: Does Π equal Π' ? In general, the question does not make sense since Π is a property of N -vertex graphs, whereas Π' is a property of t -vertex graphs.⁶ But the question does make sense if these properties can be expressed in a uniform way that is independent of the size of the graph. Two specific frameworks for such properties were presented in [8] and [2], respectively:

1. The framework of graph partition problems [8, Section 9] is briefly reviewed in the introduction (see also Appendix A), and is further studied in Section 5. Goldreich, Goldwasser, and Ron [8] showed that every problem Π in this framework can be tested by checking whether a random induced subgraph of $\text{poly}(1/\varepsilon)$ vertices satisfies a related graph property Π' . In general, Π' is not equal to Π , and the tester has two-sided error probability. (In fact, this is unavoidable for some properties Π ; for example, the property of the graph consisting of two equal size parts one being a clique and the other being an independent set.) Confining ourselves to properties Π that admit a one-sided error tester of query complexity independent of N , we show that, without loss of generality, Π' equals Π (see Corollary 5.9). This follows

⁵In general, if T rejects graphs having property Π with probability at most ε , then $\mathbb{E}_S(p_{G_S}) \geq 1 - \varepsilon$, which implies that $\Pr_S[p_{G_S} < \frac{1}{2}] < 2\varepsilon$.

⁶Formally, both Π and Π' are properties of all graphs, but when we fix an N for the above discussion, we actually care of the properties Π and Π' when restricted to N -vertex and t -vertex graphs, respectively.

from the characterization of properties Π that admit a one-sided error tester of query complexity independent of N (see Theorem 5.8).

2. Alon et al. [2] considered the class of graph properties that can be expressed by quantified boolean formula over the edge relation E . For example, the property of being triangle-free is represented by the formula

$$\forall x, y, z [(x, y) \notin E \vee (y, z) \notin E \vee (z, x) \notin E]$$

For graph properties Π expressible by formulae of the form $\exists\forall$, Alon et al. [2] presented a tester of query complexity that only depends on ε (but grows very rapidly with $1/\varepsilon$).⁷ Their tester checks whether a random induced subgraph of suitable size satisfies a related graph property Π' . Actually, they showed that for every such property Π can be approximated⁸ by a “graph coloring” property Π' that can be tested by checking whether or not a random induced subgraph of suitable size satisfies Π' itself. Thus, the class of general “graph coloring” properties can be tested by checking whether a random induced subgraph of suitable size satisfies the very same property.

In addition, as observed by Alon [1], for any property Π that is closed under taking induced subgraphs (i.e., if G has property Π , then so does every induced subgraph of G), if Π is easily testable, then it is easily testable by a canonical tester (as above) that uses $\Pi' = \Pi$. That is, if such a property Π is testable by $q(\varepsilon)$ queries (i.e., independent of the size of the graph), then it is so testable by inspecting a random induced subgraph of size $\text{poly}(q(\varepsilon))$ and accepting iff the subgraph has property Π . For details, see Appendix D.

4.3. Putting It Together

Combining Lemma 4.1, Claim 4.2, Claim 4.3, and Lemma 4.4, we obtain:

Theorem 4.5 (Theorem 2, restated). *Let Π be any graph property. If there exists a tester with query complexity $q(N, \varepsilon)$ for Π then there exists a tester for Π that uniformly selects a set of $O(q(N, \varepsilon))$ vertices and accepts iff the induced subgraph has property Π' , where Π' is some fixed graph property. Furthermore, if the original tester has one-sided error then so does the new tester, and, moreover, a sample of $2q(N, \varepsilon)$ vertices suffices.*

Proof. By Lemma 4.1, there exists a tester for Π that inspects a random induced subgraph of $2q(N, \varepsilon)$ vertices. In case the original tester had one-sided error, we conclude by invoking Claims 4.2 and 4.3. Otherwise, we first reduce the error probability of the resulting test to $\frac{1}{6}$ (by invoking it for a constant number of times and ruling by majority). Next, we invoke Claims 4.2 and 4.3. Finally, invoking Lemma 4.4, we are done. ■

⁷At best, their bound on the query complexity is a tower of $\text{poly}(1/\varepsilon)$ -many exponents. In general, their bound is a tower of towers.

⁸Loosely speaking, every graph having one of these properties is close to a graph having the other property.

5. ON GENERAL GRAPH PARTITION PROBLEMS THAT ARE TESTABLE WITH ONE-SIDED ERROR

We refer to the framework of **graph partition problems** [8, Section 9]. Recall that in this framework a testing problem is parameterized by a sequence of corresponding pairs of lower and upper bounds: For some (implicit) parameter k , the sequence contains k pairs of vertex-sets densities and $k + \binom{k}{2}$ pairs of edge-densities, and the problem is to determine whether there exists a k -partition of the vertices so that the number of vertices in each component of the partition as well as the number of edges within each component and between each pair of components falls between the corresponding lower and upper bounds. If such a partition exists, we call it a **witness partition**. For example, k -colorability falls into this framework by requiring that the density of edges within each of the k parts equals zero (and making no other requirements).⁹ For further details see Appendix A.

Goldreich, Goldwasser, and Ron [8] showed that every graph partition property (i.e., problem in the above framework) can be tested by making $\text{poly}(1/\varepsilon)$ queries, but in general the tester has two-sided error probability. They also gave one-sided error testers for k -colorability (which operate by checking whether a random induced $\text{poly}(1/\varepsilon)$ -vertex subgraph is k -colorable).¹⁰ Our main goal in this section is to characterize the set of graph partition properties that admit a one-sided error tester of query complexity that is independent of N .

A few technicalities. Throughout the discussion, we consider only admissible sequences of parameters (cf. [8, Definition 9.3.1]): These are sequences for which the set of graphs that have the property is infinite (i.e., contains at least one N -vertex graph, for infinitely many N 's). We avoid integrality problems by allowing up to $k - 1$ vertices to be split between the k parts of the partition (see [8, Remark 9.1]). Also, following [8], we consider vertex-densities as fractions of N , and edge-densities as fractions of N^2 [rather than of $\binom{N}{2}$]. Finally, for greater expressibility, we allow self-loops and count them as half edges: For example, using these conventions, a ρN -vertex clique in an N -vertex graph has edge density $\frac{\rho^2}{2}$, which is independent of N . (Note that the latter convention is not consistent with the exposition in the previous sections.) For further discussion of the latter two conventions, see Appendix A. We comment that the analysis can be carried out also under the alternative conventions mentioned above (and in Appendix A), but the exposition would be more cumbersome and the final result would be slightly different: The second case in Theorem 5.8 would not be possible.

The starting point. Using Theorem 4.5, we may confine ourselves to **canonical testers** (of one-sided error) that operate by inspecting the subgraph induced by a uniformly selected set of vertices, where the size of the vertex set is independent of N . Recall that

⁹Formally, the k lower-bound and upper-bound pairs on vertex-set sizes are all trivial [i.e., all equal $(0, 1)$], and so are the $\binom{k}{2}$ pairs of bounds on of edge-densities between pairs of components. The only nontrivial pairs of bounds are those referring to edge-densities within each of the k components [i.e., all these pairs equal $(0, 0)$].

¹⁰An improved bound on the size of the sample was later presented by Alon and Krivelevich [4]: They showed that a sample of size $\tilde{O}(1/\varepsilon)$ [rather than $\tilde{O}(1/\varepsilon^2)$] suffices for $k = 2$, and size $\tilde{O}(1/\varepsilon^2)$ [rather than $\tilde{O}(1/\varepsilon^3)$] suffices for any constant $k > 2$.

the inspection (of the subgraph) consists of determining whether or not it has some graph property Π' (not necessarily equal to the property Π being tested).

5.1. Trivial Graph Properties

A graph property is called *trivial* if for every $\varepsilon > 0$ and for all sufficiently large N , every N -vertex graph is ε -close to having the property. We may discard trivial graph properties from our discussion, since they have a “trivial” tester, which (provided N is big enough) accepts all N -vertex graphs. Furthermore, within the framework of graph partition problems, trivial properties are satisfied by all but finitely many graphs.

Lemma 5.1. *Let Π be any trivial graph partition property. Then, for all sufficiently large N , every N -vertex graph has property Π .*

The proof can be found in Appendix B.

5.2. Some Graph Partition Properties That Are Trivial

We first identify (and discard from the rest of the discussion) a class of graph partition properties that contains only trivial properties.

Claim 5.2. *Let Π be a graph partition property that is testable by a canonical tester with one-sided error and query-complexity independent of N . Suppose that the graph $G = ([N], E)$ has property Π , and let (V_1, \dots, V_k) be a witness partition of G . If for some i the subgraph induced by V_i is neither a clique nor an independent set, then Π is trivial.*

Proof. Let $G = ([N], E)$, (V_1, \dots, V_k) and i be as in the hypothesis. Then the number of edges with both endpoints in V_i is greater than zero and smaller than $\frac{|V_i|^2}{2}$. Suppose for a moment that for some integer t the said number were at least $\frac{t^2}{2}$ and at most $\frac{|V_i|^2}{2} - \frac{t^2}{2}$. Then, for every t -vertex graph H , there would exist a graph G_H having property Π such that G_H would contain H as an induced subgraph. (The graph G_H is derived from G by modifying the subgraph induced by V_i so that the number of edges is maintained and the subgraph induced by the first t vertices of V_i equals H . This is certainly possible, because both the number of edges and non-edges in H is at most $\frac{t^2}{2}$).

Now, suppose that (for some $\varepsilon > 0$) the canonical tester selects a sample of $t = t(\varepsilon)$ vertices. Then for every t -vertex graph H there exists an N -vertex graph G_H satisfying Π (as above) so that when given oracle access to G_H , with positive probability, the tester sees H as the induced subgraph. Since G_H satisfies Π and the tester has one-sided error, it must accept upon seeing H as the induced subgraph. But this holds for every t -vertex graph H , and so the tester must always accept (no matter which induced subgraph it sees). It follows that, for every $N' \geq t$, the tester accepts (with probability $1 > \frac{2}{3}$) any N' -vertex graph, and thus every N' -vertex graph is ε -close to having property Π . If we can repeat the above argument for every $\varepsilon > 0$, then it will follow that Π is trivial. To do so, we must show that, for every t , which will be set to equal the size of the vertex-sample selected by the tester on distance parameter ε , there exist a graph and a k -partition as in the claim's

hypothesis so that the number of edges with both endpoints in V_i is at least $\frac{t^2}{2}$ and at most $\frac{|V_i|^2}{2} - \frac{t^2}{2}$ (rather than just greater than zero and smaller than $\frac{|V_i|^2}{2}$).

Fixing t , we first consider a graph $G = ([N], E)$ with (V_1, \dots, V_k) and i be as in the claim's hypothesis. Let e denote the number of edges with both endpoints in V_i (i.e., $0 < e < \frac{|V_i|^2}{2}$). If $\frac{t^2}{2} < e < \frac{|V_i|^2}{2} - \frac{t^2}{2}$ then we are done. Otherwise we consider the graph $G' = ([t] \times [N], E')$, where

$$E' = \{(\langle i, u \rangle, \langle j, v \rangle) : (u, v) \in E \wedge i, j \in [t]\}. \quad (7)$$

Clearly, the graph G' has property Π , as is witnessed by the partition (V'_1, \dots, V'_k) , where $V'_j \stackrel{\text{def}}{=} \{\langle i, v \rangle : v \in V_j \wedge i \in [t]\}$ (which preserves the relative densities of vertices and edges with respect to the witness partition (V_1, \dots, V_k)). Furthermore, the number of edges in E' with both endpoints in V'_i is at least $\frac{t^2}{2}$ and at most $\frac{|V_i|^2}{2} - \frac{t^2}{2}$. The claim follows. \blacksquare

5.3. Nontrivial Graph Partition Properties—the Two Cases

Below, we refer to the lower and upper bound parameters that appear in the specification of property Π . Recall that there are bounds that refer to the fraction of vertices in each part, and bounds referring to the fraction of edges inside parts or between parts.

Claim 5.3. *Let Π , $G = ([N], E)$ and (V_1, \dots, V_k) be as in Claim 5.2.*

1. *If for some i the subgraph induced by V_i is an independent set, then no edge lower-bound parameter is positive.*
2. *Suppose that no edge lower-bound parameter is positive. Then, if for some i the subgraph induced by V_i is a clique, Π is trivial.*

Proof. We start with Part (1). As in the proof of Claim 5.2, we may assume that the size of V_i is greater than the size of the vertex-sample chosen by the tester. It follows that when given oracle G , with positive probability, the tester will see an induced subgraph that is an independent set. Because the tester has one-sided error, it must accept in this case, and thus it always accepts an oracle N -vertex graph that is an independent set. Repeating the argument for any $\varepsilon > 0$, it follows that, for sufficiently large N , the N -vertex independent set graph is ε -close to having property Π . This contradicts the possibility that some edge lower-bound parameter is positive (i.e., $c > 0$), because it would have meant that (for some constant $c > 0$) the independent set is too far (i.e., c -far) from having property Π .

Turning to Part (2), we first observe that if for some i the subgraph induced by V_i is a clique then (for every $\varepsilon > 0$ and all sufficiently large N) the N -vertex clique graph is ε -close to having property Π . (The proof is similar to the main part of the above argument.) Now, consider a witness partition, denoted (V'_1, \dots, V'_k) , of a graph G' satisfying Π that is ε -close to the N -vertex clique graph, and let j be such that $|V'_j| \geq N/k$. Then, the number of edges with both endpoints in V'_j must be at least $\frac{|V'_j|^2}{2} - \varepsilon N^2 \geq \frac{(N/k)^2}{2}$

– $\varepsilon N^2 > 1$ (where the last inequality holds for $\varepsilon < \frac{1}{2} k^2$ and sufficiently large N). Using the hypothesis that no edge lower-bound parameter is positive, it follows that omitting edges from G' results in a graph that also has property Π . In particular, by possibly omitting a single edge residing in V'_j , we can obtain a graph G'' satisfying Π so that (V'_1, \dots, V'_k) is also a witness partition of G'' and so that the subgraph of G'' induced by V'_j is neither a clique nor an independent set. Using Claim 5.2, Part (2) follows. ■

Corollary 5.4. *Let Π be a graph partition property that is testable by a canonical tester with one-sided error and query-complexity independent of N , and suppose that Π is not trivial. Then exactly one of the following two cases holds:*

1. *Every graph having property Π is k -colorable, and all edge lower-bound parameters in the specification of Π are zero. Furthermore, all upper-bounds referring to edges inside parts must be zero.*
2. *Every graph having property Π can be k -partitioned so that each part is a clique.*

Proof. Using Claim 5.2, the parts in a witness partition of any graph G having property Π must be either cliques or independent sets. Suppose first that for *some* G having property Π , some part of the witness partition of G is an independent set. Then, by Part 1 of Claim 5.3, all edge lower-bounds in the specification of Π are zero. Using Part 2 of Claim 5.3, in this case no part in the witness partition of *any* graph G' having property Π (regardless if $G' = G$ or not) can be a clique, and so (using Claim 5.2 again) all parts in the witness partition of every graph having property Π are independent sets (and so the graph is k -colorable). The only other case allowed for (nontrivial) Π is the one described in Item 2 of the corollary. The main part of the corollary follows.

To finish the proof we show that in the first case (i.e., Π implying k -colorability) all the upper-bounds on the number of edges inside parts must be zero. Suppose on the contrary that the i th upper-bound referring to edges inside parts equals $c > 0$, and consider a witness partition (V_1, \dots, V_k) of a sufficiently large graph $G = ([N], E)$ having property Π . Specifically, we need $\frac{(k+1)^2}{2} \leq c \cdot N^2$ and $|V_i| \geq k + 1$, where the latter can be obtained by amplifying the graph and the witness partition [as in Eq. (7)]. Indeed, each of the V_j 's is an independent set, but we can easily modify G to a graph G' that satisfies Π and yet contains a $(k + 1)$ -clique (in contradiction to the hypothesis that all graphs satisfying Π are k -colorable). This is done by putting a $(k + 1)$ -clique inside V_i , which does not violate the edge density upper-bound of the i th part (and thus guarantees that the modified graph satisfies Π). The furthermore-part of the corollary follows. ■

Below, we consider the two cases of Corollary 5.4. We refer to the second case of Corollary 5.4 (i.e., a graph is k -partitioned so that each part is a clique) as to a graph is *covered by k cliques*.

5.3.1. Graph Partition Properties That Imply k -Colorability

We call an lower-bound (resp., upper-bound) parameter *trivial* if it equals 0 (resp., 1). That is, a trivial bound parameter is satisfied by any k -partition of any graph.

The notion of a relaxation. We say that property Π' is a *relaxation* of Π if every graph satisfying property Π also satisfies property Π' . For $\varepsilon > 0$, we say that property Π' is an ε -*relaxation* of Π if Π' is a relaxation of Π and every sufficiently large graph

satisfying property Π' is ε -close to satisfying Π . We say that Π' is a 0-relaxation of Π if, for every $\varepsilon > 0$, property Π' is an ε -relaxation of Π .

The notion of 0-relaxation is related to the notion of indistinguishability defined by Alon et al. [2]. Note that if Π' is a 0-relaxation of Π , then every tester for Π' is almost a tester for Π in the following sense: For every value of $\varepsilon > 0$ the Π' -tester may behave improperly with respect to Π on finitely many graphs. Thus, as far as property testing is concerned, we may consider 0-relaxations of a property instead the property itself.

We conjecture that, within the framework of graph partition problems, all but finitely many graphs that satisfy a 0-relaxations of a property also satisfy the property itself. This conjecture was already proven for the special case of trivial properties (see Lemma 5.1),¹¹ and is established next for a more general special case that suffices for our needs.

Lemma 5.5. *Let Π and Π' be graph partition properties. Suppose that all edge lower-bound parameters in the specification of both Π and Π' are zero, and that each edge upper-bound parameter in the specification of Π' is either 0 or 1. Further suppose that Π' is a 0-relaxation of Π . Then, for every sufficiently large graph, the graph has property Π if and only if it has property Π' .*

The proof can be found in Appendix C.

Towards a characterization. The main step towards characterizing graph partition properties that imply k -colorability is the following characterization of their 0-relaxations.

Lemma 5.6. *Let Π be as in Corollary 5.4. Suppose that every graph having property Π is k -colorable, and all edge lower-bound parameters in the specification of Π are zero. Then, there exists a graph partition property Π'' that is a 0-relaxation of Π so that the only non-trivial bounds in the specification of Π'' are upper-bounds that equal zero.¹² Furthermore, these zero upper-bounds must include the upper-bounds referring to edges inside each part.*

Proof. As a first step, consider a specification of a property Π' derived from the specification of Π as follows: All edge upper-bounds that equal zero in the specification of Π are set to zero also in the specification of Π' , and all other edge bounds in Π' are trivial. (Recall that all edge lower-bounds in Π are trivial, and so this holds also for Π' .) The vertex bounds of Π are maintained in Π' . (We will deal with them at the second stage.)

Recall that in Π all the upper-bounds referring to edges inside parts must be zero. Thus, all edge bounds of property Π' are as required. Clearly, Π' is a relaxation of Π and so to establish that Π' is a 0-relaxation of Π we need to show that, for every $\varepsilon > 0$, every sufficiently large graph having property Π' is ε -close to have property Π .

Fixing any $\varepsilon > 0$, we consider a sufficiently large N so that the vertex-sample chosen by the tester on distance parameter ε is smaller than the numbers implied by all positive

¹¹The set of all graphs (i.e., the property satisfied by all graphs) is a 0-relaxation of any trivial graph property. Lemma 5.1 can be restated as saying that, for any trivial graph partition property Π and all sufficiently large N , any N -vertex graph (i.e., that is a graph and thus satisfies the 0-relaxation) satisfies property Π .

¹²In case all vertex upper-bounds in Π are positive (see Footnote 13), the only nontrivial bounds in the specification of Π'' are edge upper-bounds (which equal zero).

nontrivial edge upper-bounds of Π . That is, if $c > 0$ be the smallest positive edge upper-bound of Π , then we set $N > t(\varepsilon)/\sqrt{c}$, where $t(\varepsilon)$ is the size of the sample chosen by the tester.

Let $G' = ([N], E')$ be an arbitrary N -vertex graph satisfying Π' . Consider a vertex sample, denoted S , taken by the tester (for Π) on distance parameter ε and access to the oracle G' . We first show that the (small) subgraph of G' induced by S can be embedded in a graph G that satisfies Π , where G is derived from G' by omitting almost all edges. Specifically, we consider the graph $G = ([N], E)$ obtained by letting $E = \{(u, v) \in E' \cap (S \times S)\}$. The only bounds of Π that can be violated by a graph having property Π' are positive (nontrivial) edge upper-bounds, because all other bounds of Π equals those of Π' . But these bounds cannot be violated by G , because G has very few edges (i.e., G has less than $|S|^2$ edges, and N was chosen so that $|S|^2 < c \cdot N^2$). It follows that any induced subgraph that can be seen by the test (for Π), when given access to the oracle G' , is also seen by the test with positive probability when given access to some oracle representing a graph that has property Π . Using the one-sided error feature of the test (for Π), it follows that the test accepts G' with probability 1, and hence G' must be ε -close to having property Π . This concludes the proof that Π' is a 0-relaxation of Π .

We now turn to the next step: Starting from Π' , we obtain Π'' by possibly modifying the vertex bounds, and leaving all edge bounds intact. Specifically, we set all nonzero vertex-bounds to be trivial (i.e., 0 for lower-bounds and 1 for upper-bounds), and maintain zero vertex upper-bounds and lower-bounds (in case they are present in Π and Π').¹³ Using an argument as in the first stage, it follows that Π'' is a 0-relaxation of Π' . Specifically, Π'' is a relaxation of Π' , and for every $\varepsilon > 0$ and sufficiently large graph satisfying Π'' , it is the case that the very same graph is ε -close to Π' . [Intuitively, looking at the witness partition w.r.t. Π'' , observe that the only bounds of Π' that can be violated by that partition are nonzero (vertex) bounds of Π' ; but this cannot be detected with one-sided error from an $o(N)$ -size vertex-sample.] The lemma follows. ■

5.3.2. Graph Partition Properties That Imply a Cover by k Cliques

Lemma 5.7. *Let Π be as in Corollary 5.4. Suppose that every graph having property Π can be k -partitioned so that each part is a clique. Then, for sufficiently large N , an N -vertex graph has property Π if and only if it is an N -vertex clique.*

(Recall that we consider only properties that are satisfied by some graphs.)

Proof. Let $G = ([N], E)$ be an arbitrary graph having property Π . As shown in the proof of Claim 5.2 (see also below), the graph G can be assumed to be large enough so that some part in its witness partition is larger than the vertex-sample taken by the tester (on distance parameter ε). Let us denote the size of that sample by $t = t(\varepsilon)$. Since the tester has one-sided error (and G contains a t -vertex clique), the tester must accept when the subgraph induced by the vertex-sample is a t -vertex clique. It follows that the N -vertex clique is ε -close to Π . Below we shall show that no other N -vertex graph can be accepted

¹³Recall that zero lower-bounds are trivial, whereas zero upper-bounds on vertex-density are nontrivial but actually redundant (because they merely mean that we specify a k' -partition problem, for some $k' < k$, rather than a k -partition problem).

by the tester. One consequence of this is that the N -vertex clique must have property Π (because, otherwise no N -vertex graph has property Π).

Suppose, towards the contradiction, that $G = ([N], E)$ has property Π but is not the N -vertex clique. We consider an amplified version of G , denoted $G' = ([t] \times [N], E')$, where E' is as in Eq. (7). Then, on one hand G' has property Π (with witness partition induced by that of G).¹⁴ On the other hand, G' contains as an induced subgraph a $2t$ -vertex graph consisting of a pair of t -cliques (corresponding to any missing edge in G). However, for every $i = 1, \dots, t - 1$, with positive probability the tester given oracle access to G' sees an induced subgraph consisting of two cliques, one of size i and the other of size $t - i$ (and no additional edges). Since the tester has one-sided error (and G' has property Π), the tester must accept in each of these cases. It follows that when given oracle access to any N -vertex graph consisting of two cliques (and no additional edges), the tester will always accept. Thus, subject to the contradiction hypothesis,¹⁵ we have:

Claim 1: For all sufficiently large N and every $1 < M < N$, every graph that consists of an M -vertex clique and an $(N - M)$ -vertex clique (and no additional edges) is ε -close to having property Π .

Next, we consider the edge lower-bounds in the specification of Π ; that is, let $\mathbb{1}_{i,i}$ (resp., $\mathbb{1}_{i,j}$) denote the lower-bound referring to edge density within the i th part (resp., between the i th and the j th parts). Specifically, these lower-bounds require that the number of edges within the i th part is at least $\mathbb{1}_{i,i} \cdot N^2$ (resp., between the i th and the j th parts is at least $\mathbb{1}_{i,j} \cdot N^2$).

Motivation: For simplicity we consider the case $k = 2$ (observing that the lemma is trivial in case $k = 1$). Furthermore, for simplicity, we first assume that $\varepsilon = 0$. Using Claim 1, it follows that the graph consisting of two $(N/2)$ -vertex cliques (and no additional edges) has property Π . The only witness partition possible for this graph is the one in which each clique is in a different part, and thus each part is allowed to have at least one half of the number of vertices. Furthermore, both $\mathbb{1}_{1,1} \leq \frac{1}{8}$ and $\mathbb{1}_{2,2} \leq \frac{1}{8}$ (since each part contains a $(N/2)$ -vertex clique, and so has only $\frac{(N/2)^2}{2} = \frac{N^2}{8}$ edges). Next, we consider the graph consisting of one $(N/3)$ -vertex clique and one $(2N/3)$ -vertex clique, which (by Claim 1) also has property Π . Again, the witness partition of this graph has each clique in a different part, and it follows that either $\mathbb{1}_{1,1} \leq \frac{1}{18}$ or $\mathbb{1}_{2,2} \leq \frac{1}{18}$ (since the part containing the $(N/3)$ -vertex clique has only $\frac{(N/3)^2}{2} = \frac{N^2}{18}$ edges). Suppose, without loss of generality, that $\mathbb{1}_{2,2} \leq \frac{1}{18}$. But now it follows that also the graph consisting of one $(N/2)$ -vertex clique, one $(N/3)$ -vertex clique, and $N/6$ isolated vertices satisfies Π [e.g., consider the witness partition in which the $(N/2)$ -vertex clique is in one side and the rest of the graph is in the other]. This contradicts the lemma's hypothesis by which every graph having property Π can be covered by k cliques, and so cannot have a large independent set. Thus, the contradiction hypothesis (by which there exists a graph $G = ([N], E)$ that has property Π but is not the N -vertex clique) must be false. The analysis is easily extended to small $\varepsilon > 0$, but extending it to arbitrary $k > 2$ is more involved.

¹⁴Observe that the parts of the induced partition are cliques, and that the fraction of edges between the parts is exactly as in the witness partition of G .

¹⁵That is, assuming that $G = ([N], E)$ has property Π but is not the N -vertex clique.

The heart of the actual analysis (for arbitrary $k \geq 2$ and $\varepsilon > 0$) is stated and proven next:

Claim 2: Subject to the contradiction hypothesis (see footnote 15) there must be a set $C \subset [k]$ such that

$$\sum_{i \leq j \in C} 1_{i,j} \leq 3k^2\varepsilon. \quad (8)$$

Furthermore, for all sufficiently large N , there exists an N -vertex graph H' satisfying Π and a witness partition (V'_1, \dots, V'_k) of H' such that

$$\sum_{i \in C} |V'_i| \geq k\sqrt{\varepsilon} N. \quad (9)$$

Proof: For N as in Claim 1, consider an N -vertex graph H that consists of a pair of cliques (and no additional edges), where the smaller clique is of size $2k\sqrt{\varepsilon} \cdot N$. Since (By Claim 1) H is ε -close to having property Π , we may consider a witness partition (V'_1, \dots, V'_k) of a graph H' that satisfies Π and is ε -close to H . Since each V'_i is a clique in H' , the subgraph of H induced by V'_i misses at most εN^2 (because H is ε -close to H'). Thus, each V'_i is “dominated by one of the cliques of H ” in the sense that either it contains at most $\sqrt{\varepsilon} N$ vertices of the small clique (of H) or it contains at most $\sqrt{\varepsilon} N$ vertices of the large clique [since otherwise the subgraph of H induced by V'_i misses more than $(\sqrt{\varepsilon} N)^2$ edges]. It follows that there exist $C \subset [t]$ such that the parts with index in C contain at most $k\sqrt{\varepsilon} N$ vertices of the large clique and all but at most $k\sqrt{\varepsilon} N$ vertices of the small clique (e.g., C contain the indices of all parts that each have at most $\sqrt{\varepsilon} N$ vertices of the large clique, and so $\bar{C} \stackrel{\text{def}}{=} [t] \setminus C$ (which contains only parts with more than $\sqrt{\varepsilon} N$ vertices of the large clique) contains parts that each have at most $\sqrt{\varepsilon} N$ vertices of the small clique). In particular, the number of vertices residing in parts with an index in C is at least $2k\sqrt{\varepsilon} \cdot N - k\sqrt{\varepsilon} \cdot N = k\sqrt{\varepsilon} \cdot N$, and Eq. (9) follows.

Turning to Eq. (8), we observe that the number of edges (in H) having both endpoints residing in parts having an index in C is at most $\frac{1}{2} \cdot ((2k\sqrt{\varepsilon} \cdot N)^2 + (k\sqrt{\varepsilon} \cdot N)^2) = \frac{5k^2\varepsilon}{2} \cdot N^2$. It follows that the corresponding number in H' is at most $\frac{5k^2\varepsilon}{2} \cdot N^2 + \varepsilon N^2 < 3k^2\varepsilon N^2$, where $\frac{k^2}{2} > 1$ is due to $k \geq 2$. Since H' satisfies Π (and in particular its edge lower-bounds), Eq. (8) follows. \blacksquare

We stress that the above two claims holds for any value of $\varepsilon > 0$, subject to the contradiction hypothesis (see Footnote 15). Thus, for every sequence $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_{2^k+1}$ of positive numbers. Claim 2 holds for some set C and two different ε_j 's. That is, there exists a set $C \subset [k]$ such that for some $p < q$ and all sufficiently large N , there exist an N -vertex graph H' satisfying Π and a witness partition (V'_1, \dots, V'_k) such that

$$\sum_{i \leq j \in C} 1_{i,j} \leq 3k^2\varepsilon_q, \quad (10)$$

$$\sum_{i \in C} |V'_i| \geq k\sqrt{\varepsilon_p} N. \quad (11)$$

It follows that there exists $i \in C$ such that $|V'_i| \geq \sqrt{\varepsilon_p} N$ and $1_{i,i} \leq 3k^2 \varepsilon_q$. Selecting the sequence of ε_j 's so that $\varepsilon_j = (7k^2)^{-j}$, we have $\varepsilon_q \leq \varepsilon_{p+1} = \varepsilon_p / (7k^2)$, and so

$$\begin{aligned} \frac{|V'_i|^2}{2} &\geq \frac{(\sqrt{\varepsilon_p} N)^2}{2} = \frac{\varepsilon_p N^2}{2} \geq \frac{7k^2 \varepsilon_q}{2} \cdot N^2 > 3k^2 \varepsilon_q N^2 + \frac{(k+1)^2}{2} \\ &\Rightarrow \frac{|V'_i|^2}{2} > 1_{i,i} \cdot N^2 + \frac{(k+1)^2}{2}. \quad (12) \end{aligned}$$

Equation (12) allows us to modify H' so that the resulting graph also satisfies Π but has an independent set of size $k+1$, which contradicts the (contradiction) hypothesis that every graph having property Π can be k -partitioned into cliques. Specifically, recall that H' has property Π and that (V'_1, \dots, V'_k) is a witness partition. In particular, the subgraph induced by V'_i is a clique, but we can omit $\frac{(k+1)^2}{2}$ from it without violating any of the bounds of Π [i.e., the only relevant bound is the lower bound on the number of edges inside V'_i , but Eq. (12) asserts that this bound will continue to hold even if we omit $\frac{(k+1)^2}{2}$ edges with both endpoints in V'_i]. This results in a graph H'' that still satisfies Π (and has an independent set of size $k+1$). Thus, the graph H'' violates the lemma's hypothesis (that graphs satisfying Π can be k -partitioned into cliques). It follows that the contradiction hypothesis (by which there exists a graph $G = ([N], E)$ that has property Π but is not the N -vertex clique) must be false. The lemma follows. \blacksquare

5.4. The Characterization Theorem and a Corollary

Combining Corollary 5.4, and Lemmas 5.5–5.7, we obtain:

Theorem 5.8 (Theorem 3, restated). *Let Π be a graph partition property that is testable with one-sided error and query-complexity independent of N , and suppose that Π is not trivial. Then exactly one of the following two cases holds:*

1. *There exists a k -vertex graph H so that for all sufficiently large graphs G , the graph G satisfies Π if and only if its vertices can be k -partitioned such that there are no edges among vertices residing in the same part and so that there are edges between vertices of the i -th part and j -th part only if (i, j) is an edge of H .*
2. *For sufficiently large N , an N -vertex graph has property Π if and only if it is an N -vertex clique.*

Proof. By Theorem 4.5, we may assume that the tester is canonical, and apply Corollary 5.4. Assuming that Π is as in Case 1 of Corollary 5.4, we apply Lemma 5.6 and conclude that Π has a 0-relaxation Π' such that the only nontrivial bounds in the specification of Π' are upper-bounds that equal zero, and that all graphs satisfying Π' are k -colorable. Thus, we are allowed to apply Lemma 5.5 to this pair (Π, Π') and conclude that every sufficiently large graph having property Π is k -colorable, and (without loss of generality) the only nontrivial bounds in the specification of Π are upper-bounds that equal zero. Defining H so that $(i, j) \in [k] \times [k]$ is an edge if and only if the upper-bound referring to edges between the i th and j th part is trivial, we obtain the condition of Case 1 of the current theorem.

The only other possibility is that Π is as in Case 2 of Corollary 5.4. Applying Lemma 5.7, we obtain the condition of Case 2 of the current theorem. ■

The property checked by the canonical tester. We are now ready to address the question posed at Section 4.2.4. Using Theorem 5.8, we show that, in the context of graph partition problems, if a property is testable by a one-sided error tester of complexity that only depends on ε , then it can be so tested by a canonical tester that accepts iff the induced subgraph has the very same property (rather than some other graph property).

Corollary 5.9. *Let Π be a graph partition property that is testable with one-sided error and query-complexity independent of N . Then Π can be tested with one-sided error by checking whether or not a random $\text{poly}(1/\varepsilon)$ -vertex induced subgraph has the property Π .*

*Proof.*¹⁶ The conclusion holds vacuously for trivial properties. Thus, using Theorem 5.8, we need to consider only two cases (corresponding to the two items). In both cases, the property Π is closed under taking induced subgraphs (i.e., if G has property π then every induced subgraph of G has property Π), and is testable by $\text{poly}(1/\varepsilon)$ queries (cf. [8, Theorem 9.2]). By using Proposition D.2, the corollary follows. ■

APPENDIX A: THE FRAMEWORK OF GRAPH PARTITION PROBLEMS

A graph partition property is defined by a sequence of pairs of nonnegative numbers. For some integer k , we have $k + k + \binom{k}{2}$ pairs providing upper and lower bounds on the fraction of vertices in each part of the k -partition as well as on the fraction of edges within parts and between parts. Specifically, consider the sequence of pairs

$$(\ell_1, u_1), \dots, (\ell_k, u_k), (\ell_{1,1}, u_{1,1}), \dots, (\ell_{k,k}, u_{k,k}), ((\ell_{i,j}, u_{i,j}))_{1 \leq i < j \leq k}. \quad (13)$$

This sequence corresponds to a graph property that is satisfied by all graphs $G = ([N], E)$ having a k -partition (V_1, \dots, V_k) such that the following two conditions hold:

$$\begin{aligned} \ell_i \cdot N &\leq |V_i| \leq u_i \cdot N \quad \forall i, \\ \ell_{i,j} \cdot N^2 &\leq |E \cap (V_i \times V_j)| \leq u_{i,j} \cdot N^2 \quad \forall i < j. \end{aligned}$$

That is, ℓ_i (resp., u_i) is a lower bound (resp., upper bound) on the fraction of vertices in the i th part, $\ell_{i,i}$ (resp., $u_{i,i}$) is a lower bound (resp., upper bound) on the fraction of edges having both endpoints in the i th part, and $\ell_{i,j}$ (resp., $u_{i,j}$) is a lower bound (resp., upper bound) on the fraction of edges crossing between the i th part and the j th part, for $i < j$.

Certainly, some sequences of parameters give rise to graph partition properties that are not satisfied by any graph. We discard these cases from our discussion (calling them

¹⁶An alternative proof, which does not use Proposition D.2, can be found in our Technical Report [9]. Both proofs use the same underlying principles, which can be traced back to the proof of [8, Corollary 7.2].

nonadmissible; see [8, Definition 9.3.1]). In particular, we will consider only sequences as in Eq. (13) satisfying $0 \leq 1_i \leq u_i \leq 1$ and $0 \leq 1_{i,j} \leq u_{i,j} \leq 1$, for all i, j .

A technicality: integrality problems. Following [8, Remark 9.1], we avoid integrality problems by allowing up to $k - 1$ vertices to be split between the k parts of the partition (and count these fractional vertices and edges in the natural way). Had we not followed this convention, the set of N -vertex graphs satisfying a graph partition property could be empty for some values of N and non-empty for others.

A technicality: counting edges and self-loops. Note that the edge bounds impose bounds in terms of multiples of N^2 [rather than of $\binom{N}{2}$, which may be more natural]. This convention is adopted for greater expressibility. For example, using this convention, a full bipartite graph with $N/2$ vertices on each side has edge density $\frac{(N/2)^2}{N^2} = \frac{1}{4}$, which is independent of N . (In contrast, if we consider multiples of $\binom{N}{2}$, then such a graph will have edge density $\frac{(N/2)^2}{\binom{N}{2}}$, which is not independent of N , and consequently the corresponding condition could not have been expressed as a graph partition problem.) For similar reasons, we allow self-loops and count them as half edges.¹⁷ This way, a ρN -vertex clique in an N -vertex graph has edge density $\frac{\binom{\rho N}{2} + \frac{\rho N}{2}}{N^2} = \frac{\rho^2}{2}$, which is independent of N . Consequently, the property of having such a clique can be expressed as a graph partition problem.

Indeed, using multiples of $\frac{N^2}{2}$ rather than N^2 would be more natural, but both more cumbersome and less in agreement with the presentation in [8, Section 9]. Clearly, the last choice is immaterial.

APPENDIX B: PROOF OF LEMMA 5.1

Recall that by the lemma's hypothesis Π is a trivial graph partition property. Our aim is to show that, for all sufficiently large N , every N -vertex graph has property Π . This follows by combining the following three claims.

Claim B.1. *The specification of Π does not contain any positive lower-bound regarding edges.*

Proof. Suppose on the contrary that Π contains a positive lower-bound regarding edges. Then, for some $c > 0$, each N -vertex graph satisfying Π must have at least $c \cdot N^2$ edges, which contradicts the hypothesis that Π is trivial (because, in this case, it cannot hold that for every $\varepsilon > 0$ and all sufficiently large N , the N -vertex independent set is ε -close to having property Π). ■

¹⁷Unfortunately, the text of [8, Section 9] is unclear regarding this aspect, which is essential for the claim that the ρ -clique problem can be expressed as a graph partition problem. Furthermore, for simplicity, in [8, Section 5] self-loops are disallowed. We stress that all the results of [8] are preserved if one allows self-loops (which can be ignored by all algorithms). Finally, we note that counting self-loops as half edges is consistent with [8], where each (non-self-loop) edge is counted twice, in correspondence to its two occurrences in the adjacency matrix of the graph. (Using this correspondence justifies counting self-loops once, which is half the count relative to edges that are not self-loops.) In this paper we chose to get rid of the annoying convention of counting each (non-self-loop) edge twice, and the result is the annoying convention by which a self-loop is counted half a time.

Claim B.2. *For all sufficiently large N , the N -vertex clique has property Π .*

The proof of Claim B.2 is postponed to the end of this section.

Claim B.3. *Let Ξ be a graph partition property with a specification that does not contain any positive lower-bound regarding edges. If the N -vertex clique has property Ξ , then all N -vertex graphs have property Ξ .*

Proof. Consider a witness partition of the N -vertex clique, denoted C_N . Then this partition is also a witness partition of any N -vertex subgraph of C_N , because all edge lower-bounds are nonpositive (i.e., zero). The claim follows. ■

Proof of Claim B.2. We consider the sequence of bounds in the specification of property Π , and refer to the notation in Eq. (13). For these bounds (i.e., $\mathbb{1}_i, u_i, \mathbb{1}_{i,j}, u_{i,j}$'s), consider the following system of equations in variables x_1, \dots, x_k :

$$\sum_{i=1}^k x_i = 1, \quad (14)$$

$$\mathbb{1}_i \leq x_i \leq u_i \quad \forall i, \quad (15)$$

$$\frac{1}{2} \cdot x_i^2 \leq u_{i,i} \quad \forall i, \quad (16)$$

$$x_i \cdot x_j \leq u_{i,j} \quad \forall i < j. \quad (17)$$

We first claim that the above system has a solution. Otherwise, there exists a constant $\varepsilon > 0$ so that any solution satisfying Eqs. (14) and (15) violates one of the other equations by at least ε (i.e., either $\frac{1}{2} \cdot x_i^2 \geq u_{i,i} + \varepsilon$ or $x_i \cdot x_j \geq u_{i,j} + \varepsilon$ for some i, j). We will show that the latter hypothesis contradicts the hypothesis that, for sufficiently large N , the N -vertex clique is $(\varepsilon/2)$ -close to having property Π : Let (V_1, \dots, V_n) be a witness partition of a graph G having property Π and being $(\varepsilon/2)$ -close to the N -vertex clique, and set $x_i = |V_i|/N$ for $i = 1, \dots, k$. Then this setting satisfies Eqs. (14) and (15). Also, for every i , we must have $\frac{|V_i|^2}{2} - \frac{\varepsilon}{2} \cdot N^2 \leq u_{i,i} \cdot N^2$ (because, on one hand, the witness partition of G respects all bounds of Π , and on the other hand G may miss at most $\frac{\varepsilon}{2} \cdot N^2$ edges [because it is $(\varepsilon/2)$ -close to the N -vertex clique]). Thus, we have $\frac{|V_i|^2}{2} < \varepsilon \cdot N^2 + u_{i,i} \cdot N^2$, and $\frac{1}{2} \cdot x_i^2 < u_{i,i} + \varepsilon$ follows. Similarly, $|V_i| \cdot |V_j| - \frac{\varepsilon}{2} \cdot N^2 \leq u_{i,j} \cdot N^2$ and $x_i \cdot x_j < u_{i,j} + \varepsilon$ follows for every $i < j$. This contradicts the above hypothesis that any solution satisfying Eqs. (14) and (15) violates one of the other equations by at least ε , and we conclude that the system of Eqs. (14)–(17) has a solution.

Let (x_1, \dots, x_k) be a solution to Eqs. (14)–(17). Then avoiding integrality problems (see [8, Remark 9.1]), we consider for each N a partition (V_1, \dots, V_k) of the N -vertex clique so that $|V_i| = x_i N$. Clearly, this partition satisfies all vertex bounds [because these correspond to Eq. (15)]. By Claim B.1, the only remaining nontrivial bounds are the edge upper-bounds (because all $\mathbb{1}_{i,j}$'s are zero). But these are shown to be satisfied as follows:

For every i , the number of edges with both endpoints in V_i is $\frac{|V_i|^2}{2}$, and we have $\frac{|V_i|^2}{2} = \frac{x_i^2}{2} N^2 \leq u_{i,i} \cdot N^2$, where the last inequality is due to Eq. (16). Similarly, for every $i < j$, the number of edges crossing between V_i and V_j is $|V_i| \cdot |V_j|$, and we have $|V_i| \cdot |V_j| = x_i x_j \cdot N^2 \leq u_{i,i} \cdot N^2$, where the inequality is due to Eq. (17). The claim follows. \blacksquare

APPENDIX C: PROOF OF LEMMA 5.5

Recall the hypotheses of the lemma:

1. Π is a graph partition property such that all edge lower-bound parameters in the specification of Π are zero.
2. Π' is a graph partition property such that all edge lower-bound parameters in the specification of Π' are zero, and each edge upper-bound parameter in the specification of Π' is either zero or one.
3. Π' is a 0-relaxation of Π .

Our aim is to show that, for every sufficiently large graph, the graph has property Π if and only if it has property Π' . (In fact, we only need to show the “if”-direction.) Our proof generalizes the proof of Lemma 5.1 (given in Appendix B). In particular, we do not need to establish an analogue of Claim B.1, because this is already guaranteed in our first hypothesis. The role of the N -vertex clique will be played by each member of a family of certain extremal (dense) graphs [and the type of the upper-bounds in Π' seem important to allow us to focus on this family (and thus perform this extension)]. We thus prove analogues of the two other claims of Appendix B: Analogously to Claim B.2, we first show that each member of the extremal family that has property Π' also has property Π . Then, analogously to Claim B.3, we show that this extends to each subgraph of the extremal graphs, and that the latter subgraphs are all the graphs that may have property Π' . The lemma will follow.

Pivotal to the above plan, is the definition of *extremal graphs* for Π' . Since the edge lower-bounds of Π' are trivial and each edge upper-bound in Π' is either 0 or 1, the extremal graphs are determined by the density of the vertex sets allowed by the vertex bounds (and the type of the edge upper-bounds). That is, using notations as in Eq. (13), let $l'_i, u'_i, l'_{i,j}, u'_{i,j}$ be the bounds in the specification of Π' . (Recall that all $l'_{i,j}$ equal zero and each $u'_{i,j}$ equals either zero or one.) Then a sequence of vertex-set densities, denoted $\bar{\rho} = (\rho_1, \dots, \rho_k)$, is **permitted** by Π' if $\sum_{i=1}^k \rho_i = 1$ and $l'_i \leq \rho_i \leq u'_i$ for all $i = 1, \dots, k$. For each such permitted sequence $\bar{\rho}$ and every N , we consider the **extremal graph** $G^{(N, \bar{\rho})} = ([N], E^{(N, \bar{\rho})})$ defined by

$$V_i^{(N, \bar{\rho})} \stackrel{\text{def}}{=} \left\{ v \in [N] : \sum_{j=1}^{i-1} \rho_j N < v \leq \sum_{j=1}^i \rho_j N \right\}, \quad (18)$$

$$F \stackrel{\text{def}}{=} \{(i, j) : 1 \leq i \leq j \leq k \wedge u'_{i,j} = 1\}, \quad (19)$$

$$E^{(N, \bar{\rho})} \stackrel{\text{def}}{=} \bigcup_{(i,j) \in F} (V_i^{(N, \bar{\rho})} \times V_j^{(N, \bar{\rho})}). \quad (20)$$

That is, Eq. (18) specifies a (canonical) k -partition that satisfies the vertex bounds of Π' (i.e., is according to the permitted sequence $\bar{\rho}$), Eq. (19) indicates the part-pairs among which edges are allowed, and Eq. (20) mandates all possible edges among each allowed pair of parts. Since all edge lower-bounds in Π' are zero (and thus trivial), they are satisfied by the above k -partition of $G^{(N,\bar{\rho})}$. Since each edge upper-bound in Π' is either 0 or 1, it is also satisfied that partition (because there are edges between the i th and j th part iff $u'_{i,j} = 1$). Thus, $G^{(N,\bar{\rho})}$ satisfies Π' , and $(V_1^{(N,\bar{\rho})}, \dots, V_k^{(N,\bar{\rho})})$ is a witness partition. [Note that we do not rule out the possibility that F contains pairs of the form (i, i) , although the text and notation may suggest otherwise. We comment that in our application of Lemma 5.5, $u'_{i,i} = 0$ and so $(i, i) \notin F$ for all i .]

Recall that each extremal graph $G^{(N,\bar{\rho})}$ satisfies Π' and that $(V_1^{(N,\bar{\rho})}, \dots, V_k^{(N,\bar{\rho})})$ is a witness partition. Next, in Claim C.1, we show that such extremal graph also satisfies Π . Later (see Claim C.2) we show that each N -vertex graph satisfying Π' is a subgraph of an N -vertex (graph that is isomorphic to an) extremal graph, and that this subgraph also satisfies Π . The latter claim (i.e., Claim C.2) is much easier.

Claim C.1. *Every extremal graph satisfies property Π .*

Proof. Let $G^{(N,\bar{\rho})}$ be an extremal graph (for Π'), where $\bar{\rho} = (\rho_1, \dots, \rho_k)$ is permitted (by Π'). We consider the sequence of bounds in the specification of property Π , and denote these bounds by $\perp_i, u_i, \perp_{i,j}, u_{i,j}$'s. Recall that by our first hypothesis, all $\perp'_{i,j}$'s equal zero. [Typically, the $\perp_i, u_i, \perp_{i,j}, u_{i,j}$'s are (possibly) more stringent than the corresponding $\perp'_i, u'_i, \perp'_{i,j}, u'_{i,j}$'s. We stress that the following proof does not refer at all to the bounds $\perp'_i, u'_i, \perp'_{i,j}, u'_{i,j}$.]

Consider the following system of equations in the variables $x_{1,1}, \dots, x_{k,k}$:

$$\sum_{j=1}^k x_{ij} = \rho_i \quad \forall i, \quad (21)$$

$$\perp_j \leq \sum_{i=1}^k x_{ij} \leq u_j \quad \forall j, \quad (22)$$

$$\frac{1}{2} \cdot \sum_{(i,i) \in F} x_{ij}^2 + \sum_{(i,i') \in F \setminus \{(i,i) : i \in [k]\}} x_{ij} x_{i'j} \leq u_{jj} \quad \forall j, \quad (23)$$

$$\sum_{(i,i') \in F} x_{ij} x_{i'j'} \leq u_{jj'} \quad \forall j < j'. \quad (24)$$

We first claim that the above system has a solution. Otherwise, there exists a constant $\varepsilon > 0$ so that any solution satisfying Eqs. (21) and (22) violates one of the other equations by at least ε . We will show that this contradicts the hypothesis that, for sufficiently large N , the graph $G^{(N,\bar{\rho})}$ (which satisfies property Π') is $(\varepsilon/2)$ -close to having property Π : Let (V_1, \dots, V_n) be a witness partition of a graph G having property Π and being $(\varepsilon/2)$ -close to the graph $G^{(N,\bar{\rho})}$, and let $V_{i,j} = V_i^{(N,\bar{\rho})} \cap V_j$ for $i, j \in [k]$. Now, set $x_{i,j} = |V_{i,j}|/N$

for $i, j \in [k]$. This setting satisfies Eq. (21), because $\cup_j(V_i^{(N, \bar{\rho})} \cap V_j) = V_i^{(N, \bar{\rho})}$ and $|V_i^{(N, \bar{\rho})}| = \rho_i N$. It also satisfies Eq. (22), because $\cup_i(V_i^{(N, \bar{\rho})} \cap V_j) = V_j$ and $\perp_j N \leq |V_j| \leq \cup_j N$. We next consider the number of edges in the subgraph of G induced by V_j . Since G may miss at most $\frac{\varepsilon}{2} \cdot N^2$ edges of the extremal graph, the number of edges in that induced subgraph is at least

$$\sum_{(i,i) \in F} \frac{|V_{ij}|^2}{2} + \sum_{(i,i') \in F \setminus \{(i,i): i \in [k]\}} |V_{ij}| \cdot |V_{i'j}| - \frac{\varepsilon}{2} \cdot N^2.$$

But this number must be at most $\cup_{j,j} N^2$ (because the witness partition of G respects all bounds of Π). Thus, for every j ,

$$\frac{1}{2} \cdot \sum_{(i,i) \in F} x_{ij}^2 + \sum_{(i,i') \in F \setminus \{(i,i): i \in [k]\}} x_{ij} x_{i'j} \leq \cup_{jj} + \frac{\varepsilon}{2} < \cup_{jj} + \varepsilon$$

Similarly, for $j < j'$, the number of edges (in G) having one endpoint in V_j and the other endpoint in $V_{j'}$ is at least

$$\sum_{(i,i') \in F} |V_{ij}| \cdot |V_{i'j'}| - \frac{\varepsilon}{2} \cdot N^2,$$

and it follows that $\sum_{(i,i') \in F} x_{ij} x_{i'j'} \leq \cup_{j,j'} + \frac{\varepsilon}{2}$. This contradicts the above hypothesis that any solution satisfying Eqs. (21) and (22) violates one of the other equations by at least ε , and we conclude that the system of Eqs. (21)–(24) has a solution.

Let $(x_{1,1}, \dots, x_{k,k})$ be a solution to Eqs. (21)–(24). Then avoiding integrality problems, we consider for each N a partition (V_1, \dots, V_k) of the extremal graph $G^{(N, \bar{\rho})}$ so that $V_{i,j}$ is a partition of $V_i^{(N, \bar{\rho})}$ satisfying $|V_{i,j}| = x_{i,j} N$, and $V_j = \cup_i V_{i,j}$. Clearly, this partition satisfies all vertex bounds [because these correspond to Eq. (22)]. By the first hypothesis, the only remaining nontrivial bounds are the edge upper-bounds (because all $\perp_{i,j}$'s are zero). But these are shown to be satisfied as follows: For every j , the number of edges with both endpoints in V_j equals

$$\begin{aligned} \sum_{(i,i) \in F} \frac{|V_{ij}|^2}{2} + \sum_{(i,i') \in F \setminus \{(i,i): i \in [k]\}} |V_{ij}| \cdot |V_{i'j}| \\ = \frac{1}{2} \sum_{(i,i) \in F} x_{ij}^2 \cdot N^2 + \sum_{(i,i') \in F \setminus \{(i,i): i \in [k]\}} x_{ij} \cdot x_{i'j} \cdot N^2 \leq \cup_{jj} \cdot N^2, \end{aligned}$$

where the last inequality is due to Eq. (23). Similarly, for every $j < j'$, the number of edges crossing from V_j to $V_{j'}$ equals

$$\sum_{(i,i') \in F} |V_{ij}| \cdot |V_{i'j'}| = \sum_{(i,i') \in F} x_{ij} \cdot x_{i'j'} \cdot N^2 \leq \cup_{j,j'} \cdot N^2,$$

where the inequality is due to Eq. (24). The claim follows. \blacksquare

We call a graph *iso-extremal* for Π' if it is isomorphic to an extremal graph for Π' . In the following claim we rely on the hypothesis that Π does not contain positive lower-bound regarding edges.

Claim C.2 (On subgraphs of graphs that are iso-extremal for Π').

1. Every N -vertex graph that satisfies Π' is a subgraph of an N -vertex graph that is iso-extremal for Π' .
2. Every N -vertex graph that is a subgraph of an N -vertex graph that is iso-extremal for Π' , satisfies Π .

Combining the two parts of the above claim, Lemma 5.5 follows.

Proof. Let $G = ([N], E)$ be a graph satisfying Π' , and let (V_1, \dots, V_n) be a witness partition. Then the sequence (ρ_1, \dots, ρ_k) , where $\rho_i = |V_i|/N$ is permitted by Π' . Furthermore, there are edges in the subgraph induced by V_i only if $u'_{i,i} > 0$ (which holds iff $u'_{i,i} = 1$). Similarly, there are edges between V_i and V_j only if $u'_{i,j} = 1$. Part 1 follows by using any isomorphism that maps V_i to $V_i^{(N, \bar{\rho})}$.

Turning to Part 2, we first observe that every graph that is iso-extremal for Π' satisfies Π [because, by Claim C.1, every graph that is extremal for Π' satisfies Π , and (being a graph property) Π is preserved under isomorphism]. Since all edge lower-bounds in Π are zero, omitting edges does not violate Π . The claim follows. ■

APPENDIX D: ON PROPERTIES CLOSED UNDER INDUCED SUBGRAPHS

In continuation to the discussion in Section 4.2.4, we present the following result of Alon [1].

Definition D.1. A graph property Π is said to be closed under taking induced subgraphs if, for every graph G having property Π , it holds that every induced subgraph of G has property Π .

For example, k -colorability is closed under taking induced subgraphs, whereas connectivity is not.

Proposition D.2. Let Π be a graph property that is closed under taking induced subgraphs, and supposed that Π is testable by $q(\varepsilon)$ queries, independent of the size of the graph. Then Π is testable by inspecting a random induced subgraph of size $O(q(\varepsilon))$ and accepting if and only if the said subgraph has property Π .

We stress that the tester referred to in the second part of the hypothesis does not necessarily have one-sided error, whereas the tester derived in the conclusion has one-sided error.

Proof. By Theorem 2, Π is testable by inspecting a random induced subgraph of size $s(\varepsilon) \stackrel{\text{def}}{=} O(q(\varepsilon))$ (and accepting if and only if the said subgraph has a property Π' , where

Π' some graph property). Let us denote by T the tester obtained this way, and assume¹⁸ that its error probability is bounded by $\frac{1}{4}$. We claim that an algorithm, A , that inspects a random induced subgraph of size $s(\varepsilon)$ and accepts iff the said subgraph has a property Π is a valid tester for Π .

Clearly, if G has property Π then A accepts it with probability 1 (because all induced subgraphs have property Π). We now turn to the case that G is ε -far from Π . Our aim is to prove that, with probability at least $\frac{2}{3}$, a random induced subgraph of size $s(\varepsilon)$ of G does not have property Π . The main idea, originating in the proof of [8, Corollary 7.2], is that T cannot distinguish the case in which it is given oracle access to G from the case it is given oracle access to a random induced subgraph of size $s(\varepsilon)$ of G . It follows that the latter random subgraph (of a graph being far from having property Π) is unlikely to have property Π , or else a contradiction is reached. Details follow.

Let S be a random set of $s(\varepsilon)$ vertices in G , and let G_S denote the corresponding induced subgraph. Consider what happens when T is invoked with oracle access to G_S : If G_S has property Π , then the tester T must accept with probability at least $\frac{3}{4}$. Thus,

$$\Pr_S[T^{G_S(\varepsilon)} = 1] \geq \frac{3}{4} \cdot \Pr_S[G_S \in \Pi]. \quad (25)$$

Now consider what happens when T is invoked with oracle access to G . On one hand, since G is ε -far from Π , the tester T must reject with probability at least $\frac{3}{4}$. On the other hand, all that T does is select a random S of size $s(\varepsilon)$ and inspect G_S . But inspecting G_S (as an induced subgraph of itself), for a random S of size $s(\varepsilon)$, is exactly what T does when given oracle access to G_S . Thus,

$$\Pr_S[T^{G_S(\varepsilon)} = 1] = \Pr[T^G(\varepsilon) = 1] \leq \frac{1}{4} \quad (26)$$

Combining Eqs. (25) and (26), it follows that $\Pr_S[G_S \in \Pi] \leq \frac{1}{3}$. The claim follows. \blacksquare

ACKNOWLEDGMENTS

We are grateful to several researcher that have contributed to our work and its presentation. In particular, we thank Yevgeniy Dodis for raising the question addressed in Section 3, Michael Krivelevich for making a comment that initiated the work presented in Section 4, and Noga Alon for insisting that a better transformation of adaptive testers to non-adaptive ones should be possible in the context of graph properties. We also thank Dana Ron for suggesting a simplification to the proof of Lemma 4.1, Amir Shpilka for suggesting a simpler proof of Claim 3.1 and an anonymous reviewer for suggesting the proof of Lemma 4.4 (which is simpler and stronger than our original proof). Finally, we wish to thank Noga Alon, Michael Krivelevich, Dana Ron, Salil Vadhan and several anonymous reviewers for various helpful comments.

¹⁸This assumption can be justified by first reducing the error probability of the original tester.

REFERENCES

- [1] N. Alon, On properties closed under taking induced subgraphs, private communication, 2001.
- [2] N. Alon, E. Fischer, M. Krivelevich, and M. Szegedy, Efficient testing of large graphs, *Combinatorica*, Vol. 20, pages 451–476, 2000.
- [3] N. Alon, O. Goldreich, J. Hastad, and R. Peralta, Simple constructions of almost k -wise independent random variables, *J Random Struct Alg* 3(3) (1992), 289–304.
- [4] N. Alon and M. Krivelevich, Testing k -colorability, *SIAM J Discrete Math* 15(2) (2002), 211–227.
- [5] N. Alon and J. H. Spencer, *The probabilistic method*, Wiley, New York, 1992.
- [6] Z. Bar-Yossef, R. Kumar, and D. Sivakumar, Sampling algorithms: Lower bounds and applications, *Proc 33rd STOC*, Hersonissos, Greece, 2001, pp. 266–275.
- [7] O. Goldreich, “Combinatorial property testing—a survey,” *Randomization methods in algorithm design* (Eds., P. Pardalos, S. Rajasekaran and J. Rolin), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, AMS, Providence, Vol. 43, 45–59. 1998.
- [8] O. Goldreich, S. Goldwasser, and D. Ron, Property testing and its connection to learning and approximation, *J ACM* Vol. 45, No. 4, 1998, 653–750.
- [9] O. Goldreich and L. Trevisan, Three theorems regarding testing graph properties, TR01-010, *ECCC*, 2001, available from <http://www.eccc.uni-trier.de/eccc/>.
- [10] J. Naor and M. Naor, Small-bias probability spaces: Efficient constructions and applications, *SIAM J Comput* 22 (1993), 838–856.
- [11] D. Ron, “Property testing (a tutorial),” *Handbook on Randomized Computing* (Eds., S. Rajasekaran, P. M. Pardalos, J. H. Reif, and J. Rolim), Vol. II, 597–649. Kluwer Academic, Dordrecht, 2001.
- [12] R. Rubinfeld and M. Sudan, Robust characterization of polynomials with applications to program testing, *SIAM J Comput* 25(2) (1996), 252–271.